



Quality of Service Configuration Guidelines, Cisco IOS XE Gibraltar 16.10.x (Cisco NCS 520 Series)

Quality of Service Configuration Guidelines 2

Quality of Service 2

Quality of Service Configuration 2

Global QoS Limitations 2

Port-Channel with EFP 5

Sample Hierarchical Policy Designs 6

Ingress Hierarchical Policing 7

QoS Policer and Shaper Calculation 7

Classification 8

QoS Marking 10

Traffic Policing 11

Traffic Shaping 13

Congestion Management 14

Scheduling 15

Additional References 15

Revised: July 3, 2019

Quality of Service Configuration Guidelines

This document outlines Quality of Service features and limitations available on the Cisco NCS 520 Routers and contains the following sections:

Quality of Service

QoS refers to the ability of a network to provide improved service to selected network traffic over various underlying technologies including ATM, Ethernet and 802.1 networks, SONET, and IP-routed networks. In particular, QoS features provide improved and more predictable network service by implementing the following services:

- Supporting guaranteed bandwidth
- Improving loss characteristics
- Avoiding and managing network congestion
- Shaping network traffic
- Setting traffic priorities across the network

Quality of Service Configuration

This document provides details on the platform-dependent implementation of QoS on the Cisco NCS 520 Series Router.

Global QoS Limitations

The following limitations apply to multiple QoS features for the router:

- When EVCs under a physical interface have a QoS policy attached, the following limitations apply:
 - The port-level policy is limited to the class-default class
 - Only the **shape** command is supported in the port-level policy
- The router supports up to 64 unique QoS classification service instances in a given bridge domain. QoS service instances refer to ports, VLAN classes, EFPs associated with a QoS classification policy.
- Modification of class-map definitions while applied to an interface or Ethernet Flow Point is not supported.
- Policy validation—Some QoS policy configurations are not validated until you apply the policy-map to an interface or Ethernet Flow Point. If a QoS configuration is invalid, the router rejects the configuration when you apply it to an interface. In some cases, a QoS configuration may be rejected due to hardware resource exhaustion or limitations. If you receive such an error message, detach the policy and adjust your QoS configuration.
- The **match-all** keyword is supported only for QinQ classification.
- QoS is supported on POS interfaces on optical interface module.

- QoS does not account for CRC values on an interface and assumes that the value is 2 bytes. CRC differences can cause accuracy issues for 2 to 3% of the 128-byte traffic.
- The ICMPv4 packets classification based on ACL attached on interface is not supported.
- Even though the support for precedence marking is not claimed for the Cisco IOS XE 16.8.1 release, it works except for conditional marking.
- If policers are configured as flat ones, then 1000 TCAM entries can be utilized. If the policers are hierarchical, then only 768 entries can be utilized.
- When a shaper is applied on two interfaces that receive unknown multicast traffic, then the total rate of traffic moving out of both the interfaces is equal to the configured shaper value.
- In EFP policies, the class queue (cosq) value with priority level 1 should be greater than that of the class queue (cosq) value with priority level 2.
- ICMP packets that are generated locally, are sent through data queues, and hence are subjected to egress QoS process as a regular egress traffic.
- A Layer 4 port-based ACL, when used as a classifier in a policy-map does not classify the traffic if the packets are fragmented, even though the traffic hits the class-map.
- For ports mapped to internal buffer, the total available buffer is 600 KB.
- Only 768 TCAM rules can accommodate policing actions in case of Hierarchical policer.



Note

768 is the number of child policer, 768 parent policer.

If a flat policer is used in the policy-map then all 1024 TCAM entries can accommodate policing actions.

Restrictions for Ingress QoS

Restrictions for Ingress QoS in the Cisco IOS XE Release 16.8.x:

- EC main interface
 - Only policing and marking are supported.
 - A class-map can have any type of filter, including the **match vlan** and **match service instance** commands.
- EC EVC/TEFP
 - Only policing and marking are supported.
 - Match service instance is *not* supported.
- Member links
 - Only policing and marking are supported.
 - Policy-map on a member link is *not* supported with EVC configured at the port-channel level.
- Policy-map application is allowed only on the EC main interface, EC member link, or EC EVC.

•

Following are the general ingress QoS restrictions:

- When ingress policer is applied with a two rate Three Color Marker (trTCM), Layer 2 header packets received with the Drop Eligible Indicator (DEI) or Canonical Format Identifier (CFI) bit set will be dropped.
- Violate packet or byte statistics is not supported. The exceed statistics include both exceed and violate actions.
- In a two-level policer cases, the higher-level policer can be a single-rate three-color marking only.
- If a policer is configured on a logical interface (for example match on VLAN or EFP), then single-rate three-color marking is only supported.
- The layer 4 port operator NEQ is not supported.
- The layer 4 port range is only 32 system wide. The port range includes both source and destination ports and for system wide only 32 different ranges can be configured.
- The range configuration for source and destination ports is not supported in the same rule.
- Even if the policy map action is set to cos and not qos-group, then application configures the action as both cos and qos-group.
- In the MAC ACL QoS, only source MAC and destination MAC addresses are supported.
- In the IP ACL QoS, source IP, destination IP, IP, DSCP, layer 4 source and destination ports are supported.
- The policer accuracy is 8 KB or 1 percent, whichever is higher.
- In an ACL-based PHB or a layer 2 classification, the L1 class can be set to EFPs only but not outer or inner VLAN.
- Table map configuration is not supported and therefore remarking based on table map profile is also not supported.
- IPv6 classification is not supported.
- Conditional marking and marking actions are not concurrent across different levels or at the same level.
- Due to ASIC limitation, the ACL-based rules have higher precedence over non-ACL classification within a class map that contains both the ACL and non-ACL classifications.
- MAC ACL will not work for reserved MAC addresses.

Restrictions for Egress QoS

- The maximum number of classes supported on the policy map is 8, which includes class class-default; 7 user-defined classes and class class-default is supported.
- The maximum number of port-channel interfaces that can be created and supported for QoS on the router is 8.
- WRED is not supported.
- EC main interface
 - Classification statistics for the policy-map on a port-channel main interface are *not* supported as no queues are allocated for a port-channel main interface.
 - Queuing actions are *not* supported.
- EC EVC/TEFP

- Classification statistics for the policy-map on port-channel EVC/TEFP are *not* supported as no queues are allocated for port-channel EVC/TEFP.
- Queuing actions are *not* supported.
- EC Member links
 - Queuing action is supported on port-channel member links.
 - The running configuration displays the first member link on the first policy applied in a service-policy configuration.
 - The **show policy-map interface brief** command only displays the policy-map applied on the running configuration.
 - Applying a same policy again on other member links where a policy-map was already applied will not display any error. A differently named policy if applied again will display an error.
 - Policy-map application is allowed only on either EC main interface, or EC member link, or EC EVC.
- Hierarchical QoS
 - Single class-default based policy only is allowed on interface (port-level).
 - Nested policy is allowed under EFP provided that parent policy matches on class-default and child policy matches on QoS groups.
- When a class-of-service queue is congested, locally generated control plane packets such as ICMP are expected to be dropped as they are subjected to the egress policy using the class-of-service queues.

Port-Channel with EFP

Port-channel with EFP

The following features are supported for ingress policy-map for port-channel with EFP:

- Marking
- Policing
- Conditional marking
- marking and policing
- The classification criteria is VLAN, EFP, DSCP, ACL, or Cos in child.

Example for port-channel with EFP

```
policy-map cos_child
class cos0
set cos1
!
policy-map efp_pc_ingress
class vlan100
police cir 10m
service-policy cos_child
!
end
```

EFP of Port-channel with EFP Configuration

- The following features are supported for ingress policy-map for EFP on port-channel:
 - Marking
 - Policing
 - Conditional marking
 - marking and policing
 - The classification criteria is VLAN, EFP, DCSP, ACL, or Cos in child.



Note Match EFP is cannot be configured.

Member Links of Port-channel with EFP Configuration

- The following features are supported for ingress policy-map on member links of port-channel on EFP:
 - Marking
 - Policing
 - Conditional marking
 - marking and policing
 - The classification criteria is VLAN, EFP, DSCP, or ACL, Match VLAN and Cos in child.

Restrictions for Hierarchical Policies

The router supports hierarchical QoS policies with up to three levels, allowing for a high degree of granularity in traffic management.

There are limitations on the supported classification criteria at each level in the policy-map hierarchy. The following limitations apply when configuring hierarchical policy-map classification:

- The topmost policy-map in a three-level hierarchy only supports classification using class-default.

Sample Hierarchical Policy Designs

The following are examples of supported ingress policy-map configurations:

- Three-Level Policy—You can only apply a three-level policy to a physical port on the router. A three-level policy consists of:
 - Topmost policy: class-default
 - Middle policy: outer vlan, inner vlan, or service instance
 - Lowest policy: outer cos, inner cos, DSCP, IPv4 ACL, or MAC ACL

Sample Policy

The following sample policy uses a flat class-default policy on the port and VLAN policies on EFP interfaces to unique QoS behavior to each EFP.

```
Policy-map port-shaper
Class class-default
Shape average percent 70
Service-policy child
```

```
Policy-map child
Class qos-1
Bandwidth percent 20
Shape average 200m
Service-policy child1
Class qos-2
Bandwidth percent 75
```

- Two-Level Policy
 - Topmost policy: match vlan or match service instance
 - Lowest policy: match cos, match dscp, match IP ACL or match MAC ACL
- Two-Level Policy
 - Topmost policy: class-default
 - Lowest policy: match vlan or match service instance
- Flat policy: match ip dscp
- Flat policy: match vlan inner
- Flat policy: class-default

Ingress Hierarchical Policing

Ingress policing is supported at two levels of the policy-map and the third level is supported only for remarking.

- Ingress policing
 - Port and EFP level
 - EFP and Class level
 - Port and Class level

QoS Policer and Shaper Calculation

Table below summarizes the packet accounting information used to make policer and shaper calculations on the Cisco NCS 520 Router.

Table 1: QoS Accounting Calculation

Feature	Direction	Traffic Type	Values Counted
Policing	Ingress	IPv4	L2 overhead, VLAN tag, CRC
Shaping	Egress	IPv4	L2 Ethernet overhead, VLAN tag, CRC, preamble, IPG

The following considerations also apply when understanding QoS policer and shaper calculations:

- Egress shaping is applied at layer 1.
- Ingress packet length accounting is performed at egress.

Classification

Classifying network traffic allows you to organize packets into traffic classes or categories on the basis of whether the traffic matches specific criteria. Classifying network traffic (used in conjunction with marking network traffic) is the foundation for enabling many quality of service (QoS) features on your network.

Ingress Classification Limitations

The following limitation apply to QoS classification on the Cisco NCS 520 Series Router:

- IPv6 QoS ACL is not supported.

Egress Classification Limitations

- When applying a QoS policy to a link aggregation group (LAG) bundle, you must assign the policy to a physical link within the bundle; you cannot apply the policy to the LAG bundle or the port channel interface associated with the bundle.
- MPLS classification using EXP values in an egress policy are applied to normal IP packets in an MPLS core network. When Egress classification for EXP values are converted to equivalent IP precedence values, the first 5 bits in the DSCP values will be used to classify the MPLS packets. However, normal IP packets will be classified as well.

It is recommended to move the EXP classification to ingress policy and egress classification to be moved to the QoS group set from ingress policy to avoid classification of normal IP packets.

Additional Classification Limitations

- The topmost policy-map in a three-level hierarchy only supports classification using class-default.

Classifying Traffic using an Access Control List

You can classify inbound packet based on an IP standard or IP extended access control list (ACL). By default, TCAM optimization or expansion method is used. Both Security ACL and QoS ACL can be configured on the same interface. Follow these steps to classify traffic based on an ACL:

1. Create an access list using the **access-list** or **ip access-list** commands

2. Reference the ACL within a QoS class map using the **match access-group** configuration command
3. Attach the class map to a policy map

Limitations and Usage Guidelines

The following limitations and usage guidelines apply when classifying traffic using an ACL:

- QoS ACLs are supported only for IPv4
- QoS ACLs are supported only for ingress traffic
- You can use QoS ACLs to classify traffic based on the following criteria:
 - Source and destination host
 - Source and destination subnet
 - TCP source and destination
 - UDP source and destination
- Named and numbered ACLs are supported.
- You can apply QoS ACLs only to the third level class (bottom-most).
- The following range of numbered access lists are supported:
 - 1-99—IP standard access list
 - 100-199—IP extended access list
 - 1300-1999—IP standard access list (expanded range)
 - 2000-2699—IP extended access list (expanded range)
- You must create an ACL before referencing it within a QoS policy.
- Deny statements within an ACL are ignored for the purposes of classification.
- Classifying traffic based on TCP flags using an ACL is not supported.
- Classifying traffic using multiple mutually exclusive ACLs within a **match-all** class-map is not supported.
- Classifying traffic on a logical/physical level using an ACL is not supported.
- Applying QoS ACLs to MAC addresses is not supported.
- A given command can consume multiple matching operations if you specify a source and destination port, as shown in the following examples:
 - **permit tcp any lt 1000 any**—Uses one port matching operation
 - **permit tcp any lt 1000 any gt 2000**—Uses two port matching operations
 - **permit tcp any range 1000 2000 any 400 500**—Uses two port matching operations
- Only the following combination of matches are currently supported for Ingress policies:
 - Combination A: DSCP, Outer COS, UDP/TCP Source and Destination port number, IP SA/DA

- Combination B: IP SA/DA, Outer COS, Inner COS, DSCP
- Combination C: MAC DA, Outer COS, Inner COS, DSCP

Configuring Multiple Match Statements

Support for a single **match** or **match-any** command in a given QoS class-map is shown in the following example:

Example for a Single **match** or **match-any** Class Map

```
class-map match-any my-restrict-class_01
  match qos-group 2

class-map match-any my-restrict-class_03
  match cos 3
```

Traffic Classification Using Match EFP Service Instance Feature

Service Provider configurations have various service instances on the PE. QoS policy-maps are applied on these service instances or group of service instances. Cisco IOS XE Release 3.9S introduces the Match EFP Service Instance feature. The benefits of this feature are:

- Identify the various types of service-instances like EFP, Trunk EFPs
- Apply policies on these service instances at the port
- Manage bandwidth and priority across the service instances on the port and across classes within the service instance
- Apply policies on a group of transport service instances such as applying similar policies to a group of EFPs.

Restrictions for Configuring Match Service Instances

- Ethernet service instances configured under the interface can be classified in a class of a policy-map. The class can match on a group or set of match service instance statements.

```
class-map match-any policeServiceInstance
  match service instance ethernet 100
  match service instance ethernet 200
```

- Match service instance supported at both Ingress level.
- match service instance and match PHB per flows classification are defined at respective levels in the policy hierarchy under the port.
- Match EFP policy-map can be configured only on the port and *not* under the service instance.

QoS Marking

QoS marking allows you to set a desired value on network traffic to make it easy for core devices to classify the packet.

Overview of Marking

The Cisco NCS 520 Series Router supports the following parameters with the **set** command:

- **set cos**
- **set discard-class**
- **set dscp**
- **set ip dscp**
- **set qos-group**

CoS Marking Limitations

The following limitations apply when configuring CoS marking:

- **set cos**—This set action has no effect unless there is a egress push action to add an additional header at egress. The COS value set by this action will be used in the newly added header as a result of the push rewrite. If there are no push rewrite on the packet, the new COS value will have no effect.
- The **set cos inner** command is not supported.

Ingress Marking Limitations

The following limitations apply to QoS marking on the Cisco NCS 520 Series Router:

- The Cisco NCS 520 Series Router does *not* support hierarchical marking.
- You can configure marking and policing for any number of classes on any one of the three levels of the policy-map hierarchy. If you configure marking on one level, you can configure policing without marking (transmit, drop) on another level.
- In the flow of the packet, if both ingress and egress markings are needed, you must classify the packet with the ingress marked phb class at egress and remark it to preserve the ingress marking. Marking in class-default of the ingress marked packets will not preserve the ingress markings.

Additional Marking Limitations

The following additional marking usage guidelines apply:

- Marking is supported on Etherchannel interfaces and individual member links; however, you cannot configure marking on both interface levels at once.

Traffic Policing

Traffic policing allows you to control the maximum rate of traffic sent or received on an interface, and to partition a network into multiple priority levels or class of service (CoS). This section describes the policing limitations and configuration guidelines for the router.

The router supports the following policing types:

- Single-rate policer with two color marker (1R2C) (color-blind mode)

- Two-rate policer with three color marker (2R3C) (color-blind mode)

Supported Commands

The Cisco NCS 520 Router supports the following policing commands on ingress interfaces:

- **police** (percent)—**police cir percent percentage [burst-in-msec] [bc conform-burst-in-msec ms] [be peak-burst-in-msec ms] [pir percent percentage] [conform-action action] [exceed-action action] [violate-action action]]]**
- **police** (policy map)—**police cir bps [[bc] normal-burst-bytes [maximum-burst-bytes | [be] [burst-bytes]]] [pir bps [be burst-bytes]] [conform-action action] [exceed-action action] [violate-action action]]]**
- **police** (two rates)—**police cir cir [bc conform-burst] [pir pir] [be peak-burst] [conform-action action] [exceed-action action] [violate-action action]]]**

The Cisco NCS 520 Router supports the following policing commands on egress interfaces:

- **bandwidth** (policy-map class)—**bandwidth {bandwidth-kbps | remaining percent percentage | percent percentage} [account {qinq | dot1q} aal5 subscriber-encapsulation]**
- **bandwidth remaining ratio**—**bandwidth remaining ratio ratio [account {qinq | dot1q} [aal5] {subscriber-encapsulation | user-defined offset}]**
- **police** (policy map)—**police cir bps [[bc] normal-burst-bytes [maximum-burst-bytes | [be] [burst-bytes]]] [pir bps [be burst-bytes]] [conform-action action] [exceed-action action] [violate-action action]]]**
- **priority**—**priority {bandwidth-kbps | percent percentage} [burst]**

Egress policing is not supported for the Cisco IOS XE 16.8.x release.

Supported Actions

The Cisco NCS 520 Series Router supports the following policing actions on ingress interfaces:

- **transmit**
- **drop**
- **set-qos-transmit**
- **set-cos-transmit**
- **set-dscp-transmit**
- **set-discard-class-transmit**

Percentage Policing Configuration

The router calculates percentage policing rates based on the maximum port PIR rate. The PIR rate is determined as follows:

- Default—Port line rate
- Speed command applied—Operational rate
- Port shaping applied to port—Shaped rate

Ingress Policing Limitations

The following limitations apply to QoS policing on the Cisco NCS 520 Series Router:

- If you configure a policer rate or burst-size that the router cannot achieve within 1% accuracy, the configuration is rejected. The command output presents recommendations for the closest possible lower and higher configuration value.
- You can configure marking and policing for any number of classes on any one of the three levels of the policy-map hierarchy. If you configure marking on one level, you can configure policing without marking (transmit, drop) on another level.
- If you configure marking using the **set** command, you can only configure policing on that level using the transmit and drop command.
- If you configure a policer using a **set** command, you cannot use the **set** command at other levels of the hierarchical policy-map.

Traffic Shaping

Traffic shaping allows you to control the speed of traffic that is leaving an interface in order to match the flow of traffic to the speed of the receiving interface. Percentage-based policing allows you to configure traffic shaping based on a percentage of the available bandwidth of an interface. Configuring traffic shaping in this manner enables you to use the same policy map for multiple interfaces with differing amounts of bandwidth.

Additional Shaping Limitations

The following additional shaping usage guidelines apply in Release 3.9:

- Policies using shaping are supported only on individual member links of an etherchannel. Applying a shaping policy directly on an etherchannel interface is not supported.
- Class-based shaping is supported at all levels.

Configuring Egress Shaping on EFP Interfaces

Configuring an EFP port shaper allows you to shape all EFPs on a port using a port policy with a class-default shaper configuration, as in the following partial sample configuration:

```
policy-map port-policy
  class class-default
    shape average percent 50
policy-map efp-policy
  class class-default
    shape average percent 25
    service-policy child-policy
policy-map child-policy
  class phb-class
    <class-map actions>
```

The following configuration guidelines apply when configuring an EFP port shaping policy:

- When the configuration specifies a shaper rate using a percentage, the router calculates the value based on the operational speed of a port. The operational speed of a port can be the line rate of the port or the speed specified by the **speed** command.
- The rates for **bandwidth percent** and **police percent** commands configured under a port-shaper are based on the absolute rate of the port-shaper policy.
- You can combine a port shaper policy (a flat shaper policy with no user-defined classes) with an egress EFP QoS shaping policy.

- Configure the port shaper policy before configuring other egress QoS policies on EFP interfaces; when removing EFP QoS configurations, remove other egress EFP QoS policies before removing the port shaper policy.

Congestion Management

Congestion management features allow you to control congestion by determining the order in which packets are sent out an interface based on priorities assigned to those packets. Congestion management entails the creation of queues, assignment of packets to those queues based on the classification of the packet, and scheduling of the packets in a queue for transmission.

Ingress Queuing Limitations

The Cisco NCS 520 Router does not support queuing on ingress interfaces.

Egress Queuing Limitations

The Cisco NCS 520 Router supports tail drop queuing on egress interfaces using the **queue-limit** command. The following limitations apply to egress queuing:

- Egress QoS can be applied to a total of 91 EFPs at a system level.
- If you configure a queue size that the router cannot achieve within 1% accuracy, the configuration is rejected. The command output presents recommendations for the closest possible lower and higher configuration value.
- Egress policy-map with queuing action is *not* supported on port-channel interface(LAG). The policy must be applied to the policy-maps on the member links.
- The maximum **bytes** value of the **queue-limit** *number-of-packets* [*bytes* | *ms* | *packets*] command is 2 MB.
- The **show policy-map interface** command displays the default queue-limit.
- The **queue-limit percent** command is supported.

Support for Low Latency Queuing on Multiple EFPs

The Cisco NCS 520 Router supports QoS policies that allow for low-latency queuing (LLQ) across multiple EFPs. For more information about this feature, see http://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/xr-3s/qos-plcshp-ehqos-pshape.html.

Additional Queuing Limitations

The following additional queuing usage guidelines:

- The router supports QoS policies that allow for low-latency queuing (LLQ) across multiple EFPs. For more information about this feature, see http://www.cisco.com/en/US/docs/ios-xml/ios/qos_plcshp/configuration/xr-3s/qos-plcshp-ehqos-pshape.html.
- CBWFQ is supported only on third level class.
- Queue-limit is supported only in leaf-level (per-hop behavior) classes.
- Queue-limit can not be configured without first configuring a scheduling action (bandwidth, shape average, or priority).
- Queue-limit can not co-exist with queue-limit percent.
- Queue-limit policy can be applied only on egress interface.

- Queue-limit can be configured in bytes or microseconds, or percent per class in the egress-policy.
- Default queue-limits for 1 and 10 G are 80 and 120 KB, respectively.
- Maximum queue-limit that can be configured in bytes is 200 KB.
- Ensure that you configure the queue-limit to a value greater than the default allocation value.

When a minimum value is configured for queue-limit, for example, lesser than 11000 bytes, then the frame-size of outgoing traffic should be lesser than that of the configured queue-limit value.

Scheduling

This section describes the scheduling limitations and configuration guidelines for the Cisco NCS 520 Series Router.

Ingress Scheduling Limitations

The Cisco NCS 520 Router does not support scheduling on ingress interfaces.

Egress Scheduling Limitations

- If you configure a CIR, PIR, or EIR rate that the router cannot achieve within 1% accuracy, the configuration is rejected. The command output presents recommendations for the closest possible lower and higher configuration value.
- You can only configure one **priority** value on each parent class applied to a QoS class or logical interface.
- You can only configure priority on one class in a QoS policy.
- You can not configure **priority** value and a policer in the same class.

The following limitations apply when configuring a 3-level scheduling policy on an egress interface configured as an EFP:

- Only two of the three levels can contain scheduling actions such as **bandwidth**, **shape**, or **priority**.
- Class-based excess bandwidth scheduling is supported on 2nd and 3rd level QoS classes.
- One of the levels containing scheduling actions must be the class (bottom) level.

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mcl/allreleasemcl/all-book.html

Standards and RFCs

Standard/RFC	Title
No specific Standards and RFCs are supported by the features in this document.	—

MIBs

MIB	MIBs Link
—	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

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.