Cisco on Cisco Best Practices
Cisco IP Addressing Policy

TABLE OF CONTENTS

1 INTRODUCTION

This document outlines the IP address usage policy within Cisco's enterprise network. The aim of this document is to provide a set of guidelines for use by IT engineers in allocating and assigning addresses appropriately for infrastructure needs within Cisco.

This document does not cover detailed address planning strategies for any particular technology solution, although such strategies must be in line with this policy document. Such strategies should be part of an architecture and technology document for a specific solution, and as such are outside the scope of this policy document.

The key words "MUST," "MUST NOT," "REQUIRED," "SHALL," "SHALL NOT," "SHOULD," "SHOULD NOT," "RECOMMENDED," "MAY," and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

"Must" implies an absolute technical need where harm would occur if it were not followed. "Should" implies a standardization need where there is the possibility that exceptions are warranted. IASC and/or the Cisco Infrastructure Network Review Board (INRB) should be the only body to grant exceptions to a "should." All requests for exceptions should be accompanied by a well-documented business case justifying why an exception needs to be made.

This publication describes how Cisco deploys its own products. Many factors may have contributed to the results and benefits described. Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties; therefore, this disclaimer may not apply to you.

2 TYPES OF ADDRESSES

Public Addresses

The Internet Assigned Numbers Authority (IANA) allocates globally unique IP address blocks to various Regional Internet Registries (RIRs) globally. The RIRs then further allocate those address blocks to National Internet Registries (NIRs), Local Internet Registries (LIRs), or in some cases, directly to large organizations that apply for them. Cisco received a number of such allocations directly from the American Registry for Internet Numbers (ARIN), the RIR covering the Americas, as a multihomed organization.

In general, public IP addresses (those assigned by RIRs to Cisco) are appropriate anywhere in Cisco IT's infrastructure because of their global uniqueness. The flexibility of using public addressing is important, because it eliminates the need for special infrastructure required to translate IP addresses for access to the Internet. Public addressing should be used in any case where general-purpose access to the Internet from a particular network is planned or likely.

Demilitarized Zone

The demilitarized zone (DMZ) is the network or networks situated between an ISP edge router and Cisco's corporate firewalls. Public addresses must be allocated to all production networks in the DMZ. A public address must be assigned to all active interfaces on single-homed and multihomed hosts, except the loopback, to which a private address described in RFC 1918 and discussed below ("Private Addresses") may be assigned.

In DMZ labs, all devices that require connectivity to other devices outside the lab or on the Internet must be assigned a public address. Hosts that connect to other devices purely within the lab may be assigned private addresses. Using a Network Address Translation (NAT) gateway to translate a private address to a public address should not be used, because the limitations associated with tracking connections made through a translation device make proper security auditing difficult or impossible for such vulnerable areas of the network.

"Guest" networks that are tunneled through Cisco's internal network to the DMZ must also be assigned public addresses, because they are typically used by non-Cisco personnel who are virtually located (from a network topology perspective) outside of Cisco's internal network.

Data Centers

Public addresses should be assigned to hosts in Cisco's data centers. The NAT gateway infrastructure on the Cisco DMZ is not designed to provide a high-performance, highly available NAT service to the corporation. Assigning private addresses to data center hosts and relying on the NAT gateways to perform address translation compromises the reliability of Internet connectivity for mission-critical hosts that require it. In some cases, RFC 1918 space may be used to create private network-attached storage (NAS) networks between the server and the storage system.

Desktop Subnets

In a Cisco campus, building, or field office protected by badge readers, public addresses should be assigned to all desktop subnets for both wired and wireless LANs.

Network Infrastructure Connections (Router-to-Router Links)

Any router-to-router links connecting to areas of the network with public addressing should be addressed with public IP addresses. Routers serving specific areas of the network using and continuing to use only private addresses may use private addresses on the router-to-router links.

This requirement enables and helps ensure that path MTU discovery (RFC 1191) works properly; routers must be able to send "packet-too-big" errors and must be assured that the packets are likely to arrive at the original source host. If router-to-router links are addressed with RFC 1918 addresses, the Internet Control Message Protocol (ICMP) messages generated by the router will come from an RFC 1918 address. Networks filtering out incoming packets with RFC 1918 source IP addresses, or using unicast reverse path forwarding (uRPF), will likely drop these packets, breaking TCP for those applications. This will cause large packet transfer across a TCP connection to fail completely or perform suboptimally.

MPLS VPN CE-PE links are generally addressed with public space. In some cases, the service provider provides appropriate addressing. In other cases, Cisco supplies the addresses. Where possible, addressing should be provided where supported for ease of maintenance and troubleshooting.

Router-to-router links should be assigned a /30 subnet mask. The /31 subnet mask should be avoided in general, because support for it is not yet widely available for all software systems. However, the /31 subnet mask may be assigned to physical serial interfaces participating in a Multilink Point-to-Point Protocol (MLPPP) bundle for availability monitoring purposes.

Private Addresses

As defined in RFC 1918, IANA has reserved a number of network ranges that are marked as private and should never be used for routing in the public Internet. Theses network ranges are reserved for enterprises that want to build an internal network infrastructure based on TCP/IP without the need to access the Internet directly.

The reserved network ranges are as follows:

10.0.0.0 – 10.255.255.255 (10/8 prefix)
172.16.0.0 – 172.31.255.255 (172.16/12 prefix)
192.168.0.0 – 192.168.255.255 (192.168/16 prefix)

Universally recognizing these ranges as private and unroutable in the Internet means each organization can use these ranges internally without causing a conflict with public Internet addresses. If any organization attempts to route these networks externally towards the Internet, traffic will most likely be filtered and dropped by the ISP.

At Cisco, these network ranges are all in active use and are routable throughout Cisco’s enterprise network. These private ranges are to be assigned to networks that, during normal operation, do not require connectivity to the Internet. Note, however, that one particular large block at Cisco is reserved for non-routable labs and should not be assigned to any active networks that are intended to be routed throughout Cisco.

In general, RFC 1918 addresses should be used sparingly, under special circumstances where a considerable amount of address space usage in a portion of the network is likely, and general-purpose access to the Internet is not required. Without special circumstances to warrant private address deployment, Cisco IT uses public IP addresses for most infrastructure in the corporate network.

A downside to private addressing is that address space collision is far more likely to occur for two networks that are joined, such as IT infrastructure from an acquired company. Most companies are encouraged to use such private addressing for internal use instead of randomly choosing registered IP address space. In many cases, temporary NAT functionality must be put in place, or infrastructure must be renumbered immediately for the two networks to be interconnected.

Should Internet connectivity be required from a network addressed with RFC 1918 space, NAT capabilities or a proxy-based mechanism must be used. Cisco IT has implemented a basic NAT capability at most Internet access points of presence (POPs). NAT and other proxy-based mechanisms SHOULD be avoided, and public addressing SHOULD be used where possible for networks that require general-purpose access to the Internet, for a number of reasons:

- Some applications exchange IP addresses or dynamic TCP or UDP port numbers in the protocol payload. NATs or proxies must understand these applications and protocols, and must apply special handling to them to provide support for their operation over a NAT infrastructure. Although most NAT and proxy software gains support for common protocols relatively quickly (within a few software releases), it represents another dependency for the deployment of new applications.
- Multiplexing many connections into the same IP address through Port Address Translation (PAT) can result in strange behavior for applications that associate per-host state with a host's IP address.
- To provide an audit trail capability for reports of abuse or tracking a particular connection, the translations made by a NAT device must be exported and stored for future reference.
- To offer a redundant NAT capability, state must be shared among NAT devices and a failover-aware load-balancing mechanism must be put in place, adding greater complexity.
- There is a denial-of-service risk (whether malicious or unintended) in that the resources required to maintain the NAT state can be exhausted, causing existing connections to terminate or the NAT service to fail completely.

Private addressing may be considered for portions of the network using large amounts of address space where Internet access would be limited to a small subset of protocols. For example, lab networks rarely need to communicate using anything but basic HTTP and FTP protocols outside of Cisco.

Remote Access

Remote access at Cisco was initially established with public addressing, but the growth of telecommuting at Cisco has required the deployment of private addressing internally. The result is a combination of both public and private addresses within the Cisco remote-access architecture: public addresses at locations touching the Internet and private addresses within the Cisco remote-access service.

Public addresses. Most remote-access services are contracted individually by Cisco employees, who find the best possible service for their area. Their ISP gives them one or more public addresses for their home IP equipment. Within the Cisco network, Cisco has deployed fully redundant Internet access points consisting of multiple ISPs that require the use of independent or public IP addressing so that the VPN headend is always available via all the ISPs at that site. Cisco IT thus assigns a public address to the VPN concentrator within Cisco.

Private addresses. Within the IP Security (IPsec) tunnel built between the two endpoints, Cisco uses private addressing (see Figure 1).

Figure 1. Cisco Employees Use Public Addresses Within a Privately Addressed IPsec Tunnel

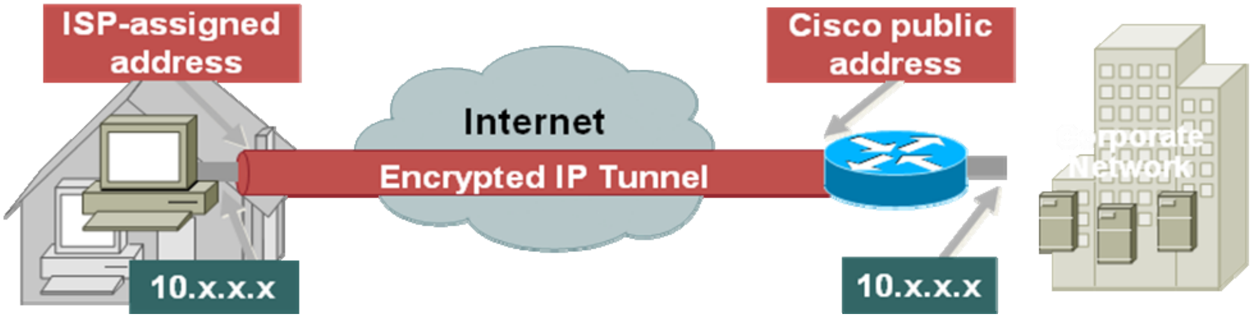

Private addressing was first assigned to remote-access users several years ago with the significant growth of home ISDN subnets. In most cases, home users were given a /30 prefix, consuming four IP addresses, to connect one host to the Cisco network. In other cases, because home broadband access to the Internet was not widespread, employees were permitted to place their home systems on the ISDN connection for fast access to the Internet and consumed even more address space. More recently, however, the RIRs had placed considerably greater restrictions on IP address space allocation. Since then, the landscape of remote access at Cisco has changed, allowing for most remote-access users to use only one IP address with software VPN.

Service provider-managed remote-access user subnets, and Cisco's in-house access server user subnets, should be assigned public IP addresses. However, all softwareVPN pools and Cisco Virtual Office spoke router subnets are assigned from RFC 1918 address space. With the limited public IP address space, and the growth of devices, assigning public address space is not a realistic option. For example, at the Cisco San Jose campus, Cisco IT uses a shared /24 pool of routable address space to service more than 12,000 concurrent connections during normal business days. For the same reason that DMZ labs, and in fact most labs, are recommended to use public IP addresses, remote-access users need to have their IP address tracked. In the event that a host is compromised or some suspicious activity must be investigated, it may be difficult to track connections made through a NAT gateway (especially historically).

Today many Cisco employees telecommute from home offices, using Cisco Virtual Office (a hardware router-based VPN solution). In future deployments of Cisco Virtual Office for these telecommuters, Cisco IT may implement split tunneling, but at present the added security risks outweigh the minimal gain. In small branch offices where Cisco will be deploying Cisco Virtual Office to provide WAN access over a VPN connection to the Internet, split tunneling is under consideration but not yet deployed.

Remote-access users with split tunneling enabled SHOULD use private addressing, because users with split tunneling enabled will not use the Cisco network to access the Internet. The user's VPN device will provide NAT translation services for access to the Internet.

If private addresses are used on the LAN interfaces of VPN headend routers that supply VPN users with private addresses, those private addresses on the LAN interfaces must be directed through the appropriate NAT infrastructure to enable path MTU discovery to work. The NAT infrastructure must translate the contents of the ICMP error message so that a remote server may recognize that a packet was too big to send through the VPN tunnel. Still, path MTU discovery is highly dependent upon the correct configuration of routers in the Internet to pass the correct MTU settings. Due to the unreliability of this feature, the software VPN client sets the MTU on its VPN interface to 1300 bytes, and the Cisco Virtual Office routers are similarly configured.

IP Telephony

Private addresses should be assigned to IP phones, because they do not require Internet connectivity.

Because of the future potential of integration with service provider IP networks, or access via the Internet, other devices supporting IP telephony such as Cisco Unified Communications Manager servers, H.323 gatekeepers, media gateways, media proxies, and voicemail bridges should receive public addressing. The relatively small number of addresses used by these devices will not significantly impact IP addressing.

Out-of-Band Network

Out-of-band (OOB) networks are solely used internally within Cisco and are not required to reach the Internet. As such, they should be assigned private addresses wherever practicable.

Network Management Loopback Interfaces

Loopback addresses used for network management systems to monitor devices are not used for connecting to the Internet and, therefore, should be assigned private IP addresses.

Some loopback addresses, such as those required for the operation of a protocol with external networks (e.g., MSDP), may need to be assigned public addresses.

Labs (Internal)

Internal labs, those located inside Cisco's corporate firewalls, should be assigned private addresses. Depending on business requirements, traffic from internal labs may be translated by the corporate NAT gateways if Internet connectivity is required. Alternatively, one of the corporate proxy servers may be used for connecting to the Internet from such a lab.

Extranet

An extension of Cisco's internal network, the Cisco extranet is deployed physically on the premises of a partner company. When Cisco started the extranet (with manufacturing sites), there was no need for Internet access. Cisco IT thus assigned private addresses to sites on the extranet network. For the few occasional sites where Internet connectivity is required, NAT may be implemented subject to approval by Cisco's internal Information Security organization. In the future, Cisco IT plans to deploy the Ironport S series web security appliance for URL filtering, and to support the IP spoofing into Internet routable addresses for web traffic. There are currently only a few cases where non-web Internet access is needed for extranet partners.

Provider Aggregatable Addresses

A provider aggregatable address is assigned by an ISP and not directly allocated to Cisco by an Internet registry. Occasionally, Cisco IT receives an ISP-assigned address space for various reasons, such as lab demos, experimental research, or via an acquisition. In locations that have Internet access using the VPN-only DMZ design, provider-assigned addresses are used because of the small number of IP addresses required to support labs and VPN concentrators.

Multicast

An IP packet with a multicast group address has the destination of an arbitrary group of IP hosts that have joined the group and wish to receive the traffic sent to the group. The multicast group address range assigned by IANA is subdivided into smaller ranges for specific purposes summarized in the following table.

Description	Group Address Range
Multicast global range	224.0.0.0 - 238.255.255.255
Link-local multicast addresses are local in scope; routers will not forward them regardless of the TTL	224.0.0.0 - 224.0.0.255
Reserved range for network protocols or applications that is forwarded by routers	224.0.1.0 - 224.0.1.255
Reserved for Source Specific Multicast (SSM)	232.0.0.0 - 232.255.255.255
Administratively scoped multicast addresses for use in private multicast domain	239.0.0.0 - 239.255.255.255

The multicast address block 239.0.0.0/8 is used internally within Cisco for administratively scoped addressing purposes.

3 ADDRESS ALLOCATION

IPv4

As the available pool of public IPv4 addresses is limited, Cisco IT should play a part in allocating addresses sensibly and appropriately. Cisco IT should allocate address space in appropriately sized blocks to allow for a good balance of summarization capability and to avoid wasting IP address space. Address block sizes and subnet sizes should be fit to the number of networks or hosts expected, while accounting for a 12- to 18-month growth projection. For more on allocation and assignment guidelines required for public IP address space usage, see Section 2, “Public Addresses.”

Address blocks SHOULD be allocated to regions based upon Internet route points; because public address space is announced to the Internet in aggregated chunks, care should be taken to build a plan that accounts for future expansion while avoiding waste. Address blocks allocated to a particular Internet route point must be a minimum size of /20, the current minimum allocation size of RIRs; these blocks should be allocated with a minimum size of /19 wherever possible (the former minimum allocation size). In some cases, blocks allocated during the “classful era” must be announced as a whole network block of their original class size.

For example, xxx.xxx.xxx.xxx/16 must not be broken up into smaller blocks of address space for different Internet route points. It must be announced as a /16 from only one route point.

A natural hierarchy exists in which addressing should be allocated. Addresses should first be organized by geographical region (see Figure 2), then by Internet route point, and then by any other hierarchy reflective of network topology to enable proper summarization.

Likewise, private addresses should also be allocated to each region in multiples of /19 blocks as a minimum.

Address space does not have to be allocated in blocks of /19 or /20; where required, larger blocks should be allocated and should be reflected as a single block within the Cisco IT Address Management repository.

All addresses should be summarized at key aggregation points before being advertised to other regions. Key aggregation points are at the site, regional, and theater levels, forming a naturally summarized addressing hierarchy.

Figure 2. Regional IP Address Allocation Within Cisco

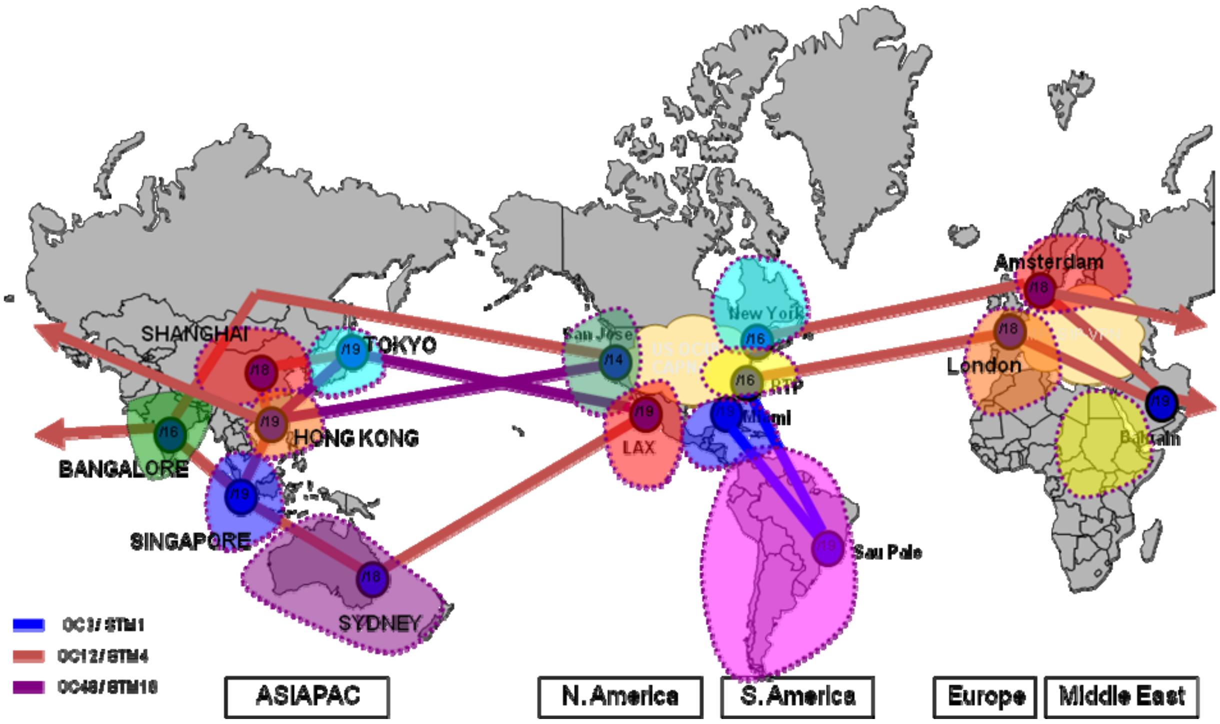

IPv6

In the same way that Cisco IT has deployed IPv4 addresses, there are guidelines that Cisco IT must follow when allocating and assigning IPv6 addresses (see Figure 3). Because IPv6 deployment is still in its infancy, Cisco IT expects to revise these guidelines as IPv6 deployment becomes more prevalent.

Figure 3. Current IPv6 Deployments Within Cisco

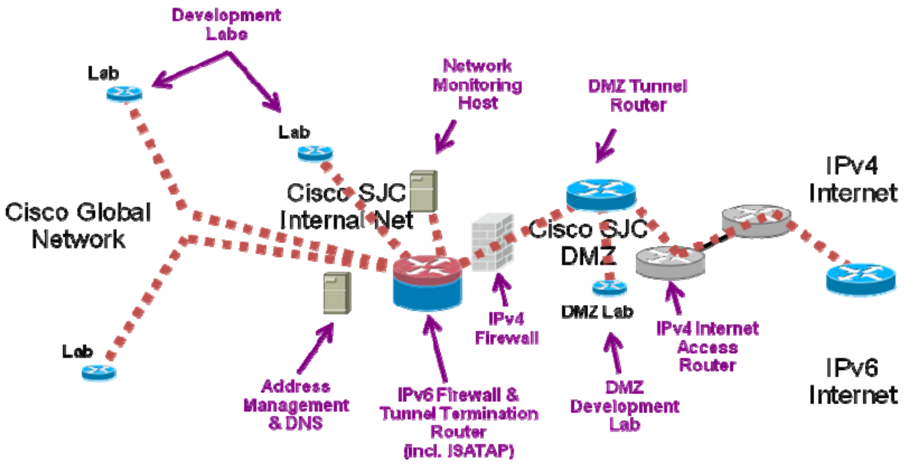

4 ADDRESS SPACE MANAGEMENT

The Cisco IT Address Management (AM) repository MUST be used by all Cisco IT engineers when allocating address blocks, assigning subnets, and assigning addresses to interfaces. Any address that is either configured or reserved to be configured on an active interface of a device connected to Cisco's production network MUST be reflected accordingly in AM.

Primary allocation of address space is made in AM by a security group called "CIDR Block Admins" who have the ultimate authority to effect address space changes in the AM repository. This group must be kept as small as possible to avoid accidental deletion of important portions of Cisco IT address information. The CIDR Block Admins group must not be granted to users who require only wide operational access to address space. The security user/group model available with AM is strong and flexible enough to support proper address space change authority for address space assigned to specific groups.

5 NON-CISCO IP ADDRESS SPACE IN THE CISCO IT ROUTING TABLE

Address space allocated or assigned to any entity other than Cisco or its acquired companies should not be introduced into the Cisco IT routing table. The primary reason for this policy is to avoid routing confusion for another corporation's Internet-reachable services and to avoid address space collisions in RFC 1918 space. For example, if an extranet partner's address blocks are injected into the Cisco routing table, the extranet design will need to account for all traffic normally routed via the Internet in its ACLs, including web servers, mail servers, and other Internet resources.

This policy primarily affects the extranet portion of the Cisco IT network. NAT is routinely used in these cases to translate specific traffic for networked transactions between Cisco and its partner organizations.

Exceptions to this policy must be approved by the INS Architecture Steering Committee (IASC).

6 ADDITIONAL IP SPACE ACQUISITION FROM REGIONAL INTERNET REGISTRIES

The CIDR Block Admins group is responsible for justifying the top-level use of Cisco's address space to ARIN (the RIR).

The CIDR Block Admins team, authorized by the Network and Data Center Architecture and Design review body, is the only team authorized to request address space for Cisco; all address space requests from any regional registry must be handled by the CIDR Block Admins team. Allocation of address space to an organization depends upon utilization of all space allocated to that organization, regardless of internal organization or controls.

Address space requests made of ARIN should be made in line with ARIN's procedures, including submitting proper IP address utilization reports and future IP address space allocation plans. The person representing Cisco in such requests should have experience in prior allocation requests from an RIR.

7 REFERENCES

LOCALADDR Hinden, R., Haberman, B., "Unique Local IPv6 Unicast Addresses," http://www.ietf.org/internet-drafts/draft-ietf-ipv6-unique-local-addr-04.txt

RFC 1191 Mogul, J., Deering, S., "Path MTU Discovery," November 1990

RFC 1918 Rekhter, Y., et al., "Address Allocation for Private Internets," February 1996

RFC 2050 Hubbard, K., et al., "Internet Registry IP Allocation Guidelines," November 1996

RFC 2119 Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels," March 1999

For More Information

For additional Cisco IT best practices, visit Cisco on Cisco: Inside Cisco IT www.cisco.com/go/ciscoit

Note

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.

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.**