White Paper

Cisco UCS and Dell PowerEdge M1000e: A Comparison

June 2016

What You Will Learn

Cisco Unified Computing System™ (Cisco UCS®) is more than a traditional blade solution. Unbound from traditional products and thinking, Cisco UCS is the first truly unified data center platform that combines industry-standard, x86-architecture servers with networking and storage access in a single system. This radically simplified solution is intelligent infrastructure that is automatically configured through integrated, model-based management to accelerate deployment of all your enterprise-class applications and services running in bare-metal, virtualized, and cloud-computing environments.

Traditional blade architectures–such as the Dell PowerEdge M1000e–originally evolved from the idea of repackaging rack servers and switching into a smaller form factor. As density and performance increased, so did complexity. As a result, these traditional blade solutions have failed to deliver the promised agility. Cisco revolutionized blade servers through the unification of networking and management fabrics and delivers on the original vision of true consolidation and agility of resources, time, and energy. This document shows you how.

Contents

Chassis Compromises 3

Cisco Unified Fabric 4

Dynamic Sharing 5

Deterministic Latency 6

Greater Availability and More Flexibility 8

Right-Size Networking Bandwidth for Less Cost 9

Revolutionary Management 9

Conclusion 11

For More Information 11

Cisco UCS and Dell PowerEdge M1000e: A Comparison

White Paper
June 2016

Imagine your IT department adapting easily to rapidly changing business needs. In contrast to vendors of traditional approaches, we deliver a state-of-the-art architecture that makes your dream a reality.

Chassis Compromises

The problem with traditional blades and with the PowerEdge M1000e architecture begins with the blade chassis itself. Rack-in-a-box blade architecture takes all the complexity of a rack (top-of-rack [ToR] switches, management modules, etc.) and squeezes it into every blade chassis. When you replicate this complexity multiple times per rack, you get even more complexity. In fact, you get far more complexity than a rack full of traditional servers with ToR switches. The PowerEdge M1000e architecture turns each chassis into an individual island that doesn't allow bandwidth to be shared outside each chassis. Changing applications or scaling resources requires you to reconsider the way that all components in the chassis are connected, leading to a cabling mess and networking errors that prolong deployments. Customers must determine when a new Dell switch module is necessary and what kinds of I/O adapters are used on each affected blade. Changing networking modules or updating blades to take advantage of new Intel architecture requires administrators to continually rethink and manage each infrastructure island. Dell's aging architecture limits flexibility and innovation in several ways:

- Multiple local management modules are necessary, and an external software overlay is required to present these individual components under a common interface.

- Multiple internal switching devices are necessary to connect each chassis to the outside world. This dependency, results in two switches for every 16 blades, or one switch for every 8 blades.

Support for the full range of Intel® Xeon® processors is limited. Even with some Intel Xeon processor E5-2600 family CPUs, Dell blade servers can use only 20 DIMMs instead of the full 24 because they have to add a large heat sink for the higher-wattage CPUs, and that occupies the space of 4 DIMMs. The end result is that you can't load the full number of DIMMs to support your applications (and may be forced to use more expensive DIMMs to meet your memory needs). The use of 10 DIMMs per CPU (instead of the 12 per CPU that you would normally have) results in an unbalanced configuration and has negative performance implications. A blade chassis struggling to keep up with the latest technology from Intel will likely continue to struggle to power and cool the system as Intel CPU generations continue to evolve. Dell does support the powerful Intel Xeon processor E7 v3 CPUs running up to 165 watts (W) in the Dell R930 rack server, but not in any of its blade servers, limiting performance and choice.

When we developed Cisco UCS, we took an approach that solved these problems. By simplifying the chassis with unified fabric and Cisco® fabric extender technology, we removed the complexity, heat, and multiple touchpoints, providing architecture ready for future technology generations.

Cisco UCS Management Streamlines Operations

Bluelock staff use automated profiles and provisioning in Cisco UCS Manager to quickly configure environments. "Not only can we get new clients started faster, but we can shift resources on the fly if a client experiences a sudden surge in usage."

– Pat O'Day, CTO, Bluelock

The Cisco UCS chassis:

- **Has no management modules:** Cisco UCS was designed from the beginning using a model-based unified management architecture that is fully redundant; has revolutionary Cisco UCS service profiles for immediate, secure configuration and reconfiguration; and adapts to changes in the environment, including the use of multiple blade generations and types.
- **Has no switching in the chassis:** Each chassis incorporates zero-management, low-power fabric extenders that simply forward all blade traffic to the system's fabric interconnects–and they do so faster and more efficiently than switches do.
- **Enables incremental scaling:** Cisco UCS eliminates the "seventeenth blade server" problem, in which lots of expensive infrastructure must be purchased to support one more blade server than a single Dell chassis holds. Instead of a 16-blade domain as in Dell's system, Cisco UCS allows hundreds of blades per domain.
- **Reduces network and management complexity:** Cisco UCS offers flexibility, scalability, and generational growth based on application needs.

The result is lower total cost of ownership (TCO) per server and greater business agility.

Cisco Unified Fabric

Cisco UCS is the only integrated system that reduces the number of hardware components and combines both blade and rack servers in a single unified fabric and management domain (Figure 1). Our approach eliminates management and networking devices in every chassis, reducing the cost of powering, cooling,

cabling, configuring, managing, monitoring, and maintaining the infrastructure. Cisco UCS places all management functions and configuration information in the fully redundant and highly available Cisco UCS Manager. Cisco unified fabric, with its wire-once capability, helps you scale data centers easily, quickly, and efficiently without requiring you to reevaluate your networking infrastructure every time you add new servers. With Cisco UCS, the network is established once, with no changes necessary as it scales to 160 servers per domain (or to multiple domains of up to 10,000 servers). Through aggregation of management and connectivity in the fabric interconnects, every server in the domain is automatically connected northbound to the LAN or SAN without time-consuming and risky reconfiguration at the chassis and server levels.

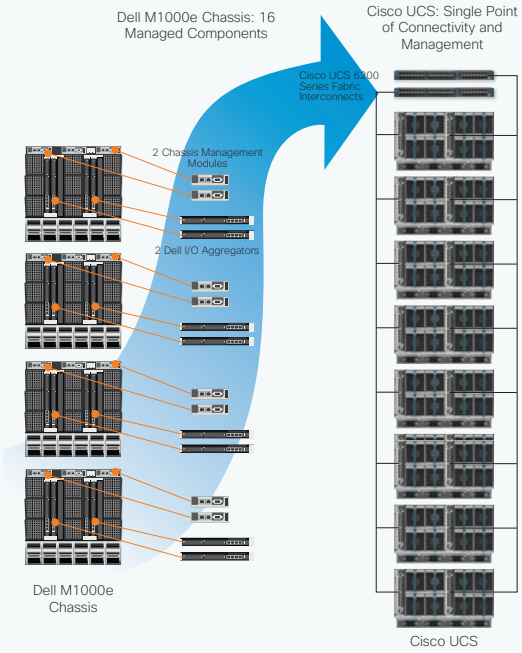

Figure 1 To Support 64 Blades, Dell PowerEdge Requires 16 Managed Components Compared to a Single Point of Connectivity and Management in Cisco UCS

Dynamic Sharing

PowerEdge M1000e architecture is more complex than Cisco UCS architecture and limits networking and scaling flexibility. Dell's minimum redundant configuration requires customers to purchase a pair of networking switches for every chassis. Without the capability for each chassis to share I/O bandwidth, customers are forced to overprovision and overpurchase hardware and port licenses every time they add another 16 blades. Customers must purchase, configure, maintain, power, and cool one switch for every 8 blades (with a minimum of two Dell modules per chassis).

Unlike Dell's network, in which bandwidth choices are made at deployment time and cannot be reconfigured without recabling the network, the Cisco UCS network is unified, and bandwidth is shared dynamically. This approach provides more effective

use of bandwidth and more headroom in the event that the application requires bandwidth from other modalities. Cisco UCS uses configurable quality-of-service (QoS) controls so you can assign and reassign bandwidth priorities, and it enables your applications to dictate connectivity and bandwidth.

For example, Figure 2 shows three types of traffic: virtual machine movement (green), SAN (yellow), and LAN (blue). Suppose the total available bandwidth is 10 Gbps, and you set a bandwidth guarantee of 3 Gbps for both the LAN and SAN and 2 Gbps for virtual machine movement traffic. If each modality offers 3 Gbps of actual traffic, the bandwidth allocation would look like time-step one (t1) in Figure 2, giving you a spare 1 Gbps. If the LAN has a burst of traffic to 4 Gbps as shown in time t2 and the bandwidth is available, LAN traffic will be allowed to use that bandwidth. Time t3 shows a burst of 6 Gbps of LAN traffic. Because the SAN is using only 2 Gbps of its allocated bandwidth, the LAN traffic is allowed to use 6 Gbps because it can "borrow" unused bandwidth from the SAN.

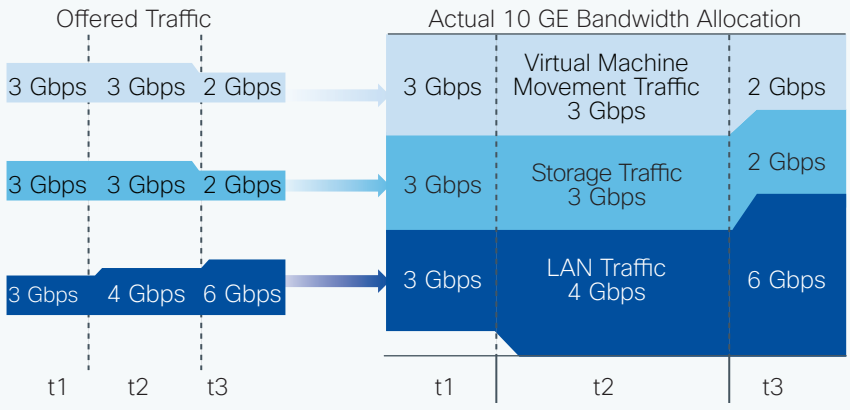

Figure 2 Cisco UCS Dynamic Traffic Allocation Example

Cisco's test results are summarized in the performance brief "Cisco UCS Outperforms Dell Blade Servers on East-West Latency."

Deterministic Latency

Cisco UCS unified fabric reduces latency and makes latency times more consistent. Cisco UCS fabric interconnects centrally manage network traffic within Cisco UCS as well as the traffic in and out of the system. For traffic that is moving from one chassis to another in the same Cisco UCS domain, there is no need to exit the system and send packets through another external switch (see path A in Figure 3). This central connectivity demonstrates one of the ways in which Cisco UCS functions as a single virtual chassis and provides the flexibility to place workloads anywhere in the system with assurance of consistent network performance.

This capability is important because it simplifies the placement of virtual machines in virtualized and cloud environments: regardless of location, virtual machines experience the same latency, removing the network as a placement constraint. When trying to reduce latency between servers, the best-case scenario for any vendor is a single network hop for data communicated between servers within

the same chassis. The worst-case scenario is three network hops required to move data between servers in different chassis. With PowerEdge M1000e, traffic between each chassis must first travel through the local Dell switch modules, through an external ToR switch, and back through the second chassis' Dell switch module, resulting in three network hops and greater latency. Although Dell supports pass-through modules, which can reduce the number of hops, this approach greatly increases the cabling (and recabling) requirements associated with the configuration. It also increases the number of ToR switch ports needed.

Actual east-west traffic tests of these identically configured systems show that:

- The PowerEdge M1000e with PowerEdge M I/O Aggregators has between 11 and 45 percent more latency than Cisco UCS. Cisco UCS demonstrated lower latency than the PowerEdge M1000e with PowerEdge M I/O Aggregators for every test case and every packet size (User Datagram Protocol [UDP], TCP, and TCP round-trip times).
- As packet sizes increased in each test, the PowerEdge M1000e disadvantage also increased compared to Cisco UCS.
- Performance is almost identical for both single-chassis tests and multichassis tests for Cisco UCS. With the Dell configuration, after traffic leaves the chassis, latency increases dramatically.

Cisco remains the leader in application-specific integrated circuit (ASIC) and network design optimization for the end-to-end network stack, enabling business applications and virtual environments to perform better. The new, third-generation Cisco UCS fabric interconnects allow Cisco UCS to have true 40-Gbps end-to-end bandwidth, which Dell does not have, and can provide 360 Gbps of total bandwidth to a chassis.

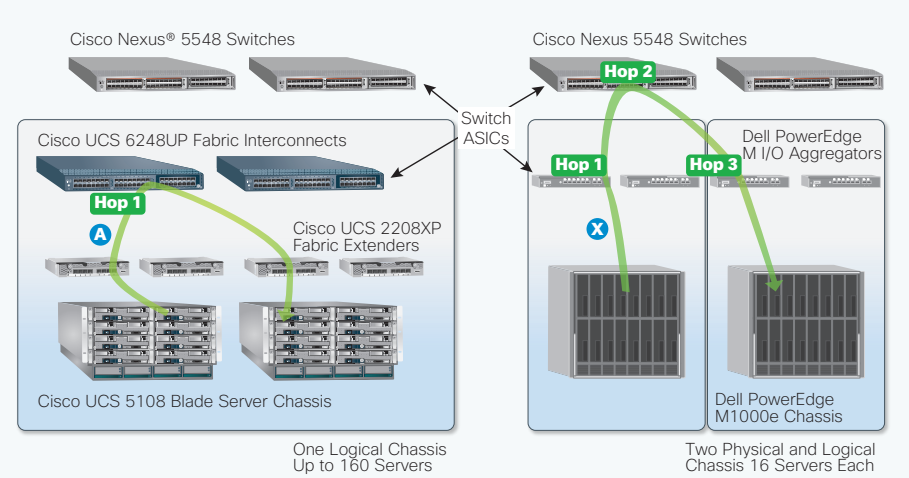

Figure 3 Cisco UCS and Dell PowerEdge M1000e Traffic Flow Between Blade Servers in Different Chassis

Greater Availability and More Flexibility

Cisco UCS not only provides an active-active data fabric out of the box, but it also provides a fully redundant unified management subsystem. As with other similar traditional designs, Dell blade server I/O connections are fixed and physically mapped in the midplane and cannot be modified in any way. For example, Dell network interface card (NIC) port 1 goes to interconnect bay A1, and Dell NIC port 2 goes to interconnect bay A2. Any changes in connectivity between a blade and the outside world forces customers to manually remove or add and recable the physical Dell switches and server I/O adapters. This recabling results in downtime and is prone to errors. In contrast, Cisco UCS is truly a wire-once technology that allows any changes, including port mapping and I/O assignment, to be fully programmatically implemented and user defined.

Using Cisco UCS virtual interface cards (VICs), Cisco UCS servers can be configured for any workload in minutes, without touching the server. Cisco UCS VICs present up to 256 PCI Express (PCIe) standards-compliant interfaces to the host that can be dynamically configured as either NICs or host bus adapters (HBAs). These adapters can be accessed by the operating system or hypervisor, and they also can be attached directly to virtual machines, accelerating performance. Cisco UCS VICs support hardware fabric failover so that if one fabric fails, the operating system never knows about it. You have the option to use OS-based NIC teaming or Cisco UCS hardware fabric failover, or both, depending on application requirements. You can configure this option on a per–virtual NIC (vNIC) basis.

As Figure 4 illustrates, NIC 1 is configured with a primary path to fabric A, with failover configured on fabric B. Similarly, NIC 2 is configured with a primary path to fabric B, with failover configured on fabric A.

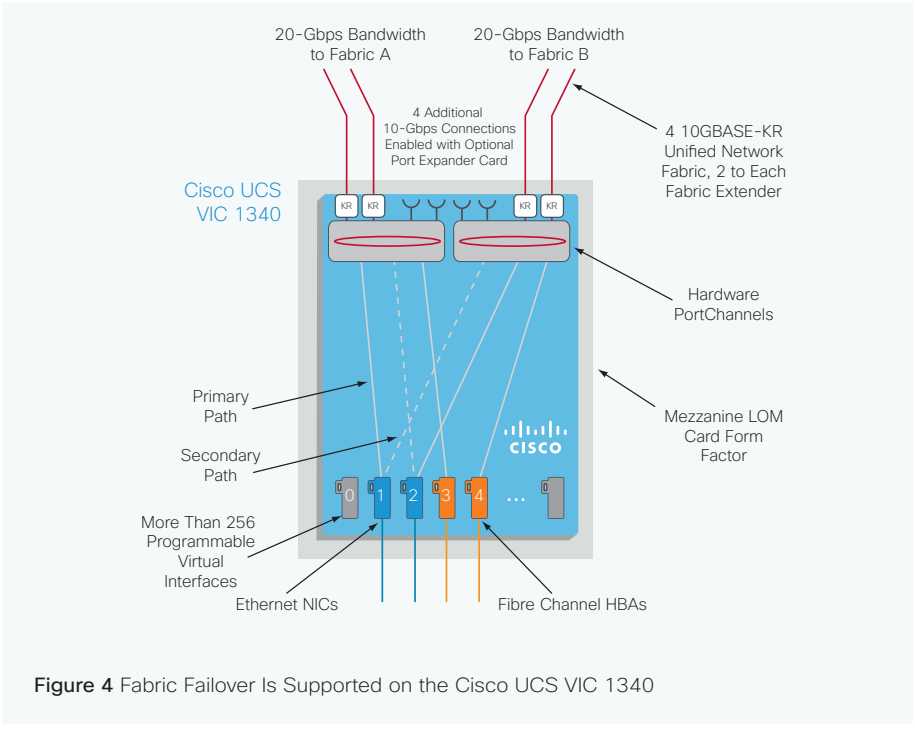

Figure 4 Fabric Failover Is Supported on the Cisco UCS VIC 1340

Right-Size Networking Bandwidth for Less Cost

Network bandwidth is crucial to getting information to the right place at the right time. Not only do companies need the flexibility to meet current application bandwidth requirements, but they must also be prepared for future bandwidth growth. Rather than a rigid, fixed-I/O topology that requires customers to add more in-chassis switches, Cisco unified fabric allows any server in the domain to access the total uplink bandwidth to accommodate traffic bursts. For example, you can double the bandwidth simply by increasing the number of uplink cables between the chassis and the fabric interconnects and enable them with no downtime or application impact. No planning, configuration, or cabling to the chassis switch is required, as would be the case with Dell switches. Cisco UCS can scale blade bandwidth as applications demand: up to 40 Gbps with the mezzanine LAN-on-motherboard (LOM)–format Cisco UCS VIC 1340, or up to 80 Gbps with the Cisco UCS VIC 1340 plus a port expander card. You are not forced to purchase and overprovision in-chassis switches regardless of application requirements. Because Cisco UCS architecture requires fewer components to scale your bandwidth for peak traffic flows, you don't pay for bandwidth or the associated components that you don't need.

Revolutionary Management

With Cisco UCS, servers, connectivity, and management are inseparable. The complete abstraction of configuration information creates an on-demand, zero-touch environment. Cisco UCS was designed from the beginning with embedded, all-inclusive, model-based management through Cisco UCS Manager. Cisco UCS is intelligent infrastructure that is self-aware and self-integrating. Every server connected to Cisco UCS, whether it is a blade server or a rack server, is automatically detected and placed in a resource pool and even automatically configured if you so desire. The system is built from the foundation so that every aspect of server identity, personality, and connectivity is abstracted and can be applied through software using a Cisco UCS service profile. With Cisco UCS, servers are configured automatically, eliminating the manual, time-consuming, error-prone assembly of components into systems. With Cisco VICs, even the number and type of I/O interfaces are programmed dynamically, making every server ready to power any workload at any time. Cisco service profiles can be quickly created from templates, enabling fast configuration of one or 100 blade and rack servers in just a few minutes.

Cisco UCS Manager is integrated, model-based management. With Cisco UCS Manager, administrators manipulate a model of a desired system configuration and associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation accelerates provisioning and workload migration, delivering accurate and rapid scalability. For the first time, you have an automated, policy-based mechanism for aligning server configuration with workload. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco UCS Manager can be accessed through a GUI, a command-line interface (CLI), or an open, standards-based XML API that is used by a large ecosystem of management tools. Cisco UCS Central Software, an extension of Cisco UCS

Manager that uses this XML API, enables management of multiple Cisco UCS deployments across geographic locations, with the first five Cisco UCS Central Software licenses free.

The PowerEdge M1000e traditional blade chassis design forces you to think of each chassis as an island, and this concept extends to management as well (Figure 5). You can manage each individual chassis with a combination of individual blade server integrated Dell Remote Access Controllers (iDRACs) and the Dell Chassis Management Controller (CMC) management board physically residing in each chassis. But to gain even limited template and policy capabilities, you must also purchase Dell Active System Manager (ASM).

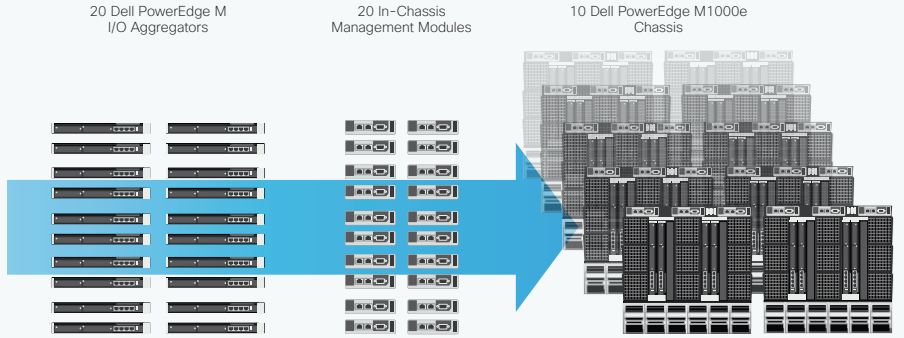

Figure 5 Dell PowerEdge M1000e Requires 20 In-Chassis Switches and 20 Management Boards for a Total of 40 Management Points to Support 160 Dell Blade Servers

ASM is a licensed, top-down software tool that sends scripted commands to management touchpoints in the chassis. Like other management tools layered on after product development, ASM struggles to adapt to changing configurations and server types. Unlike Cisco UCS, which was designed from the beginning with policy-based management at its core, ASM has a number of limitations:

- ASM requires you to perform many steps to deploy, license, and implement the initial configuration.
- You can define only seven BIOS settings, so you are tied to your hardware configuration.
- No QoS support is available for network traffic management.
- No role-based access control (RBAC) is available for enhanced security or multitenancy. ASM provides only three default roles.
- ASM is not redundant, so no high-availability capability is available as it is with Cisco UCS Manager.
- All managed resources must be manually configured for discovery through the IP address, which requires an endpoint set up for IP addresses and local authentication.
- Server profile mobility is constrained:
  - Dell recommends that you migrate between identically configured hardware.

  - All I/O devices must be in the same slots and must be the same models.
  - Different operating systems may not boot correctly on different hardware.
- ASM does not fully support hardware state abstraction. It uses scripts and specifically defined parameters to create and move profiles between servers.

Conclusion

Cisco reinvented the server market five years ago when it delivered Cisco UCS. It surpassed the capabilities of Dell, HP, and IBM at the time. Although Dell has attempted to incorporate some of the innovations that Cisco has brought to market, these attempts have been haphazard and have forced Dell into many compromises that Cisco did not have to make. Although Dell has attempted to improve blade server management over the years, its offerings require complex licensing and fail to eliminate the limitations of the underlying traditional architecture. The result is that even Dell's most recent offering is still far behind Cisco UCS.

Cisco has become the number-one x86-architecture blade vendor in the Americas and number two worldwide. According to IDC, Cisco UCS forms the basis of the some of the top integrated infrastructure solutions: Virtual Computing Environment (VCE) coalition Vblock™ System and FlexPod. This achievement is the result of the innovation used to design, build, and integrate the Cisco UCS hardware and management software. This integration radically simplifies the infrastructure, helping reduce both capital and operating costs, and delivers a solution with excellent agility, flexibility, scalability, resiliency, and manageability and greatly reduced TCO.

For More Information

- For more information about Cisco UCS, please visit http://www.cisco.com/go/ucs.
- For more information about Cisco UCS award-winning performance, please visit http://www.cisco.com/go/ucsatwork.
- For more information about the Cisco UCS VIC 1340, please visit http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.html.
- For more information about Cisco UCS Manager, please visit http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html.
- For more information about Cisco UCS Central Software, please visit http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-central-software/index.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**.

LE-43506-01 07/16