Cisco 7600 Series Solution and Design Guide—Consolidated POP

The Cisco 7600 Series Internet Router is a solution for service providers, offering high-performance optical networking with uncompromising line-rate packet services. This document addresses one of the key applications of Cisco 7600 in the service provider market—consolidated point of presence (POP).

Consolidated POP

There are three main functional areas within a POP architecture: access aggregation, core backbone (including Internet peering), and local server farms. Top-tier service providers typically deploy a segregated architecture at their primary POPs, where each functional area is implemented with a separate set of networking equipment. In other words, there is a set of routers for access aggregation, a set of routers as core backbone, and a set of routers and switches for the local server farm. Such a POP architecture is appropriate for large POPs with a lot of access and backbone connections and plans for significant growth.

For smaller POPs where scalability requirements are not as intensive, a segregated POP architecture may not be the most efficient or cost-effective solution, because the amount of connections to be supported may not justify a dedicated set of networking equipment for each of the three main functional areas. Instead, one can consolidate the three functional areas into a single set of networking equipment, with each piece of equipment involved in supporting multiple functions, to optimize equipment utilization and cost-effectiveness. Such an architecture is referred to as a consolidated POP.

A well-designed consolidated POP architecture can support a wide range of connectivity requirements. It can scale and grow from a simple single-box solution to a multibox structure, with flexibility to support a wide range of interface media and speeds, while maintaining the same basic architecture for consistent network operations and modular growth. The key to achieving such a powerful solution is the proper choice of products. The Cisco 7600 is a unique product that is ideally suited for building efficient, scalable, and high-performance consolidated POP solutions.

This document presents a consolidated POP architecture using the Cisco 7600. It first provides a configuration overview of the Cisco 7600 to describe the product components and configuration guidelines. Then it presents the Cisco 7600 solution in consolidated POP, starting from a single-box solution to double-box and multibox solutions with a systematic and efficient interconnect architecture, and discusses traffic engineering and guidelines for sizing uplinks. The document provides technical details and reasoning for the architecture, to help readers to adapt the architecture components and design techniques to meet their customers' networking requirements.

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

Cisco 7600 Configuration Overview

The Cisco 7600 is an integrated router solution leveraging the Catalyst® 6500 chassis and enhanced core components, together with a new set of high-performance WAN linecards called Optical Services Modules (OSMs). The following is a summary of the Cisco 7600 system components and configuration options.

Chassis

There are three chassis available: a six-slot chassis and a nine-slot chassis both with horizontal slots, and a nine-slot chassis with vertical slots. The six-slot chassis provides six horizontal slots with side-to-side airflow and redundant power supplies. The slots are numbered 1 to 6 from top to bottom. The nine-slot chassis provides nine horizontal slots with side-to-side airflow and redundant power supplies. The slots are numbered 1 to 9 from top to bottom. The other nine-slot chassis provides nine vertical slots with front to back airflow and redundant power supplies. The slots are numbered 1 to 9 from right to left. In all three chassis, slot 1 is for the Supervisor Engine 2 (SUP2), while all other slots can support all linecards. Optionally, slot 2 can support a redundant SUP2, slot 5 can support a Crossbar Switch Fabric Module, and slot 6 can support a redundant Crossbar Switch Fabric Module.

Core Components

- The Supervisor Engine 2 (SUP2) is the processor card, and it has two built-in GE switch ports. An optional redundant SUP2 is supported, and the two GE ports on the redundant SUP2 can be active all the time.
- The Policy Feature Card 2 (PFC2) is a daughter card on SUP2. It is the layer 3 hardware-switching engine, and it also provides hardware-based access-list checking, QoS classification, and input policing.
- The Multilayer Switch Feature Card 2 (MSFC2) is a daughter card on SUP2 that provides layer 3 routing and control functions. MSFC2 has its own CPU and memory to run routing protocols and build the CEF forwarding table, which is then downloaded to the PFC2. MSFC2 is not part of the data path; in other words, it does not perform any per-packet lookups. Per-packet lookups are performed by the PFC2 in hardware.
- The Switch Fabric Module is a crossbar switch and provides two 8-Gbps full-duplex connections to each slot on the chassis for an aggregate switching capacity of 256 Gbps. The Switch Fabric Module is an internally buffered non-blocking crossbar switch. A redundant Switch Fabric Module is supported.

Optical Services Modules

These are high-performance fixed-configuration optical linecards equipped with the Cisco patented Parallel Express Forwarding (PXF) technology to support high-touch line-rate packet services such as destination-sensitive accounting/billing and QoS. The PXF IP Services Processor is an ASIC which contains a 4-by-4 array of CPU cores to perform per-packet service tasks in a parallel and pipelined fashion. The PXF technology offers feature performance as if the features were implemented in hardware, and at the same time provides feature upgrade flexibility via software. The following OSMs will be available:

- 8-port OC-3c / STM-1c POS with 4 GE ports
- 16-port OC-3c / STM-1c POS with 4 GE ports
- 2-port OC-12c / STM-4c POS with 4 GE ports
- 4-port OC-12c / STM-4c POS with 4 GE ports
- 1-port OC-48c / STM-16c POS with 4 GE ports
- 2-port OC-12c / STM-4c ATM with 4 GE ports
- 4-port Gigabit Ethernet WAN

Other Cisco 7600 Components

The Cisco 7600 supports all existing interface modules from the Catalyst 6000 family, such as the FlexWAN module (a module with two slots to support two port adapters from the Cisco 7000 router family), 10/100 Ethernet modules, GE modules, ATM LANE/MPOA modules, and voice modules.

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

Cisco 7600 Configuration Notes

The following are some basic guidelines on the configuration and operation of the Cisco 7600.

- SUP2 with PFC2 and MSFC2 are mandatory in the Cisco 7600. Redundant SUP2 is optional. The two GE switch ports on a redundant SUP2 can be used just like any GE switch ports on the chassis.
- Crossbar Switch Fabric Module is optional but strongly recommended. The Cisco 7600 has two backplane connections – the crossbar with a 256 Gbps capacity, and the data bus with a 32 Gbps capacity. Without the Switch Fabric Module, there is no crossbar, and all linecards will use the data bus. With the Switch Fabric Module, linecards that are fabric-enabled will use the crossbar, while non-fabric-enabled linecards will continue to use the bus. In other words, a packet going from one fabric-enabled linecard to another fabric-enabled linecard will use the crossbar, while a packet going from a non-fabric-enabled linecard to any linecard, or from any linecard to a non-fabric-enabled linecard, will use the data bus. Currently the only fabric-enabled linecards are all the OSMs and the new 16-port fabric-enabled GE switch modules. Both the WAN ports and the GE switch ports on an OSM use the crossbar. The two GE switch ports on SUP2 also use the crossbar.
- GE switch ports are found on three types of components on the Cisco 7600: OSMs, GE switch modules, and SUP2. GE switch ports on OSMs, SUP2, and the new 16-port fabric-enabled GE switch modules use the crossbar, while GE switch ports on non-fabric-enabled modules use the data bus. Other than which backplane connection they use, all GE switch ports on a Cisco 7600 are the same, regardless of whether they are on the SUP2, on an OSM, or on a GE switch module. GE switch ports are L2 ports, with L3 support via VLANs. All GE switch ports on a Cisco 7600, regardless of slot location, can be used together to form VLANs and Gigabit EtherChannels.
- Each OSM has a PXF-based feature performance capacity of up to 6 Mpps full duplex for high touch packet services such as QoS. The PXF provides packet services only for the WAN ports on the OSMs, including the 4 GE WAN ports on the 4-port GE WAN OSM, but not the four GE switch ports on the OSMs.
- Aside from the availability of PXF services, another major difference between GE WAN ports and GE switch ports is that the GE WAN ports are layer 3 ports and do not support layer 2 functions such as bridging (VLAN) on the GE switch ports. In other words, a group of GE switch ports on the Cisco 7600 can be configured as a VLAN, but GE WAN ports cannot be part of any VLAN. OSM WAN ports (including the GE WAN ports) are controlled and configured by MSFC2, while all GE switch ports (regardless of whether they are on OSMs, GE switch modules, or SUP2) are controlled and configured by SUP2.
- Current OSMs do not perform packet switching; all packet switching is performed by the PFC2.
- Performance of Cisco 7600 depends on which linecards are present. When all linecards are fabric-enabled, and with the presence of the Switch Fabric Module, Cisco 7600 performance capacity is 30 Mpps. When one or more of the linecards are not fabric-enabled, regardless of the presence of the Switch Fabric Module, performance capacity is 15 Mpps. These performance numbers refer to packet switching by SUP2/PFC2. Some of the GE switch modules (not the OSMs listed in this document) support distributed forwarding, and that can increase the overall performance capacity of the Cisco 7600. As a quick reference to gauge the performance requirement of a given configuration, line-rate pps for various interfaces are shown in the following table. Line-rate pps refers to the packet rate of a given packet size that will load up an interface to 100%. So these are "worst case" reference values, as the average utilization of real world networks is far below 100%. From the table, 30 Mpps can saturate 20 GE interfaces with 64-byte packets, or 66 GE interfaces with 256-byte packets.

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

Table 1

	Line-rate pps (Uni-directional)		
Interface	**64-byte* pkts**	**128-byte* pkts**	**256-byte* pkts**
GE	1,488 Kpps	844 Kpps	453 Kpps
POS OC3c/STM-1c	353 Kpps	160 Kpps	76 Kpps
POS OC12c/STM4c	1,413 Kpps	640 Kpps	306 Kpps
POS OC48c/STM-16c	5,651 Kpps	2,560 Kpps	1,223 Kpps
ATM OC3c/STM-1c	177 Kpps	118 Kpps	59 Kpps
ATM OC12c/STM-4c	706 Kpps	471 Kpps	235 Kpps

* Packet sizes are based on Ethernet encapsulations. For example, a 64-byte Ethernet frame has 46 bytes of IP payload, and when encapsulated in PPP becomes a 53-byte packet when carried over POS; the POS pps numbers in this table reflect that. The same 46 bytes of IP payload becomes two 53-byte ATM cells because of ATM overhead, and that is why the ATM pps numbers in the table are roughly half of the corresponding POS pps numbers with 64-byte packets.

Consolidated POP with Cisco 7600 – Single-Box Solution

The Cisco 7600 with its POS and ATM OSMs, FlexWAN modules, GE and FE switch modules together can support all kinds of WAN connectivity and LAN switching requirements, making it possible to implement a consolidated POP in a single box. The POS and ATM OSMs provide high-bandwidth high-performance and feature-rich connections at OC3 and OC12 speeds (up to OC48 for POS OSM), to support core backbone/peering as well as access connectivity. The GE WAN OSMs provide high-performance and feature-rich GE connections to support GE over dark fiber in metropolitan areas. For lower speed access aggregation, the FlexWAN works with a wide range of WAN port adapters from the Cisco 7000 router family to support T3/E3 and T1/E1 connections, both clear channel and channelized, as well as POS and ATM connections at up to OC3 speeds. The GE and FE switch modules connect local servers directly to the Cisco 7600 without external switches. VLANs can be configured as needed.

With the use of FlexWAN, which is not a fabric-enabled module, the switching capacity of the Cisco 7600 is 15 Mpps. This performance capacity is adequate for such configurations, because the presence of FlexWAN modules generally means that there are fewer high-bandwidth (such as OC-12 and GE) interfaces on the chassis, therefore requiring less aggregate performance capacity. The amount of performance capacity required on a chassis can be estimated from Table 1. As an example, 15 Mpps can saturate more than 40 OC3 interfaces with 64-byte packets, or almost 50 OC12 interfaces with 256-byte packets.

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

Figure 1 shows a Cisco 7600 in a single-box consolidated POP configuration.

Figure 1

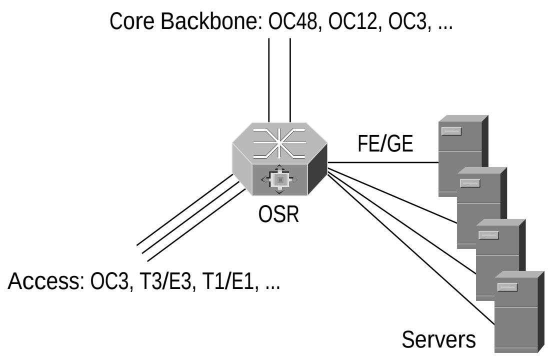

Consolidated POP with Cisco 7600 – Double-Box Solution

The use of two Cisco 7600s together is a natural next step to increase the overall capacity of a consolidated POP, and it adds redundancy to the solution. The two Cisco 7600s can be interconnected by a GE link or a Gigabit EtherChannel, depending on the bandwidth requirement. This can be accomplished using the GE switch ports on the OSMs without having to purchase additional interface modules.

Both routers will support backbone and peering connections, as well as access connections and local servers. This is a characteristic of consolidated POPs, where all routers perform the same set of functions. As mentioned earlier in the configuration overview section, the switching capacity of a Cisco 7600 is 15 Mpps with non-fabric-enabled modules such as FlexWAN, and 30 Mpps when only fabric-enabled modules are present in the chassis. Therefore, the use of non-fabric-enabled modules can be confined to one router, when it is desirable to have 30 Mpps of switching capacity in the other router to support more high-bandwidth connections.

The splitting of local servers between the two Cisco 7600s can be accomplished simply with a single VLAN configured across the two routers, each connecting half of the servers. For enhanced availability, each server can be dual-homed to both routers via two separate VLANs (one on each router) or with HSRP.

Figure 2 shows two Cisco 7600s in a double-box consolidated POP configuration.

Figure 2

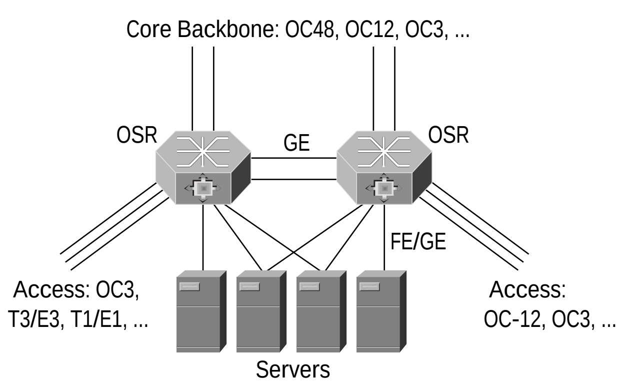

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

Consolidated POP with Cisco 7600 – MultiBox Solution

The consolidated POP architecture described so far can easily scale to multiple Cisco 7600s as the POP grows in size. Each router will still terminate both access and backbone/peering connections as well as connect to local servers. As has been mentioned earlier in the two-box solution, the switching capacity of a Cisco 7600 is higher when configured with only fabric-enabled modules. Therefore, it is recommended that FlexWAN modules be confined to one or a few routers, so that more switching capacity can be available in the other routers to support additional high-bandwidth (GE, OC12, and above) interfaces.

In a multibox consolidated POP solution, a systematic and efficient architecture to interconnect the routers is needed for performance, scalability, operational efficiency, and modular growth. Figure 3 shows an interconnection architecture using the GE switch ports on the OSMs of Cisco 7600s. The first two Cisco 7600s, A and B, are configured as the roots of two VLANs (VLAN-A with solid lines and VLAN-B with dotted lines in Figure 3). Additional Cisco 7600s are then dual-homed to the two roots as shown. In this topology, all Cisco 7600s are at most two switch hops away from each other, even if one of the two roots fails. Each active GE link carries traffic to or from at least one of its terminating routers, so no bandwidth is wasted in carrying purely transit traffic. This is a significant factor in cost-performance design to optimize the utilization of system resources. With the two VLANs interconnecting the routers, intra-POP traffic (such as traffic from router C to router D) can be load-balanced (at Layer 3) between path-A (solid lines in Figure 3) and path-B (dotted lines in Figure 3), and so there are no idle links in this configuration.

To optimize performance, the new GE modules with Distributed Forwarding Card can be used on root routers A and B to support the interconnection VLANs, and the two root routers should not be configured with FlexWANs. For additional interconnection bandwidth, Gigabit EtherChannels can be used instead of single GE links.

Local servers can be split among the multiple Cisco 7600s simply with a single VLAN configured across the multiple routers, or with additional VLANs as needed. For enhanced availability, each server can be dual-homed (or multi-homed) to two (or more) routers, via separate VLANs (one on each router) or with HSRP. It is also possible to use Server Load Balancing (a post-FCS feature on the Cisco 7600) on each router to send as much traffic to directly connected servers as possible, in order to reduce inter-router traffic in the POP for performance optimization.

Figure 3

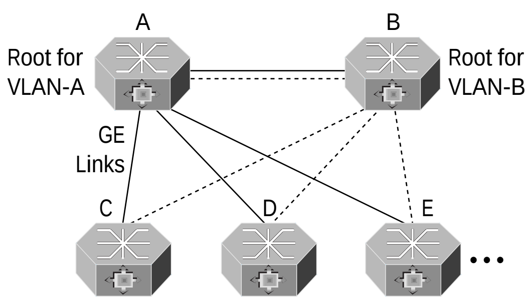

Traffic Engineering in a Consolidated POP

Each router in the consolidated POP performs multiple roles, one of which is access aggregation. When there is access aggregation, there is "uplink"; and when there is uplink, there is the issue of over-subscription. Uplink here consists of inter-POP links (backbone/peering connections on each router) as well as intra-POP links (GE switch links from each router to the other routers in the consolidated POP). Over-subscription refers to provisioning a total amount of uplink bandwidth to be less than the total amount of access bandwidth in order to reduce overall network costs. The justification is based on the assumption that access links are not 100% utilized, therefore total uplink bandwidth can be less than total access

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

bandwidth. For example, if there is an aggregate access bandwidth of 10 Gbps with an average utilization of 50%, then 5 Gbps of uplink bandwidth is needed. The 5 Gbps of uplink bandwidth will be fully loaded with traffic from the 10 Gbps of aggregate access bandwidth at 50% average utilization.

A common mistake of over-subscription is double over-subscription. If some of the access links being aggregated are already uplinks from downstream, then those "intermediate uplinks" should not be candidates for any further over-subscription. This is simply because those "intermediate uplinks" will be heavily loaded as they are already over-subscribed with access traffic from downstream.

The first step to size an uplink is to know how much access bandwidth is being aggregated. Table 2 shows the total interface bandwidth for each OSM, counting only the WAN ports (i.e. GE switch ports on the OSM are not counted). For example, a 16-port OC3 POS OSM has 16 x 0.155 Gbps = 2.5 Gbps. Summing up the per-OSM WAN bandwidth for all OSMs on a Cisco 7600 gives the total amount of access bandwidth that needs to be uplinked. For example, a chassis with four 16-port OC3 POS OSMs will have a total of 10 Gbps of access bandwidth. With no over-subscription, a total uplink bandwidth of 10 Gbps is required. With an over-subscription ratio of 2:1, a total uplink bandwidth of 5 Gbps is required, and with an over-subscription ratio of 4:1, a total uplink bandwidth of 2.5 Gbps is required. To determine the appropriate over-subscription ratio requires understanding of the access traffic load characteristics, and it is important to find out if any of the access circuits are uplinks from downstream. Traffic pattern does change, so monitoring traffic volumes is an important routine in network operations.

Table 2

Optical Services Module	Total WAN Interface BW* (Approx.)
8-port OC3c / STM-1c POS	1.25 Gbps
16-port OC3c / STM-1c POS	2.50 Gbps
2-port OC12c / STM-4c POS	1.25 Gbps
4-port OC12c / STM-4c POS	2.50 Gbps
1-port OC48c / STM-16c POS	2.50 Gbps
2-port OC12c / STM-4c ATM	1.00 Gbps
4-port GE WAN	4.00 Gbps

*For simplicity, bandwidth here is based on nominal interface bandwidth (for example OC-3c/STM-1c is 155 Mbps), and does not account for full duplex, because both access links and uplinks are full duplex. For ATM interface, a 20% adjustment has been included to account for ATM cell overhead.

In the consolidated POP described in this paper, each router has a set of access links to aggregate, and the uplinks consist of inter-POP links (backbone/peering connections) as well as intra-POP links (GE switch links to the other routers in the POP). Uplink sizing involves sizing both the peering/backbone connections and the GE switch links. For ease of subsequent description, let x denote the amount of backbone/peering bandwidth on a given router in the consolidated POP, and let y denote the total amount of backbone/peering bandwidth on all routers in the consolidated POP. Without any specific knowledge of traffic patterns, one can generally assume that a packet entering the POP from one of its access links has equal probability of going to any one of the backbone/peering links in the POP, weighted by the relative bandwidth of those

Copyright © 2001 Cisco Systems, Inc. All Rights Reserved.

backbone/peering links. Based on this assumption, the fraction *x/y* of the access traffic from the given router will go on its own inter-POP backbone/peering links. The rest *(1 – x/y)* of its access traffic will go on its intra-POP GE switch links to get to the other routes for their backbone/peering links.

The above assumption and estimation do not take into consideration the amount of local (intra-POP) traffic among all the access links of the POP, as well as the effect of local servers and cache engines in reducing the amount of inter-POP traffic. If a certain percentage of traffic is estimated to be local (intra-POP), that amount should be treated separately from the remaining inter-POP traffic in uplink sizing. So the backbone/peering connections are sized based on inter-POP traffic alone, while the intra-POP GE links are sized based on both the inter-POP and the intra-POP traffic. Over-subscription factors and redundancy considerations are then applied accordingly to determine the needed uplink bandwidth for each router in the POP.

The intra-POP GE links in this architecture use the GE switch ports on the OSMs, so one should make sure there is an adequate amount of GE switch ports on each chassis configuration to support GE uplinks. From Table 2, it is obvious that the POS and ATM OSMs have more than enough GE switch ports for uplink, while the 4-port GE WAN OSM has no GE switch ports. The proper combination of OSMs should have enough GE switch ports from the POS and ATM OSMs to cover uplink requirements for the GE WAN OSMs (if any), and to provide GE server farm connectivity. For example, a Cisco 7600 configured with five 4-port GE WAN OSMs and four 16-port OC3 POS OSMs has a total access bandwidth of 30 Gbps. Assuming an over-subscription ratio of 2:1, 15 GE uplinks are needed, and the 16 GE switch ports from the four POS OSMs meet this requirement perfectly. If more GE switch ports are needed, 16-port GE switch modules can be added to the Cisco 7600. When there is an abundant supply of GE switch ports in a configuration, there is less pressure for aggressive over-subscription on the intra-POP GE links. This is an advantage of the Cisco 7600 solution.

Summary

The Cisco 7600 is an ideal solution for consolidated POPs. It supports access aggregation, core backbone/peering, and local server farm connectivity with a simple and scalable architecture. This document has presented a reference architecture and design guidelines for consolidated POPs. The architecture uses the unique capabilities and features of the Cisco 7600 to achieve an efficient, scalable, and high-performance solution. The Cisco 7600 is also an ideal fit in another service provider solution, metro aggregation, which is covered in a separate white paper.

Cisco Systems

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

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the
Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • 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 • Scotland • Singapore • Slovakia Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001 Cisco Systems, Inc. All rights reserved. Catalyst, Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries. All other brands, names, or trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0011R)