

# IP Multicast Command Reference

## Global Commands

`[no] ip multicast-routing [distributed]`

Enables IP multicast forwarding. If disabled, group addressed IP packets that the router is not a member will be discarded. The default value is IP multicast routing disabled. [10.2]

The distributed keyword will enable distributed fast switching for the router. The interface command described below will enable individual interfaces for distributed fast switching. [11.1(20)CC].

`[no] ip multicast cache-headers [rtp] [<entries>]`

Allocates a circular buffer to store IP multicast packet headers that are received by the router. This command will allocate approximately a 32 kilobyte buffer. If you are low on memory, this command should not be used. Use the “show ip mpacket” command to display this buffer. This feature is used to determine 1) who is sending to what groups, 2) what the inter-packet delay and 3) if there are any duplicates or multicast forwarding loops in your network. [11.1]

When the keyword “rtp” is used, RTP headers will also be saved. This is used in conjunction with the “show ip mpacket <group> quality” command. [11.1(20)CC]

<entries> is the power-of-2 number of cache entries maintained in the circular buffer. Valid values are 10 through 20. Use caution when setting this value greater than 10 because you can use up all the memory in the router. The default value is 10, which means 1024 entry circular buffer is maintained. [11.1(22)CC, 12.0S]

`[no] ip pim rp-address <ip-address> [<group-access-list>] [override]`

Configures the PIM Rendezvous Point (RP) address for a particular group. The RP address is used by first-hop routers to send Register packets on behalf of source multicast hosts. The RP address is also used by routers on behalf of multicast hosts that want to become members of a group. These routers send Join and Prune messages toward the RP. A single RP can be configured for multiple groups described by the access-list pointer. [10.2]

If there is no RP configured for a group, the router will treat the group as dense using the dense-mode PIM techniques. If the RP for a group is learned through a dynamic mechanism, such as Auto-RP, then this command may not be required. If there is a conflict between the RP configured with this command and one learned by Auto-RP, the Auto-RP information is used. Unless, the “override” keyword is specified. [11.1]

```
[no] ip pim accept-rp {<address> | auto-rp} [<acl>]
```

When this command is entered, the router will only accept (\*,G) Joins with an RP address of <address> if G is in the group range specified by <acl>. When this command is entered with an address equal to one of the system's addresses, the system will be the RP only for the specified group range specified in <acl>. When not in the group range, the RP will not accept Joins or Register messages and will respond immediately to Registers with Register-Stop messages. There is no default setting for this command. [10.2]

When the keyword "auto-rp" is specified, Join and Register messages will only be accepted for RPs that are in the Auto-RP cache. [11.1]

If <address> is 0.0.0.0, the filter will accept any RP for any group accepted by <acl>, and deny any RP rejected by <acl>.

When multiple "ip pim accept-rp" filters are configured, they must be configured in the following order (in 11.3(4.5), 12.0(1) and newer images the following order is guaranteed by the IOS, irrespective of the configuration sequence):

```
ip pim accept-rp <specific RP address> <acl>
ip pim accept-rp auto-rp
ip pim accept-rp 0.0.0.0 <acl>
```

The following example will accept 171.69.58.88 as the RP for groups in 239.0.0.0/8, and RPs for groups in the Auto-RP cache. If the RP and group don't match the first two filters, the 3rd filter is in effect, i.e., any RP is accepted for groups permitted by ACL 2, and no RP for 224.0.1.39 and 224.0.1.40 is accepted.

```
ip pim accept-rp 171.69.58.88 1
ip pim accept-rp auto-rp
ip pim accept-rp 0.0.0.0 2
access-list 1 permit 239.0.0.0 0.255.255.255
access-list 2 deny 224.0.1.39
access-list 2 deny 224.0.1.40
access-list 2 permit any
[no] ip pim send-rp-announce <interface-unit> scope <ttl> group-list <acl>
[interval <num-seconds>]
```

This command sends an Auto-RP RP announcement message to the well known group CISCO-RP-ANNOUNCE (224.0.1.39). This command should be used in a router you want to be the RP. The RP address field inside the announcement message will contain the IP address from the <interface-unit>. <ttl> is the time-to-live in the IP header which is set. This allows for the announcements to stay inside a ttl scoped boundary. <group-list> is an access-list describing the group ranges this system is willing to be the RP for. Note that the deny clauses in the <group-list> are ignored. [11.1]

Starting with IOS versions [12.0(1.1)], if the access-list contains a "deny" entry, auto-rp will maintain a negative entry for those group ranges. This will make it easier to configure group ranges to be dense-mode only groups. An RP announcement with a denied group prefix overrides any positive announcements for the same prefix from other RPs. However, IOS versions prior to this required a deny clause. The access-list must be changed to remove these deny clauses to obtain the correct RP map.

When "interval" is specified, the interval between RP announcements is set to <number of seconds>. The total holdtime of the RP announcements is automatically set to 3 times <interval>. The default interval is 60 seconds. Tuning this interval down can reduce the time required to fail over to a secondary RP, at the expense of generating more Auto-RP messages through the entire region covered by the ttl scope. [11.2(18), 11.3(8), 12.0(3.1)]

```
[no] ip pim send-rp-discovery [<interface>] scope <ttl>
```

This command is entered in the PIM RP-mapping agent. The RP-mapping agent listens on well-known group address CISCO-RP-ANNOUNCE (224.0.1.39) that candidate RPs send to. The RP-mapping agent will send RP-to-group mappings in an Auto-RP RP discovery message to the well known group CISCO-RP-DISCOVERY (224.0.1.40). PIM DRs will listen to this group and use the RPs they learn from the RP discovery message. <ttl> is the time-to-live in the IP header which is set. This allows for the discovery messages to stay inside a ttl scoped boundary. [11.1]

When <interface> is specified, RP discovery messages will be sourced from the IP address assigned to <interface>. [12.0]

```
[no] ip pim rp-announce-filter rp-list <acl> group-list <acl>
```

This command is entered in the PIM RP-mapping agent. This command configures an incoming filter for RP announcement messages. Parameter “rp-list <acl>” configures an access-list of RP addresses that, if permitted, will be filtered or rejected for the group ranges supplied in the parameter “group-list <acl>.” If this command is not configured, all RP announcements are accepted. If you are going to use more than one RP-mapping agent, it is suggested that the filters are consistent among them so there is no conflicts in mapping state when the announcing agent goes down. [11.1]

```
[no] ip pim spt-threshold <kbps> | infinity [group-list <acl>]
```

Configures when a PIM leaf router should join the shortest path source-tree for the specified group. <kbps> is the traffic rate in kilobits per second. If a source sends at a rate greater than or equal to <kbps>, a PIM Join message is triggered towards the source to construct a source-tree. If “infinity” is used, all sources for the specified group will use the shared-tree. Specifying a “group-list” indicates what groups the spt-threshold applies to. <acl> is a reference to a simple IP access-list. When a value of 0 is specified or the group-list parameter is not used, the threshold applies to all groups. The default setting (when this command is not used), is to join the shortest path tree immediately after the first packet arrives from a new source. [11.1]

```
[no] ip pim bsr-candidate <interface> <hash-mask-len> [<priority>]
```

Configures the router to send bootstrap messages with <interface>'s address as the bootstrap-router (BSR) address, if no better bootstrap router is found. <interface> must be a PIM enabled interface. <hash-mask-len> is the mask length used by the PIMv2 hash function. This hash mask length is accepted by all routers within the same PIM domain when selecting an RP. <priority> is an integer whose value is between 0 and 255. It is 0 by default. BSRs with larger preference values are preferred over those with smaller values.

This command should only be used in “backbone” routers with good connectivities to all parts of the PIM region. E.g., a stub router that relies on an on-demand dial-up link to connect to the rest of the PIM domain is not a good candidate BSR. [11.3T]

```
[no] ip pim rp-candidate <interface> [group-list <acl>]
```

Configures to send pim version 2 candidate RP advertisement to the bootstrap RP. The IP address associated with <interface> will be advertised as the candidate RP address. The group prefixes defined by simple access-list <acl> will also be advertised in association with the RP address. RP-candidates should also be placed in the well-connected “backbone” part of the PIM domain. [11.3T]

```
[no] ip pim register-rate-limit <pps>
```

Sets a limit on the maximum number of data registers/second sent for each (S,G). If this is configured on a PIM domain border, a recommended <pps> is 2. [11.3T, 11.1(20)CC]

```
[no] ip dvmrp routehog-notification <route-count>
```

This configures the number of routes allowed within an approximate one minute interval before a syslog message is issued warning that there maybe a route surge going on in the MBONE. This is typically used to detect quickly when someone has misconfigured their routers to inject a large number of routes into the MBONE. The default value is 10,000. You can find a running count in the “show ip igmp interface” display. When the count is exceeded, you’ll see an “\*\*\* ALERT \*\*\*” string appended to the line. [10.2]

```
[no] ip dvmrp route-limit <route-count>
```

This command limits the number of DVMRP routes advertised over an interface enabled to run DVMRP. That is a DVMRP tunnel, an interface where a DVMRP neighbor has been discovered, or an interface configured to run “ip dvmrp unicast-routing.” The default value is 7000. This command will be automatically generated to the configuration file when at least one interface is enabled for multicast routing. This command is necessary so misconfigured “ip dvmrp metric” commands don't cause massive route injection into the MBONE. The “no” version of the command configures no limit. [11.0]

```
[no] ip dvmrp distance <admin-distance>
```

Configures the default administrative distance for received DVMRP routes. This command should be used so routes advertised from the unicast routing table that are reflected back through DVMRP cause the original unicast routes to continue to be advertised. The “ip dvmrp accept-filter” command may override this value when specified on an interface. [11.2]

```
[no] ip mroute <source> <mask> [<protocol><as-number>] [route-map <map>]
      <rpf-address> | <interface> [<distance>]
```

Configures a multicast static route (called a “static mroute”). When a source range is specified, the mroute applies only to those sources. When <protocol><as-number> is specified, the mroute applies to those sources that have been learned by the corresponding routing process. If route-map <map> is specified, further classification can be accomplished by the match clauses from <map>. If the mroute is selected, the <rpf-address> address dictates the incoming interface for the source that matches the mroute. If the <rpf-address> is a PIM neighbor, PIM Joins, Grafts, and Prunes will be sent to it. The <rpf-address> can be a host address of a directly connected router or a route. When it is a route, a recursive lookup is done from the unicast routing table to find a directly connected neighbor. If <rpf-address> is not specified, <interface> is used as the incoming interface. <distance> is used to decide if a unicast route, a DVMRP route, or a static mroute should be used for the RPF lookup. The lower distances have better preference. If the static mroute has the same distance as the other two RPF sources, the static mroute will take precedence. There are only two exceptions to this rule, directly connected routes and the default unicast route. Default <distance> is 0. Static mroutes are local to the router and are not redistributed by any dynamic routing protocol. [11.0]

```
[no] ip sdr cache-timeout <minutes>
```

The amount of time an sdr cache entry stays active in the cache. A value a 0 indicates the entry will never timeout. The default value is 24 hours. [11.2]

## Interface Subcommands

```
[no] ip igmp join-group <group-address>
```

Informs the router to join group <group-address> on the interface. IP packets that are addressed to this group address will be passed up to the IP client process in the router as well forwarded out the interface. If you do not want packets forwarded out the interface, join the group on a loopback interface. Packets are not sent on the loopback interface. [10.2]

```
[no] ip igmp static-group <group-address>
```

Configures a static connected member for the group <group-address> on the interface. Packets to the group will get fastswitched, unlike the “ip igmp join-group” command. If “ip igmp join-group” command is used for the same group address, it will take precedence and the group will behave like a locally joined group. [11.2]

```
[no] ip igmp query-interval <time-in-seconds>
```

Configures the frequency of IGMP Host-Query packets transmitted. A designated router for a LAN is the only router that transmits queries. For IGMPv1, the designated router is elected according to the multicast routing protocol that runs on the LAN. For IGMPv2, the designated querier is the lowest IP addressed multicast router on the subnet. The default value is 60 seconds. [10.2]

```
[no] ip igmp last-member-query-interval <interval>
```

Configures the “last member query interval” to be <interval>, in milliseconds. The default value is 1000 ms, or 1 second. A value below 1 second can result in faster IGMP leave actions. [12.0]

```
[no] ip igmp access-group <access-list>
```

Configures what groups are allowed on the interface. The default is all groups are allowed. [10.2]

```
[no] ip igmp version 2 | 1
```

Interface subcommand to change IGMP version. Default is version 2. [11.1]

```
[no] ip igmp query-timeout [timeout value in secs]
```

This command is valid when IGMP v2 is running. This command specifies the timeout for the router to take over as the querier for the interface, after the previous querier has stopped querying. The default value is 2 \* query-interval. If the router hears no queries for the “timeout” period, it becomes the querier. [11.1]

```
[no] ip igmp query-max-response-time [secs]
```

This command is valid when IGMP v2 is running. This command specifies the maximum query response time advertised in the IGMP queries. Default value is 10 secs. Configuring a value less than 10 seconds enables the router to prune groups faster. [11.1]

[no] ip igmp helper-address <ip-address>

This command causes all IGMP host report and leave message received on the interface to be forwarded toward the given ip-address. The reports are resent out the next-hop interface towards the ip-address with that interface's source address. This command can enable a sort of "dense-mode" join, allowing stub sites not participating in PIM to indicate membership in multicast groups. [11.3]

[no] ip cgmp

Enables the Cisco Group Management Protocol (CGMP) for IP multicast on a LAN. The command triggers a CGMP Join message. This should only be enabled on 802 media (i.e., Ethernet, Fddi, and Token Ring) or ATM. When a "no" is issued, a triggered CGMP Leave message is sent for the router's MAC address on the interface for group 0000.0000.0000. CGMP can only run on an interface if PIM is configured on the same interface. [11.1]

A Cisco will send CGMP Join messages in response to receiving IGMP reports from multicast capable members. Only the IGMP querier Cisco router sends these CGMP Join messages on behalf of hosts.

[no] ip cgmp proxy

Enables CGMP for IP multicast as well as a proxy function. Initially supported is DVMRP proxying. If a DVMRP Report is received from a router that is not a PIM router, a Cisco IGMP querier will advertise the MAC address of the DVMRP router in a CGMP Join with group address 0000.0000.0000. [11.1]

To perform CGMP proxy, a Cisco must be the IGMP querier. An IGMPv2 querier is selected based on the lowest IP addressed router on the interface. An IGMPv1 querier is selected based on the multicast routing protocol used on the interface.

When multiple Cisco routers are connected to a switched network and "ip cgmp [proxy]" is needed, it is recommended that all of them are configured 1) with the same CGMP option and 2) to have precedence of becoming IGMP querier over non-Cisco routers.

[no] ip pim version [1 | 2]

Configures PIM version for this interface. Default is 2. [11.3T]

[no] ip pim [dense-mode | sparse-mode]

Enables the PIM multicast routing protocol on the interface. Configures the interface to operate in dense or sparse mode. The default mode is dense-mode. A dense-mode interface is subject to multicast flooding by default. A sparse-mode interface is only used for multicast forwarding if a join is received from a downstream router or their are directly connected members on the interface. When "no ip pim" is entered, it disables PIM on the interface. [10.2]

[no] ip pim sparse-dense-mode

Enables the PIM multicast routing protocol on the interface. In this mode, the interface will be treated as dense-mode if the group is in dense-mode. If the group is in sparse-mode, the interface will be treated in sparse-mode. When an interface is treated in dense-mode, it will be populated in a multicast routing table's outgoing interface list when 1) there are members or DVMRP neighbors on the interface, or 2) any of the PIM neighbors on the interface have not pruned for the group. When an interface is treated in sparse-mode, it will be populated in a multicast routing table's outgoing interface list when 1) there are members or DVMRP neighbors on the interface or 2) an explicit Join has been received by a PIM neighbor on the interface. [11.1]

[no] ip pim query-interval <time-in-seconds>

Configures the frequency of PIM Router-Query messages. Router Query messages are used to elect an PIM designated router (DR). The DR is responsible for sending IGMPv1 Host-Query messages, PIM Register messages on behalf of sender hosts, and Join messages on behalf of member hosts. The default value is 30 seconds. [10.2]

[no] ip pim neighbor-filter <acl>

This command filters all PIM control messages based on the given access-list. It can be used to administratively deny a misconfigured PIM neighbor from participating in PIM, or in conjunction with "ip igmp helper-address" to be the basis for a simple stub IP multicast setup. Note: this command does not filter Auto-RP announcements and is only intended to filter neighbor-to-neighbor packets. [11.3]

**[no] ip pim multipoint-signalling**

Enables PIM to open ATM multipoint VCs for each multicast group that is joined. This command is only accepted on an ATM interface. This command allows optimal multicast trees to be built down to ATM switch granularity. This can enhance router performance and link utilization since packets are not replicated and sent multiple times over the ATM interface. The default setting is disabled. That is, all multicast traffic goes to the static map multipoint VC as long as “atm multipoint-signalling” is configured. [11.3]

**[no] ip pim vc-count <number>**

Configures the maximum number of VCs PIM opens. The default value is 200. When the router hits this maximum limit it will delete inactive VCs so it may open VCs for new groups that might have activity. [11.3]

**[no] ip pim minimum-vc-rate <pps>**

Configures the minimum traffic rate to keep VCs active. When the maximum number of VCs are opened and a new VC needs to be opened, the router will scan existing VCs. VCs that have a current 1 second rate less than or equal to <pps> are eligible for deletion. If a VC is deleted, it means that packets for its respective group do not have its own multipoint VC. However, packets will flow over a shared multipoint VC which delivers packets to all PIM neighbors. If all VCs have a 1 minute rate more than <pps>, the new group will use the shared multipoint VC. The default value is 0 packets per second. [11.3]

**[no] ip pim border**

Configures the interface to be the PIM domain border. Bootstrap messages will not be able to pass through this border in either directions. Thus effectively partitions the network into regions using different bootstrap routers. No other PIM messages are dropped by this domain border setup. Please also note that this command does not set up any multicast boundaries. [11.3T, 11.1(20)CC]

**[no] ip dvmrp metric <metric> [list <access-list>]  
{[<protocol> <process-id>] | dvmrp}**

When PIM is configured on an interface and DVMRP neighbors are discovered, the router will, by default, send DVMRP Report messages. Under certain circumstances, it may be desirable to tailor the metric used for various unicast routes. This command allows one to configure the metric associated with a set of destinations for Reports sent out this interface. The acceptable <metric> value is between 0 and 32, where 0 means that the route will not be advertised and 32 means infinity and the route is advertised unreachable. If an <access-list> is specified, only the destinations that match the access-list will be reported with the configured metric. <access-list> can be a simple or extended access-list. When extended access-lists are used, you have both address and netmask granularity. Any destinations not advertised due to split horizon will not use the configured metric.

If the <protocol><process-id> is configured, only routes learned by the specified routing protocol will be advertised in DVMRP Report messages. This parameter can be used in conjunction with <access-list> so a selective list of destinations learned from a given routing protocol may be reported. If this command is not used, only directly connected networks are advertised when DVMRP neighbors are discovered.

If the “dvmrp” keyword is configured, only routes from the DVMRP routing table will be selected to be advertised with <metric>.

This command can be used multiple times on an interface. [10.2]

**[no] ip dvmrp metric <metric> route-map <map-name>**

This command has existed since release 10.2. What has been added is the route-map keyword. Now unicast routes can be subject to route-map conditions before being injected into DVMRP. Note, route-maps cannot be used for DVMRP routes. [11.1]

**[no] ip dvmrp metric <metric> [route-map <map-name>] mbgp**

Configures redistribution of MBGP routes into DVMRP. If you supply a route-map, you can specify various match criteria options for the MBGP routes. [11.1(20)CC]

[no] ip dvmrp accept-filter <access-list> [neighbor-list <nbr-acl>]  
[<distance>]

Configures an acceptance filter for incoming DVMRP Reports. Any destinations that match <access-list> received in DVMRP reports from neighbors in the <nbr-acl> are stored in the DVMRP routing table with <distance>. The distance is used to compare with the same destination in the unicast routing table. The lower distance route (either from the unicast routing table or DVMRP routing table) will take precedence when computing the RPF interface for a source of a multicast packet. When no filters are configured on an interface, all destinations are accepted with distance configured from the "ip dvmrp distance" global command. An <access-list> value of 0, accepts all destinations. <access-list> can be a simple or extended access-list. If extended access-lists are used, you have both address and netmask granularity. An <nbr-acl> value of 0 means accept from all neighbors on interface. [10.2]

[no] ip dvmrp default-information originate | only

Indicates network 0.0.0.0 is advertised to DVMRP neighbors on this interface. By default, metric 1 will be used. This command can be used with the "ip dvmrp metric" command to tailor the metric used when advertising default. This command only takes effect when peering with mrouterd 3.4 machines. When keyword "only" is used, no other DVMRP routes are reported. When keyword "originate" is used, other more specific routes may be advertised. Do not advertise the DVMRP default route into the MBONE. [10.2]

[no] ip dvmrp metric-offset [in | out] <increment>

This is the value added to the metric of a DVMRP route advertised in a Report message. When "in" (or no keyword is supplied, the <increment> applies (is added) to incoming DVMRP reports and is reported in mroute replies. When the "out" keyword is supplied, the <increment> applies (is added) to outgoing DVMRP reports for routes from the DVMRP routing table. This is similar to the metric keyword in mrouterd configuration files. The default value for "in" is 1. The default value for "out" is 0. [11.0]

[no] ip dvmrp unicast-routing

Enables DVMRP unicast routing on the interface. This means that routes in DVMRP Report messages are cached by the Cisco in a DVMRP routing table. When PIM is running, these routes get preference over routes in the unicast routing table. This allows PIM to run on the MBONE topology, when it is different (or deviates) from the unicast topology. The default setting is off. DVMRP unicast routing can run on all interfaces including GRE tunnels. On DVMRP tunnels, it runs by virtue of doing DVMRP multicast routing. This command does not enable DVMRP multicast routing among Ciscos. However, if there is a DVMRP capable multicast router, the Cisco will do PIM/DVMRP multicast routing interaction. [10.3]

[no] ip dvmrp reject-non-pruners

This command will cause the router not to peer with a DVMRP neighbor if the neighbor doesn't support DVMRP Pruning/Grafting. If a DVMRP Probe or Report message is received without the Prune-Capable flag set, a syslog message will be logged and the message will be discarded. The default setting of this command is that all DVMRP neighbors will be accepted, regardless of capability (or lack thereof). This command only prevents peering with neighbors. If there are any non-pruning routers multiple hops away (downstream towards potential receivers) that are not rejected, then a non-pruning DVMRP network may still exist. [11.0]

[no] ip dvmrp summary-address <address> <mask> metric <value>

Configures a summary address to be advertised out the interface. If there is at least one more specific route in the unicast routing table that matches the <address>/<mask>, the summary will be advertised. Routes in the DVMRP routing table are not candidates for summarization. When the metric keyword is supplied, the summary will be advertised with metric <value>. The default metric <value> is 1. Multiple summary addresses can be configured on an interface. When multiple overlapping summary addresses are configured on an interface, the one with the longest mask takes preference. [11.2]

[no] ip dvmrp auto-summary

Enables/disables DVMRP auto-summarization. DVMRP auto-summarization occurs when a unicast subnet route is collapsed into a classful network number route. This occurs then the subnet is a different network number than the IP address of the interface (or tunnel) the advertisement is sent over. If the interface is unnumbered, the network number of the numbered interface the unnumbered interface points to is compared. The default setting for this command is enabled. [11.2]

```
[no] ip dvmrp output-report-delay <delay-time> [<burst>]
```

Configures an inter-packet delay between DVMRP reports. <delay-time>, in milliseconds, is the amount of time that elapses between transmission of a set of <burst> number packets. For example, at the periodic DVMRP report interval, if 6 packets are built, and the <delay-time> is 200 with <burst> of 2, 2 packets will be sent, then a delay of 200 milliseconds occurs, then another 2 packets are sent, then another delay of 200 milliseconds, then the final 2 packets are sent. The default value for <burst> is 2. The default value for <delay-time> is 100 milliseconds. [11.2]

```
[no] tunnel mode dvmrp
```

Configures a Cisco tunnel to encapsulate IP in IP using protocol number 4. This mode can be used when a Cisco connects to a mouted machine to run DVMRP over a tunnel. This is a popular way to connect to the MBONE. It is required to configure PIM and an IP address on a DVMRP tunnel. This mode is not used to construct a tunnel between a pair of Cisco routers. [10.2]

```
[no] ip multicast ttl-threshold <ttl-value>
```

Configures a packet TTL threshold for traffic going out the interface. Any multicast packets with a TTL less than the threshold are not forwarded out the interface. The default value is 0 which means all multicast packets are forwarded out interface. [10.2]

```
[no] ip multicast rate-limit in | out [video] | [whiteboard]
      [group-list <acl>] [source-list <acl>]
      [<kbps>]
```

Controls the rate a sender from the source-list can send to a multicast group in group-list. Any packets greater than <kbps> are silently discarded. The default value for <kbps> is 0, meaning all packets are discarded. If “in” is used, only <kbps> will be accepted on the interface. IF “out” is used, only a maximum of <kbps> kilobits per second will transmitted on the interface. If keywords “video” or “whiteboard” are used, then rate-limiting is performed based on the UDP port number used for the respective media. For this to work, “ip sdr listen” must be enabled so the port number can be obtained from the sdr cache. If sdr is not enabled, or the group address is not in the sdr cache, no rate-limiting is done for the group. The default setting of this command is disabled, meaning there is no rate-limiting in effect. [11.0]

```
[no] ip multicast boundary <acl>
```

Configures an administratively scoped boundary on the interface for multicast group addresses in the range defined by the simple IP access-list <acl>. No multicast data packets will be allowed to flow across the boundary from either directions. This allows reuse of the same multicast group address in different administrative domains. The multicast address range 239.0.0.0 to 239.255.255.255 will be designated as the administratively scoped addresses by the IANA in the near future. For example, to configure a boundary for all administratively scoped address, do:

```
access-list 1 deny 239.0.0.0 0.255.255.255
access-list 1 permit 224.0.0.0 15.255.255.255
interface ethernet 0
ip multicast boundary 1
```

If this interface is also a PIM interdomain link, the access-list that also stops unwanted Auto-RP packets would be:

```
access-list 1 deny 224.0.1.40
access-list 1 deny 224.0.1.39
access-list 1 deny 239.0.0.0 0.255.255.255
access-list 1 permit 224.0.0.0 15.255.255.255
```

If the RPF interface for a multicast route has a multicast boundary configured for that group, its outgoing interfaces will not be populated. Joins received on other interfaces will be ignored as long as the boundary remains on the RPF interface. If the RPF interface changes and the boundary no longer applies to the new RPF interface, there will be join latency introduced because of the delay in populating outgoing interfaces. [11.1]



```
[no] ip multicast helper-map {<group-address> | broadcast}
    {<broadcast-address> | <multicast-address>}
<acl>
```

When a multicast-capable cloud is between two broadcast-only clouds, this command helps to convert broadcast traffic to multicast at the first hop router, and convert it back to broadcast at the last hop router to deliver the packets to the broadcast clients. This command utilizes the multicast capability of the intermediate multicast cloud. It prevents unnecessary replication at the intermediate routers and can take advantage of multicast fast switching in the multicast cloud. <acl> is an extended access-list, when specified you can selectively configure what broadcast packets are translated based on the UDP port number. [11.1]

```
[no] ip mroute-cache [distributed]
```

Configures IP multicast fast-switching. If fast-switching is disabled on an incoming interface for a multicast routing table entry, the packet will be sent at process level for all interfaces in the outgoing interface list. If fast-switching is disabled on an outgoing interface for a multicast routing table entry, the packet is process level switched for that interface but may be fast-switched for other interfaces in the outgoing interface list. The default setting is all interfaces are multicast fastswitched. [11.0]

The “distributed” keyword will enable the interface to perform distributed fast switching on incoming packets. This command applies to the configuration of the physical interface and not to subinterfaces. Once this command is configured on the interface, all packets coming in this interface will be distributed switched. The command “ip route-cache distributed” should be configured before this command is configured on a 7500 (and not on the 12000). [11.1(20)CC]

```
[no] ip sdr listen
```

Enables/disables Session Directory Protocol V2 (sdr) from accepting packets on the interface. The router joins group 224.2.127.254 on the interface. [11.1]

Session Directory (sdr) is a multicast application for setting up conferencing sessions. It uses the protocols SDP/SAP to disseminate the names and properties of conferencing sessions over the well known session directory groups 224.2.127.254 for global scope sessions and 239.255.255.255 for administrative scope session. The session properties include contact information, session lifetime and the media being used in the session (audio, video, whiteboard and others) with their specific attributes like ttl-scope, group address and UDP port number.

When `ip sdr listen` is configured on an interface, the software will join the well known directory groups on that interface to receive and store session announcements. The announcements can be displayed with the `show ip sdr` command. IOS uses stored announcements for the `ip multicast rate-limit` command. Use the `ip sdr cache-timeout` command to configure the period of time after which received announcements are expired.

When the system has no `ip multicast routing` configured, announcements are only stored if they are received on an interfaces with `ip sdr listen` configured. When the system is configured as a multicast router, it is sufficient to configure `ip sdr listen` on just a single multicast enabled interface. The well known session directory groups are handled as local joined groups after the first `ip sdr listen` is configured (see “L—local” flag in `show ip mroute`). This causes announcements received from all multicast enabled interfaces to be routed within the system and stored.

The following router example enables a router to listen to session directory advertisements:

```
ip routing
interface loopback 0
ip address 10.0.0.51 255.255.255.0
ip pim sparse-dense mode
ip sdr listen
```

**[no] ip pim nbma-mode**

This configures a multi-access WAN interface to be in Non-Broadcast Multi-Access mode. In this case, each PIM Join message is kept track of in the outgoing interface list of a multicast routing table entry. Therefore, only PIM WAN neighbors that have joined for the group will get packets sent to as data link unicasts. This command should only be used when “ip pim sparse-mode” is configured on the interface. This is not recommended for LANs that have natural multicast capabilities. The default setting is NBMA mode disabled. [11.0]

**[no] ip multicast use-functional**

Enables the use of the RFC 1469 mandated MAC-level address, 0xc000.0004.0000, for transmission and reception of IP Multicast traffic. Valid only on Token Ring interfaces. [11.1]

**[no] ip multicast tagswitch**

Enables IP multicast tagswitching. Multicast tagswitching will stay enabled on an interface provided all neighboring routers 1) support multicast tagswitching and 2) advertise the same tag table size and router count estimate. When a member is directly connected on an attached interface, packets for that group will not be tagswitched and a single IP multicast packet is delivered on the interface. Multicast tagswitching can be active if multicast fast switching is disabled, however that router cannot tag encapsulate packets but only tag relabel packets. [11.3T]

## Debug Commands

**debug ip packet**

Logs IP packets received and transmitted. This generates lots of messages. Use with caution. [Since the beginning of day]

**debug ip mpacket [detail] [<acl>] [<group-name-or-address>]**

Logs only IP multicast packets received and transmitted. This generates lots of messages. Use with caution. An optional group address is allowed to monitor a single group's packet activity. If <acl> is specified, only multicast packets from sources described by the access-list will be logged. When “detail” is used, the packet's IP header information as well as MAC address information will be display. [10.2]

**debug ip igmp <group>**

Logs IGMP packets received and transmitted as well as IGMP host related events. Specifically IGMP protocol messages and mtrace messages. [10.2]

When <group> is specified, logging for the single group is performed. [11.3T]

**debug ip cgmp**

Logs CGMP packet/event activity. [11.1]

**debug ip dvmrp [detail [<acl>] [in|out]] [pruning]**

Logs DVMRP packets received and transmitted. Keyword “detail” will log packet contents. If <acl> is specified, only routes associated with the access-list are logged. Keywords “in” and “out” can be used to look at routes received in DVMRP reports or routes transmitted in DVMRP reports, respectively. If “pruning” is specified, only DVMRP pruning and grafting events are logged. [10.2]

**debug ip pim [<group-name-or-address>]**

Logs PIM packet received and transmitted as well as PIM related events. An optional group address is allowed to monitor a single group's PIM activity. [10.2]

**debug ip pim auto-rp**

Displays Auto-RP packet activity. [11.1]

**debug ip pim atm**

Logs PIM ATM signaling activity. [11.3]

**debug ip mrouting [<group-name-or-address>]**

Logs transaction events to and from the IP multicast routing table. An optional group address is allowed to monitor a single group's routing table activity. [10.2]



**debug ip mcache**

Enables debug logging for IP multicast fast-switching events. [11.0]

**debug ip sdr**

Enables logging of received sd and sdr announcements. [11.1]

**debug ip pim tag**

Logs PIM activity related to multicast tag allocation and distribution. [11.3T]

**debug ip mtag**

Logs multicast tag allocation events as well as insertion and deletions to the TIB. [11.3T]

**debug ip mds ipc [event | packet]**

Logs MDS IPC related activity. This command is used both on the RP and line-card consoles. [11.1(20)CC]

**debug ip mds process**

Logs MDS RP and line card events. This command is used both on the RP and line-card consoles. [11.1(20)CC]

**debug ip mds mevent**

Logs multicast FIB related events on the line card. Used to debug MFIB route creation, deletion, and update events. [11.1(20)CC]

**debug ip mds mpacket**

Logs MDS switching events. Used to debug packet drops, interface drops, and switching failures. [11.1(20)CC]

**debug ip mbgp updates**

Logs MBGP related information passed in BGP Update messages. [11.1(20)CC]

**debug ip mbgp dampening [<acl>]**

Logs route flap dampening activity. If <acl> is specified, logging occurs for the routes specified in the access-list only. [11.1(20)CC]

## Show Commands

**show ip igmp interface [<interface-unit>]**

Displays learned groups for given interface. [10.2]

**show ip igmp groups [<group-name> | <group-address>] | [<interface-unit>]**

Displays directly connected groups learned via IGMP. [10.2]

**show ip pim neighbor [<interface-unit>]**

Displays PIM neighbors discovered. [10.2]

**show ip pim rp [<group-name> | <group-address>] [mapping]**

Displays active RPs that are cached with associated multicast routing entries. If “mapping” is supplied, displays all group-to-RP mappings that are configured and/or learned from Auto-RP. [10.2]

**show ip pim interface [<interface-unit>] [count]**

Displays PIM interface information. [10.2]

**show ip pim vc [<group-or-name>] [<interface>]**

Displays ATM VC status information for multipoint VCs opened by PIM. When <group-or-name> is specified, only the single group is displayed. When <interface> is specified, only the single ATM interface is displayed. [11.3]

**show ip pim bsr**

Displays the bootstrap router (BSR) information, including the elected BSR's IP address, priority, hash mask length etc. Information is shown if the local router is a candidate bootstrap router (BSR) but not an elected BSR. This command also displays information about locally configured candidate RP advertisement. [11.3T]

**show ip pim rp-hash <group> ,**

shows which RP is being selected for the <group>. It also shows whether this RP was learned via Auto-RP or v2 bootstrap mechanism. [11.3T]

```
show ip mroute [[<group-name> | <group-address>] [<source-address>]]
[summary]
```

Displays the IP multicast routing table. When “summary” is specified, a one line abbreviated display is provided. When “count” is specified, group count, source count, and packet count statistics are provided. [10.2]

```
show ip mroute [[<group-name-or-address>] [<source-name-or-address>]]
count
```

Displays packet count per (S,G) multicast routing table entry. Also includes average packet size and data rate in kilobits per second. [10.2]

```
show ip mroute [<group-name-or-address>] active [<kbps>]
```

This command shows you the rate active sources are sending to multicast groups. You can display for all groups or specify a single <group>. <kbps> will only display sources that are sending >= <kbps>. The default setting shows all sources sending at a rate greater than or equal to 1 kilobit/second. If sd is running on the router, the sd session name is displayed. [11.0]

```
show ip mroute [<group-name-or-address>] [<source-name-or-address>] pruned
```

Displays (S,G) entries that have been pruned. [11.3]

```
show ip mroute [<group-name-or-address>] [<source-name-or-address>] [<interface>]
```

Display the multicast routing table entries that have <interface> in the outgoing interface list. [11.3]

```
show ip mroute [<group-name-or-address>] active [<kbps>] [<interface>]"
```

Display the active multicast routing table entries that have <interface> in the outgoing interface list. [11.3]

```
show ip mroute static
```

Displays statically configured multicast static routes. [11.2]

```
show ip mpacket [detail]
```

```
show ip mpacket <source-address-or-name> [detail]
```

```
show ip mpacket <group-address> [detail]
```

```
show ip mpacket <source-address-or-name> <group-address> [detail]
```

```
show ip mpacket [read-only]
```

Displays the contents of the circular cache-header buffer. Each time this command is entered, a new buffer is allocated. The summary display shows IP packet identifier, ttl, source and destination IP addresses plus a local timestamp when the packet was received. The detail display shows the rest of the IP header fields on an additional line plus the first 8 bytes after the IP header (usually the UDP port numbers). This command is only applicable when the “ip multicast cache-headers” command is in effect. [11.1]

When keyword “read-only” is specified, a new buffer is not allocated. The old buffer continues to accumulate packet entries. Useful for repeating the display with more or less detail the second time. [11.1(22)CC, 12.0S]

```
show ip mpacket <group-address> [<source-address-or-name>] quality
```

Displays RTP loss statistics for each given source sending to group <group-address> or the single source specified in <source-address-or-name>. Packets received, packets lost, packets misordered, average loss gap, and loss percentage is provided.

[11.1(20)CC]

```
show ip dvmrp route [<name-or-address> | <interface>]
```

Displays the DVMRP routing table. The DVMRP routing table contains unicast destinations only when any DVMRP is running on an interface or tunnel. If <name-or-address> is specified, a longest match lookup is performed and the entry that matches is displayed. If <interface> is specified, all routes with a next-hop interface of <interface> is displayed. [10.2]

```
show ip dvmrp route [<name-or-address>] poison
```

Display interfaces where a DVMRP report with a poisoned-reversed metric has been received for the route. [11.3]

**show ip mcache [<group> [<source>]]**

Displays the IP fast-switching cache. If <group> is specified, the fast switching cache for the single group is displayed. <group> can be either a Class D IP address or a DNS name. If <source> is specified with <group>, a single multicast cache entry is displayed. <source> can be either a unicast IP address or a DNS name. [11.0]

Displays “mds” if MDS is being used instead of the fastswitching.

Displays the “last used time” if the fast switching is being used.

Displays “never” if the fast switching entry is not used (process level switched).

Note that if MDS is not enabled on an incoming interface which is capable of doing MDS, incoming packets wont be distributed fastswitched and will be fastswitched at the RP. Also, if the incoming interface is not MDS capable, the packet will get fastswitched or process-level switched at the RP. If the incoming interface is enabled for MDS but at least one of the outgoing interfaces is not fastswitch capable, packets will be process-level switched. So it is a good idea not to disable fast switching on any interface when MDS is enabled. [11.1(20)CC]

**show ip sdr [<group> | “<session-name>” | detail]**

Displays the sdr cache. If the router is configured to be a member of 224.2.127.254 (the default sd group), it will cache sdr announcements. If no parameters are used, a sorted list of session names are displayed. If <group> is specified, the session(s) defining the multicast group will be displayed in detail format. If “<session-name>” is specified, the single session is displayed in detail format. If the keyword “detail” is specified, all sessions are display detail format. [11.1]

**show ip rpf <source-address-or-name>**

Displays how IP multicast routing does Reverse Path Forwarding. Since the Cisco can RPF from multiple routing tables (i.e., unicast routing table, DVMRP routing table, or static mroutes), this command tells you where the information is retrieved from. [11.0]

**show tech-support ipmulticast**

Provides output for a set of IP multicast related show commands. This is useful for a customer to use a single type-in to provide the necessary information so Cisco can debug an IP multicast related problem. [11.2]

**show ip pim neighbor tag [<interface>]**

Displays per neighbor multicast tag related information. [11.3T]

**show ip pim interface tag [<interface>]**

Displays multicast tag information per interface. [11.3T]

**show ip mtag [<group>] [<source>] [<interface>]**

Displays the TIB for a particular multicast route or for an entire interface. If no parameters are supplied, the display will show tag information for all multicast routes. [11.3T]

**show ip mds stats [switching | linecard]**

Displays distributed fast switching stats and linecard stats. [11.1(20)CC]

**show ip mds interface**

Displays MDS related interface information. [11.1(20)CC]

**show ip mds forwarding [<group>] [<source>]**

Displays the MFIB table on the line-card. Displays forwarding information and related flags and counts for MDS. This command is used on the line cards via the line-card console. [11.1(20)CC]

**show ip mds summary**

Displays MDS summary information on the line-card. This command is used on the line cards via the line-card console. [11.1(20)CC]

**show ip mbgp <command>**

<command> can be any of the existing BGP commands supported by “show ip bgp <command>.” The MBGP variants show multicast RIB related information. Use this command in conjunction with “show ip rpf” to determine if IP multicast routing is using MBGP routes. [11.1(20)CC]

## Clear Commands

```
clear ip mroute [<group-name> | <group-address> [<source-address>]] | [*]
```

Deletes entries from the IP multicast routing table. [10.2]

```
clear ip igmp group [<group-name> | <group-address>] | [<interface-unit>]
```

Deletes entries from the IGMP cache. [10.2]

```
clear ip cgmp [<interface>]
```

Sends a CGMP Leave message with a group address of 0000.0000.0000 and a unicast address of 0000.0000.0000. This instructs the switches to clear all group entries they have cached. If <interface> is specified, the Leave is sent only on <interface>. Otherwise, it is sent on all CGMP enabled interfaces. [11.1]

```
clear ip dvmrp route * | <route>
```

Deletes routes from the DVMRP routing table. When "\*" is used, all routes are deleted. When <route> is used, the longest match route will be deleted. [10.2]

```
clear ip sdr [<group-address> | "<session-name>"]
```

Deletes an sdr cache entry. No parameters delete the entire sdr cache. If <group-address> is supplied, all sessions associated with the IP group address are deleted. If "<session-name>" is specified, only the sdr cache entry with the supplied name is deleted. [11.1]

```
clear ip pim interface [<interface>] count
```

Clears the multicast packet counters for interface <interface>. Or clears for all interfaces when <interface> is not specified. [11.2]

```
clear ip pim auto-rp <rp-address>
```

Clears the Auto-RP cache. When <rp-address> is specified, only the entries related to RP <rp-address> are cleared. Otherwise, the entire Auto-RP cache is cleared. [11.2]

```
clear ip mtag [<group>] [<source>] [<interface>]
```

Clears the multicast Tag Information Base (TIB). If <interface> is specified, all TIB entries for the interface are cleared. If <group> is specified, all (S,G)s for the group will have their respective TIB entries cleared. If both <group> and <source> are specified, only the TIB for the single (S,G) will be cleared. [11.3T]

```
clear ip mds forwarding *
```

Clears this linecards MFIB table and resync with the RP. This command is used on the line cards via the line-card console. [11.1(20)CC]

## Exec Commands

```
mrinfo [<hostname-or-address>] [<source-address-or-interface>]
```

mrinfo is the Mbone's original tool to determine what neighboring multicast routers are peering with a multicast router. Cisco has supported responding to mrinfo requests since release 10.2. Now you can query a multicast router using the mrinfo exec command. If <hostname-or-address> is not used, the router queries itself. Otherwise, the DNS name or IP address of the multicast router is queried. The output format is identical to the mroute version. If <source-address-or-interface> is specified, that is the source address used on mrinfo requests. When it is omitted, the source address is based on the outbound interface for the destination. [11.0]

```
mtrace <source> [<destination>] [<group>]
```

mtrace is a multicast traceroute. It allows you to trace the path from a source to a destination branch for a multicast distribution tree for a given group. The trace traverses in the reverse direction from destination to source. This allows you to isolate multicast routing failures. When <group> is not specified, 224.2.0.1 is used. This is the group used for "Mbone Audio." When <destination> is not specified, the mtrace starts from the system that the command is typed from. <source> is required and can be a DNS name or the IP address of a multicast capable source. <destination> and <group> can also take DNS names. When no parameters are entered, the router will interactively ask you for the parameters. [11.0]

**mstat <source> [<destination>] [<group>]**

The form of Unix mtrace that reports packet rate/loss information. Takes the same input as mtrace. [11.0]

**ping <group-address-or-name>**

You can send ICMP Echo Requests to a multicast group address. When you supply <group-address-or-name>, a single Echo Request packet is sent. When you type ping with no arguments, you will be prompted. [10.2]

**ping tag <ipmulticast-address>**

Allows a router running multicast tagswitching to send tag encapsulated IP packets to the group <ipmulticast-address>. If a tag is not known for the group, no packet is sent. This command accepts all the same options as the traditional IP ping command. [11.3T]

## MBGP Router Subcommands

**[no] neighbor <address> remote-as <asn> [nlri unicast | multicast]**

Configures a BGP peer and associated AS number. If the “multicast” keyword only is supplied, only multicast NLRI will be sent to the neighbor. Conversely, if the “unicast” keyword only is supplied, only unicast NLRI will be sent to the neighbor. Both keywords may be supplied which indicates the neighbor will be sent both types of routes. Unicast NLRI will be sent in the conventional encoding and the multicast NLRI will be sent in the MP\_REACH and MP\_UNREACH path attribute. The default is to send unicast NLRI only. This version of BGP will negotiate NLRI in the Capabilities Option of the Open message. Therefore, both sides of a BGP connection must be configured consistently with respect to NLRI or the TCP connection will not be established.

**[no] neighbor <tag> peer-group [nlri unicast | multicast]**

Configures the peer-group to support either unicast NLRI, multicast NLRI or both. Supplying the both keywords “unicast” and “multicast” indicates both NLRI are sent. The default value is unicast only.

**[no] neighbor <address-or-tag> description <user-supplied>**

Configures a user supplied description string for the neighbor.

**[no] neighbor <address-or-tag> dont-capability-negotiate**

This hidden command configures whether Capability Parameter negotiation in BGP Open messages are used. By default, capability negotiation is performed. This command is used to interoperate with older MBGP implementations in which no Capability Parameters are used in Open messages between multicast peers. This command is not applicable when peering between two unicast-only BGP neighbors. This command will eventually go away.

**[no] distance mbgp <dist1> <dist2> <dist3>**

For configuring administrative distance for MBGP routes. <dist1> is the distance for external routes, <dist2> is the distance for internal routes, and <dist3> is the distance for local routes. The default values are the same as for BGP.

**[no] network <address> [mask <mask>] [nlri unicast | multicast]**

Determines if network in the AS should be injected into the BGP unicast RIB or the MBGP multicast RIB. If both keywords “unicast” and “multicast” are specified, it is injected in both RIBs. If “multicast” is specified, it is injected in the multicast RIB only. The default is unicast only.

**[no] aggregate-address <address> <mask> [nlri unicast | multicast]**

This command, used with the keyword “multicast” indicates if the aggregate should be applied to the multicast RIB. If the keyword is not supplied, then it applies to the unicast RIB. If you want it to apply to both RIBs, enter the both keywords “unicast” and “multicast.”

**[no] redistribute dvmrp [route-map <map-name>]**

Configures redistribution of DVMRP routes into MBGP. If you supply a route-map, you can specify various BGP attribute settings.

## MBGP Route-Map Subcommands

```
[no] match nlri unicast | multicast
```

The route-map criteria can be based on the unicast or multicast RIB (or both). If the multicast RIB entry is being processed for a route-map with a "match nlri multicast," then the route-map condition will yield true, likewise for the unicast corollary. If both keywords "unicast" and "multicast" are specified, then either RIB entry being processed will yield TRUE. The default value is both.

This command can be used in conjunction with the "neighbor <address> route-map <map> in" command so you can use one route-map reference to describe filtering policies for different NLRI types.

```
[no] set nlri unicast | multicast
```

If the route-map match criteria is met, decide if the route should be injected into the unicast or multicast RIB. If both keywords "unicast" and "multicast" are specified, the route is injected into both RIBs and advertised as separate NLRI in a BGP Update message. If "multicast" is specified, the route is injected only into the multicast RIB. The default value is unicast only in all cases except when this route-map is referenced by a "neighbor <foo> route-map <foo> out" command. This route-map subcommand is used when referencing a route-map by various router subcommands (i.e., like "redistribute," "aggregate-address," and neighbor outbound route-map references).

This command can be used in conjunction with the "neighbor <address> default-originate route-map <map>." If the "set nlri" is supplied in the route-map referenced by the neighbor command, the multicast default route can be generated independent of the unicast default route.



### 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 s.a.r.l.  
Parc Evolic, Batiment L1/L2  
16 Avenue du Quebec  
Villebon, BP 706  
91961 Courtaboeuf Cedex  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 69 18 61 00  
Fax: 33 1 69 28 83 26

### 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/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

Copyright © 1999 Cisco Systems, Inc. All rights reserved. Printed in the USA. Access Registrar AccessPath, Any to Any, AtmDirector, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCSI, CD-PAC, the Cisco logo, Cisco Certified Internetwork Expert logo, CiscoLink, the Cisco Management Connection logo, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco Systems Capital, the Cisco Systems Capital logo, Cisco Systems Networking Academy, the Cisco Technologies logo, ConnectWay, ControlStream, Fast Step, FireRunner, GigaStack, IGX, JumpStart, Kernel Proxy, MGX, Natural Network Viewer, NetSoner, Network Registrar, Packet, PIX, Point and Click Internetworking, Policy Builder, Precept, RouteStream, Secure Script, ServiceWay, SlideCast, SMARTnet, StreamView, The Cell, TrafficDirector, TransPath, ViewRunner, VirtualStream, VisionWay, VlanDirector, Workgroup Director, and Workgroup Stack are trademarks; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, The Internet Economy, and The New Internet Economy are service marks; and ASIST, BPX, Catalyst, Cisco, Cisco IOS, the Cisco IOS logo, Cisco Systems, the Cisco Systems logo, the Cisco Systems Cisco Press logo, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, FastSwitch, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Registrar, StrataView Plus, Stratm, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any of its resellers. (9907R)