

# Distributed Services on the Cisco 7500 Series Router

## Executive Summary

The Cisco 7500 Series router enables the network to support today's large-scale e-business applications. With an industry-leading combination of network services, interface density and packet-forwarding performance, the Cisco 7500 Series is the optimum choice for a large-scale WAN services router.

With the Cisco 7500 Series, hundreds of sites can be aggregated in a single router while providing the richest set of network services including security, quality of service (QoS), and multicast. The Cisco 7500 Series distributes switching and services across multiple Versatile Interface Processors (VIP), which allows the network to scale performance linearly by adding additional VIP cards.

## Content Networking

Cisco Content Networking delivers the network agility required by an enterprise to deploy new Internet business applications critical to securing a competitive advantage by increasing revenue while reducing operating costs. By creating the end-to-end intelligent network services required for Internet business applications such as e-commerce, supply chain management, and workforce optimization, Cisco Content Networking integrates the enterprise with customers, suppliers, and business partners. Cisco Content Networking is an intelligent network architecture that dynamically recognizes Internet business applications and engages network services to achieve end-to-end security, performance, and availability. This architecture has the following three components:

- Intelligent network classification and network services delivered through Cisco IOS<sup>®</sup> software
- Intelligent network devices that integrate Internet business applications with network services
- Intelligent policy management framework for configuration, monitoring, and accounting

With these three components, the dynamic architecture of Cisco Content Networking delivers the intelligent network services required to drive the next-generation Internet business model.

## Why Scalability Matters

Agile enterprises are deploying a dizzying array of new applications that support a wide variety of business processes, including workforce optimization, supply chain management, e-learning, e-commerce, and customer care. These applications are being deployed across large enterprises with hundreds and even thousands of sites. To be successful, the network services must scale with this growth.

The network, an integral component of any e-business deployment, must provide both basic connectivity between the users and servers, and a wide variety of services. With Cisco Content Networking, the network can recognize the content of data on the network and apply the appropriate service, such as QoS, security, load balancing, or caching.

Public

Copyright © 2000 Cisco Systems, Inc. All Rights Reserved.

Page 1 of 8

Growth of new e-business applications and the increasing use of intranets and extranets drive the demand for an increase in available wide area network (WAN) bandwidth. The enterprise WAN must better use existing WAN resources and cost-effectively scale available WAN bandwidth. The integration of mission-critical applications with numerous other applications across a common WAN also requires multiple user or application service levels. This ensures that each application gets the right mix of throughput and delay over finite WAN resources. In addition, most enterprise WANs mix multiple interconnect strategies such as ATM, Frame Relay, IP access, virtual private networks (VPN) and dedicated leased lines to take advantage of cost, services rendered, and performance gains in WAN technology.

For example, in an e-learning video-on-demand application, the video feed must be delivered within certain bounds to ensure a high-quality picture. The Cisco network is able to recognize the video feed and provide minimum guaranteed bandwidth to the feed using the QoS and multicast features inherent in Cisco IOS software. Imagine this scenario being played out across thousands of sites in a large enterprise. The central-site router located near the video server must be able to recognize the video stream and apply QoS services for each end-user. In addition, some users may be located across a high-speed leased line, others over Frame Relay, and others over a dialup remote access VPN. This can place a huge load on the router, and the load will only get higher as other new e-business applications are introduced.

Traditionally one would have solved this problem by segmenting the load and installing multiple routers at the central site. Each router would carry a portion of the load as determined by the network administrator at the time the router was installed. This system will work fine in many instances, however it requires redundant networking equipment at additional capital cost; is harder to manage, increasing management cost; and requires increased space to house the equipment.

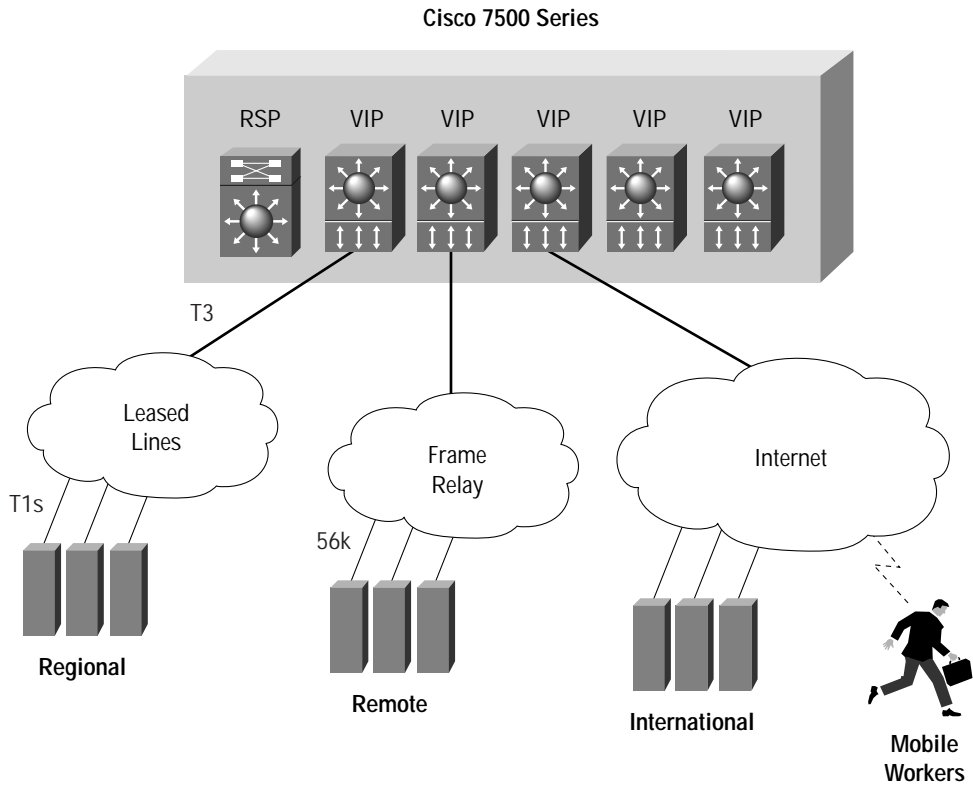
#### **Introduction to the Cisco 7500 Series**

The Cisco 7500 Series provides a superior solution to the problem of scaling the services necessary to support the growth in e-business applications. It is unique in its support of a wide range of services, with the required high port density necessary for a scalable enterprise WAN. The multiprotocol/multiple-media network capabilities of the Cisco 7500 Series enable integration of multiple WAN technologies and provide the necessary bandwidth and service requirements unique to each remote site. The Cisco 7500 Series integrates the richest set of Layer 3 services available including QoS, security, encryption, and traffic management. Advanced multichannel technology provides for a scalable and cost-effective WAN solution where high-density, leased-line services are required.

The Cisco 7500 Series distributes services across multiple VIPs to provide the scalability necessary in today's fast-paced Internet environment. The VIP is based on a RISC engine optimized for input/output (I/O) functions. One or two port adapters or daughter boards are attached to this engine. These port adapters provide the media-specific interfaces to the network.

A distributed architecture ensures high packet throughput as network speeds and route complexities increase by moving the packet-forwarding decisions from the central routing engine out to the network line cards. VIPs provide this service by offloading all packet-forwarding decisions from the central processor. This scenario allows a Cisco 7513, for example, to have 11 packet-forwarding engines, ensuring router performance in high-bandwidth applications. This is illustrated in Figure 1, which shows the Cisco 7500 Series router connecting multiple remote sites over a diverse set of WAN technologies.


Figure 1 Building a Scalable Network Using the Cisco 7500 Series Router



Cisco IOS software services are distributed to each VIP in a Cisco 7500 Series router. This enables the deployment of large-scale e-business applications by spreading the service load across multiple independent processors. As enterprise networks migrate to supporting mixed packet, video and voice services with mission-critical applications, it becomes necessary to maximize the use of network resources through advanced services. The Cisco 7500 Series utilizes many technologies to accelerate data movement in VIPs and to scale the application of network services. The Cisco 7500 Series takes advantage of the distributed Cisco IOS features within each VIP to accelerate a variety of network services such as QoS, security, data compression, and accounting.

For example, each VIP can intelligently classify traffic flows at T3 rates (90 Mbps). Using two VIPs increases this capability to two T3s, three VIPs, to three T3s, and so on. Each additional VIP linearly increases the router’s ability to provide services without adversely affecting the other VIPs already running in the system. This is illustrated in Figure 2.

Figure 2 Scaling Services with Additional VIPs



The ability to remove the main route processor from all forwarding and services activity, even in the presence of extensive network topology changes, frees more of the processor's resources for route-management (control) tasks. Even a distributed architecture can experience a significant performance loss without a complete separation of all data-handling activities from the control plane. The Cisco 7500 Series provides high performance of the distributed architecture through new software technologies, which scale both the control and forwarding processes in the router. For example, Cisco Express Forwarding (CEF) proactively updates routing information within VIPs to ensure that all packet forwarding is localized to the interface card.

## Distributed Services

The Cisco 7500 Series provides a wide variety of services that can be distributed across multiple VIP processing engines. Table 1 provides a summary of the distributed services.

Table 1 Distributed Services

Service	Features
Basic Switching	<ul style="list-style-type: none"> <li>• Cisco Express Forwarding (CEF)</li> <li>• IP fragmentation</li> <li>• Fast EtherChannel</li> </ul>
VPN	<ul style="list-style-type: none"> <li>• ACLs—extended and turbo</li> <li>• Cisco encryption</li> <li>• Generic route encapsulation (GRE) tunnels</li> <li>• IP Security (IPSec)</li> <li>• Layer 2 Tunneling Protocol tunnels (L2TP)</li> </ul>
QoS	<ul style="list-style-type: none"> <li>• NBAR</li> <li>• Traffic shaping (dTS)</li> <li>• Policing (CAR)</li> <li>• Congestion avoidance (dWRED)</li> <li>• Guaranteed minimum bandwidth (dCBWFQ)</li> <li>• Policy propagation via BGPh</li> <li>• Policy routing</li> </ul>
Multiservice	<ul style="list-style-type: none"> <li>• Low latency queuing</li> <li>• FRF 11/12</li> <li>• RTP header compression</li> <li>• Multilink PPP with link fragmentation and interleaving</li> </ul>
Accounting	<ul style="list-style-type: none"> <li>• Output accounting</li> <li>• NetFlow export</li> <li>• Precedence and MAC accounting</li> </ul>
Load Balancing	<ul style="list-style-type: none"> <li>• CEF load balancing</li> <li>• Multi-link PPP</li> </ul>
Caching	<ul style="list-style-type: none"> <li>• WCCP V1</li> <li>• WCCP V2</li> </ul>
Compression	<ul style="list-style-type: none"> <li>• L2 SW and HW compression</li> <li>• L3 SW and HW compression</li> </ul>
Multicast	<ul style="list-style-type: none"> <li>• Multicast distributed switching</li> </ul>

In the following section, we give some examples of how these services can be used to provide a complete solution for various e-business applications.

## Connectivity Services

When building an intranet or extranet, connectivity is a fundamental requirement upon which all the other services are built. The Cisco 7500 Series router provides high packet-forwarding performance by distributing the switching operation to the VIP. Distributed CEF ensures that the VIP is always able to provide the maximum packet throughput.

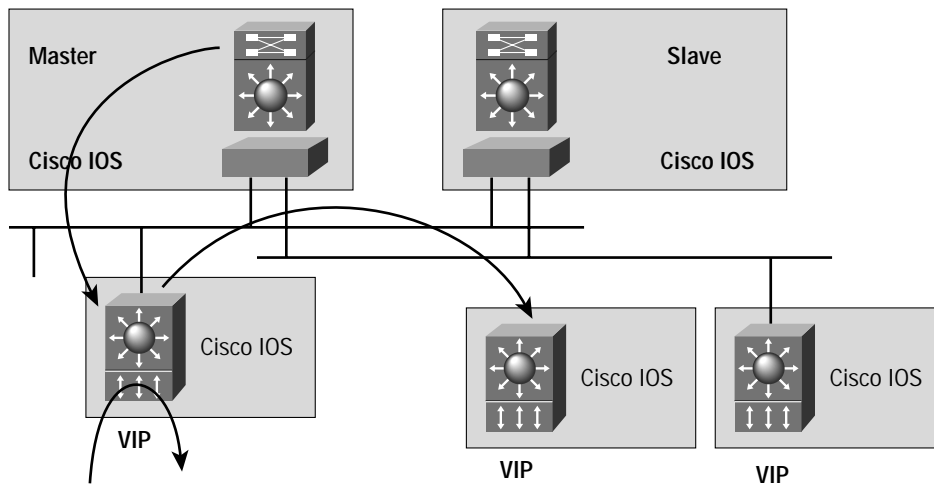
CEF is a non-cache-based switching mode for IP packets. Historically, all switching modes (except process switching) on Cisco routers have been cache-based. In cache-based switching, the first packet of a flow is sent up to process level, where its destination address is compared with the routing table to obtain forwarding information. Then a route cache entry with the corresponding forwarding information is built so those subsequent packets of the same flow can be fast-switched based on the route cache.

This cache-based scheme is more than adequate for most networks. However, the high traffic volume and dynamic traffic mix found in today's Internet backbone can sometimes reach the size limit of a route cache. In this situation, the router spends more CPU cycles than desired in flushing and rebuilding route cache entries.

With CEF, instead of building route cache entries on demand, a forwarding information base (FIB) is built. The FIB is based on the entire routing table, and then is downloaded to all the VIPs for distributed switching. Both the RSP and the VIPs use the same FIB, as CEF runs in both places. The FIB is updated only when the routing tables change, thus completely eliminating the per-flow overhead of route-cache maintenance. In addition, there is no issue of cache flushing and rebuilding, so performance is not affected by traffic mix.

This is illustrated in Figure 3 where the switching decisions are handled locally by the VIP whenever possible.

Figure 3 Distributed CEF Switching



## Accounting

When you rely on the network as a mission-critical resource, monitoring the amount and type of traffic is imperative. NetFlow is an ideal solution to this need, providing customized summaries of traffic statistics that are invaluable to network management, planning, and growth. NetFlow analyzes traffic flows, where a flow is defined by such variables as IP address, protocol type, and autonomous system (AS) number. It provides a highly granular analysis, including flow duration, packet size, and many other elements.

For large-scale networks where the core has a heavy traffic load, it is more appropriate to run NetFlow on the edge of the network so that core routers do not have to spend CPU cycles exporting lots of NetFlow data. By distributing this service on each VIP, the Cisco 7500 Series router is able to easily provide NetFlow accounting data for even the largest networks.

## Quality of Service (QoS)

QoS is a fundamental service required for next-generation networks. As enterprises shift mission-critical applications to the network and create a unified multi-service architecture for voice, video, and data, the ability to manage delivery terms becomes increasingly critical. Enterprises need QoS to deploy on-demand bandwidth-intensive applications such as videoconferencing, as well as time-sensitive information-delivery applications such as stock transactions. QoS reduces WAN costs through more efficient use of network links. With QoS

in place, network administrators worry less about healthy levels of oversubscription and concern themselves instead with managing the traffic mix to ensure that time-sensitive and mission-critical traffic is not delayed. QoS allows organizations to converge their voice, video, and data traffic by ensuring that each traffic type receives the network characteristics it requires. For example, voice requires relatively little bandwidth but low latency. File transfer applications, on the other hand, require high bandwidth, but have few latency requirements.

The Cisco 7500 Series has a rich set of Layer 3 QoS services that are distributed to the VIP. These enable the deployment of large-scale multiservice networks while ensuring that each traffic type receives the network characteristics it requires.

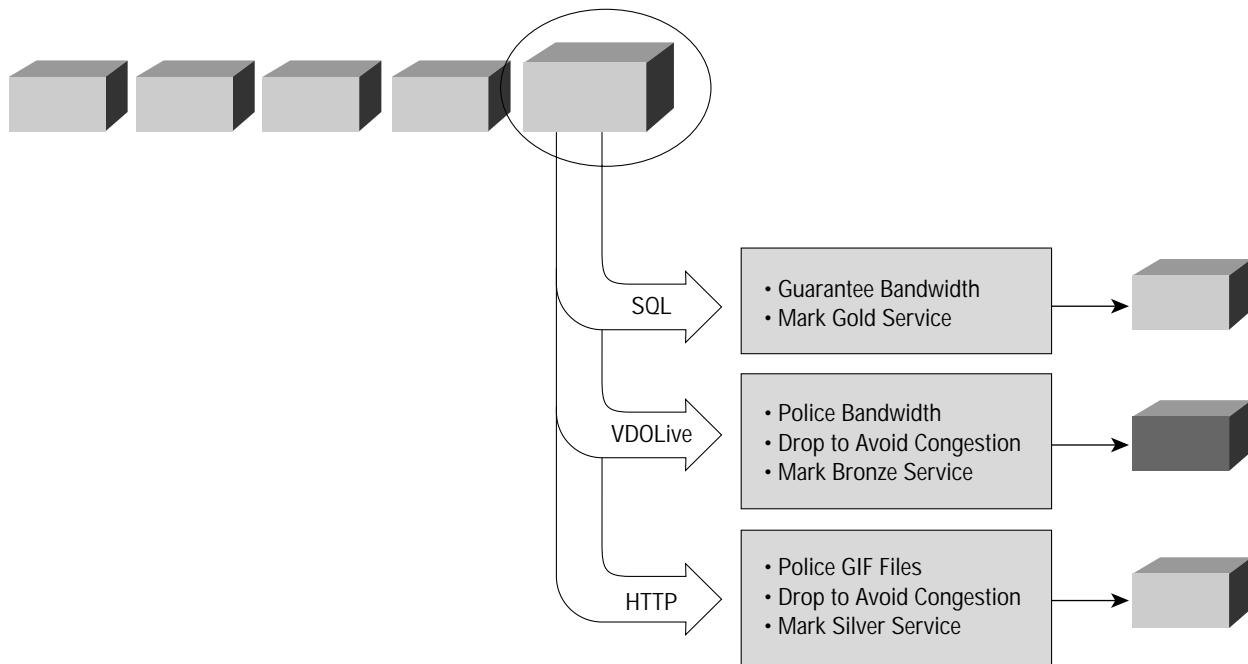
Using distributed access control lists and distributed NBAR, the router can intelligently recognize the type of data passing through the network. NBAR is a new classification engine that can recognize a wide variety of applications, including Web-based applications and client/server applications that dynamically assign TCP or UDP port numbers. Once the application is recognized, the network can invoke specific QoS services for that particular application.

The QoS services on the VIP include:

- Guaranteed bandwidth with class-based weighted fair queuing (CBWFQ). This is a mechanism to provide guaranteed bandwidth to particular traffic classes, such as voice or Enterprise Resource Planning (ERP) applications, while still fairly serving all other traffic in the network. The bandwidth assigned to a class is the minimum bandwidth delivered to the class during congestion.
- Marking for differentiated service downstream using the type of service (ToS) bits or Diff Serv code points (DSCP). Packet marking is usually done at the edges of the network so the proper policies may be applied later in the routing path. The older ToS byte markings provide three bits in the IP header, of which six combinations are usable. The newer DSCP mechanism utilizes six bits in the IP header, providing 64 separate markings. The bits used by both mechanisms overlap, so it's best to utilize either ToS or DSCP values in the network, but not both.
- Enforce bandwidth limits via policing or shaping. The policing and shaping feature allows one to limit the bandwidth available to a particular application. For example, one can limit FTP traffic to at most one mbps on a DS3 WAN link. Any traffic over the policed rate will be dropped or the ToS or DSCP values may be changed to a lower class of service. The difference between policing and shaping is that policing simply drops excess traffic while shaping buffers excess packets and plays them out at the specified rate. Also, traffic shaping can respond to congestion information provided by a Frame Relay network. The shaping function is called distributed traffic shaping (dTS) while the policing function is known as distributed committed access rate (dCAR).
- Congestion avoidance policy (distributed weighted random early detection [dWRED]). Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Without random early detection, when the output queue is full packets are dropped until the congestion is eliminated. Unfortunately, global traffic synchronization may occur as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. This occurs when multiple TCP hosts simultaneously reduce their transmission rates in response to packet dropping, and then increase their transmission rates once again when the congestion is reduced. RED takes advantage of TCP's congestion control mechanism by randomly dropping packets before the link is completely congested. This avoids the problem of dropping packets from all the hosts at one time while still signaling that they should temporarily slow down their transmission of packets. Weighted RED drops packets from lower priority flows first, ensuring that high priority flows do not get stuck behind congestion caused by less important data.

For example, in provisioning an e-business application, configure the router to classify traffic and then apply a set of these services to each traffic type. For example, your sales representatives require a fast response when they query for an order status in the corporate data warehouse hosted on an Oracle database server. Unfortunately, if others on the network are using high bandwidth applications, such as VDOLive or viewing large GIF files, then the SQL\*NET transaction to the Oracle database may be delayed. NBAR and the distributed QoS services address this problem by properly classifying the applications and then providing guaranteed bandwidth to the SQL\*NET queries while simultaneously policing the other applications. Figure 4 illustrates this solution.

Figure 4 Providing a Distributed QoS Solution



The power of the distributed architecture becomes apparent when you consider how this solution would be deployed to hundreds of sales offices connected over leased lines. With a Cisco 7513 router, one can deploy two multichannel T3 port adapters in each VIP. Each of these cards can support a mix of T1 and DS0 connections for 512 connections per VIP. This scenario can be replicated across each VIP installed in the router. When packets traverse the router, all of the QoS classification and services are handled at the local VIP. This enables up to 11 processors to be simultaneously servicing traffic, thus providing a highly scalable solution.

#### **Multicast**

Bandwidth intensive multi-point applications can be deployed today using the Cisco end-to-end IP multicast solutions in a scalable and efficient manner. With multicast distributed switching (MDS), multicast traffic that was switched at the Route Switch Processor (RSP) can now be distributed switched on the VIP. Switching multicast traffic at the RSP has disadvantages:

- The load on the RSP increased. This affected important route updates and calculations (for Border Gateway Protocol [BGP], among others) and could stall the router if the multicast load is significant.
- The net multicast performance is limited to what a single RSP can switch.

MDS solves these problems by performing distributed switching of multicast packets received at the VIP. MDS is accomplished by using a forwarding data structure called a multicast forwarding information base (MFIB), which is a subset of the routing table. A copy of MFIB runs on each line card and is always kept up to date with the central RSP's MFIB table.

#### **Using Distributed Services to Deploy E-Business Solutions**

One final example will help illustrate the strengths of the Cisco 7500 Series.

The Widget Company, a hypothetical example, has 400 remote sites of various sizes, from the home offices of sales reps up to large regional offices with hundreds of employees. The offices are scattered around the world, making it imperative that the Widget Company be able to deploy a diverse set of network technologies to cost effectively connect all of their sites. The Cisco 7500 Series router provides a complete solution to their WAN aggregation needs.

Using multichannel port adapters combined with the high-speed distributed switching capabilities of the Cisco 7500 Series, the larger remote sites can be connected over T1 lines aggregated into a T3 connection at the central site. Smaller regional offices can be connected over a Frame Relay network. International sites can be connected using a site-to-site VPN that leverages the ubiquity, flexibility, and low cost of the Internet.

On top of the physical configuration, the Widget Company can layer a rich set of services, including:

- Security, to encrypt the site-to-site VPN traffic and to protect the corporate site with the built-in Cisco Secure Integrated Software with firewalling capabilities.
- Quality of service, to protect mission-critical applications across the slower Frame Relay and VPN links while simultaneously allowing voice, video, and data to run across a common IP network.
- Multicast, to allow e-learning through video-on-demand services

The beauty of the Cisco 7500 Series is that the Widget Company's network can be built and run with the rich level of services it demands without sacrificing any performance. The distributed architecture allows the network manager to increase processing capability quickly and easily by adding additional VIP cards.

The Cisco 7500 Series enables the network to support today's large-scale e-business applications. With the industry-leading combination of network services, interface density and packet-forwarding performance, it is the optimum choice for a large-scale WAN services router.



**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems Europe  
11, Rue Camille Desmoulins  
92782 Issy Les Moulineaux  
Cedex 9  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 58 04 60 00  
Fax: 33 1 58 04 61 00

**Americas**

**Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Headquarters**

Nihon Cisco Systems K.K.  
Fuji Building, 9th Floor  
3-2-3 Marunouchi  
Chiyoda-ku, Tokyo 100  
Japan  
<http://www.cisco.com>  
Tel: 81 3 5219 6250  
Fax: 81 3 5219 6001

**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at <http://www.cisco.com/go/offices>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela