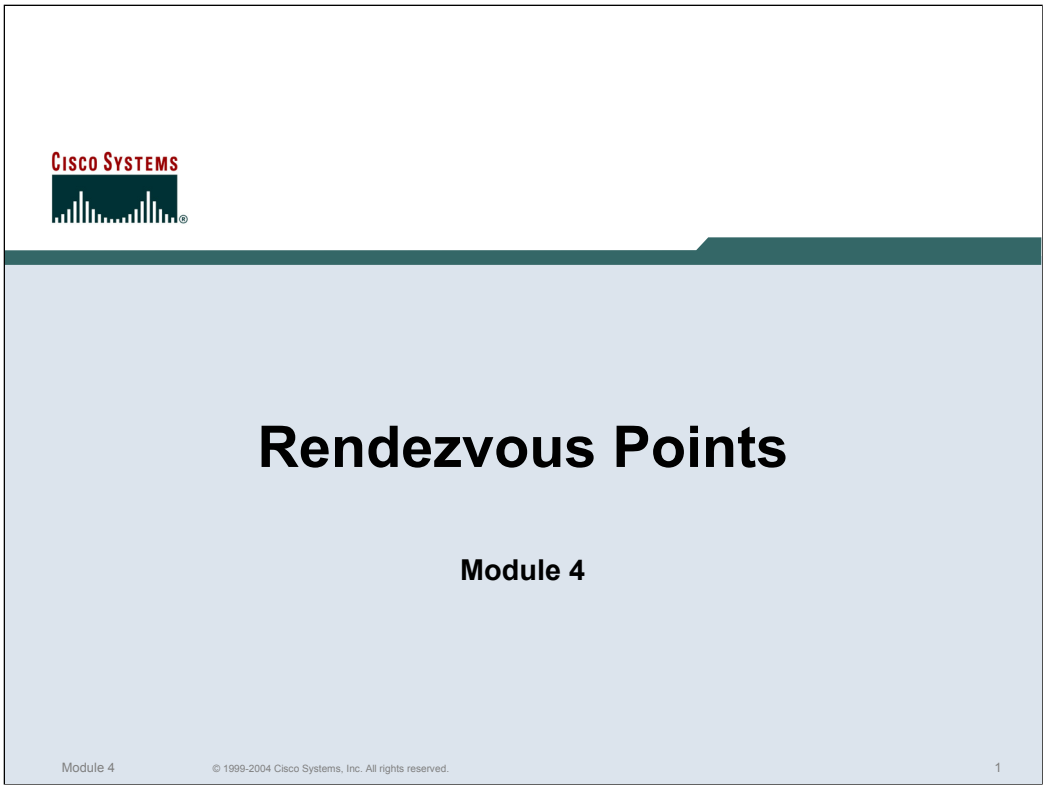
Rendezvous Points
Module 4
Module 4
© 1999-2004 Cisco Systems, Inc. All rights reserved.
1

Module Agenda

Cisco.com

- Auto RP
- PIMv2 BSR
- Static RP's
- Anycast RP's
- Tuning RP Operations

Auto-RP

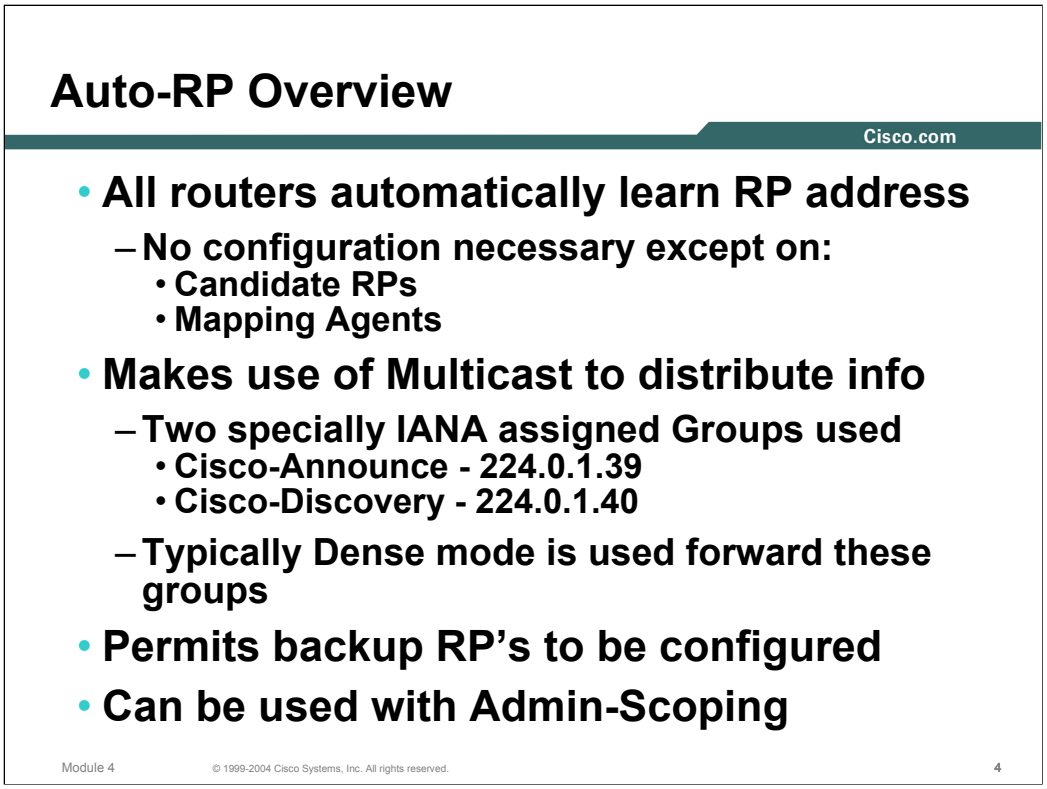

- **Auto-RP Overview**
  - Auto-RP allows all routers in the network to automatically "learn" Group-to-RP mappings.
  - There are no special configuration steps that must be taken except on the router(s) that are to function as:
    - Candidate RP's
    - Mapping Agents
  - Multicast is used to distribute Group-to-RP mapping information via two special, IANA assigned multicast groups.
    - Cisco-Announce Group - 224.0.1.39
    - Cisco-Discovery Group - 224.0.1.40
  - Because multicast is used to distribute this information, a "Chicken and Egg" situation can occur if the above groups operate in Sparse mode. (Routers would have to know a priori what the address of the RP is before they can learn the address of the RP(s) via Auto-RP messages.) Therefore, it is recommend that these groups *always* run in Dense mode so that this information is flooded throughout the network.
  - Multiple Candidate RP's may be defined so that in the case of an RP failure, the other Candidate RP can assume the responsibility of RP.
  - Auto-RP can be configured to support Administratively Scoped zones. (BSR cannot!) This can be important when trying to prevent high-rate group traffic from leaving a campus and consuming too much bandwidth on WAN links.

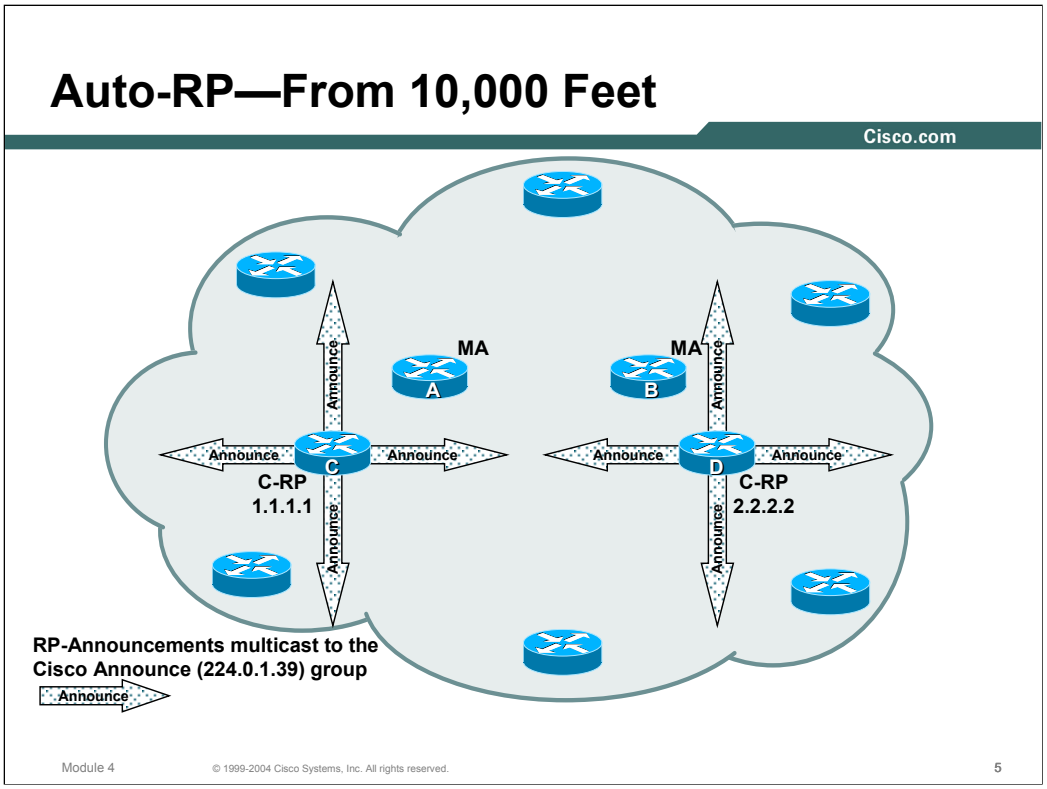

- **Auto-RP - The big picture**

  In this example, routers A and B have been configured as Mapping Agents while routers C & D have been configured as Candidate RP's.

  – Step 1

    - The Candidate RP's begin multicasting their candidacy to be the RP via RP-Announce messages which are sent via the Cisco-Announce group, 224.0.1.39.

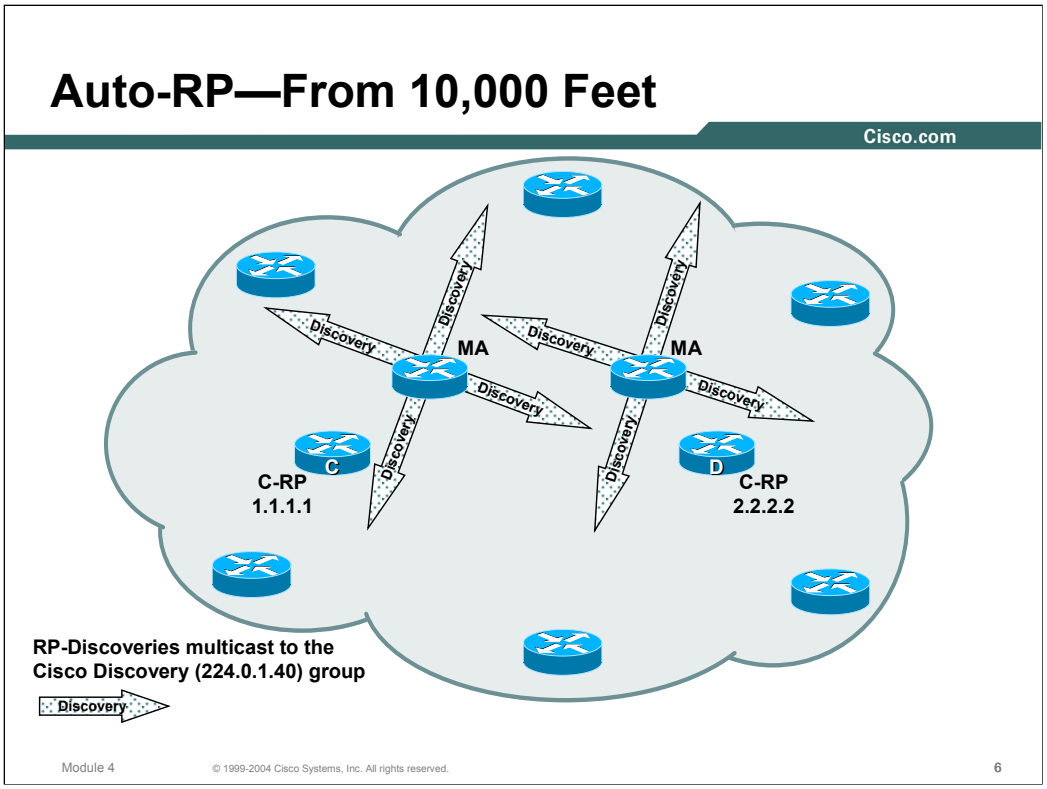

- **Auto-RP - The big picture**
  - Step 2
    - The two Mapping Agents (routers A & B) receive the RP-Announce messages from the two Candidate RP's (routers C & D).
  - Step 3
    - The C-RP with the highest IP address (in this case, router D) is stored in the Group-to-Mapping cache of the Mapping Agents.
  - Step 4
    - The Mapping Agents *both* multicast the contents of their Group-to-RP Mapping Cache to the Cisco-Discovery group, 224.0.1.40.

      Note: All Mapping Agents are transmitting this Group-to-RP Mapping information simultaneously. The originally published specification on Auto-RP implied that there was a Master-Slave relationship between Mapping Agents and that only the Master would transmit while the Slave(s) were quiet until the Master failed. This specification is in error and this is not how Auto-RP has been implemented. As long as both Mapping Agents are transmitting identical information, there is no need to add the complexity of a Master-Slave failover scheme.
  - Step 5
    - The RP Discovery messages are received via multicast by all routers in the network. The Group-to-RP mapping information contained in these messages is stored in the router's local Group-to-RP mapping cache. This information is subsequently used by the router to determine the IP address of the RP for a given group.

Auto-RP Fundamentals

Cisco.com

- **Candidate RPs**
  - **Multicast RP-Announcement messages**
    - **Sent to Cisco-Announce (224.0.1.39) group**
    - **Sent every rp-announce-interval (default: 60 sec)**
  - **RP-Announcements contain:**
    - **Group Range (default = 224.0.0.0/4)**
    - **Candidate's RP address**
    - **Holdtime = 3 x <rp-announce-interval>**
  - **Configured via global config command**
    `ip pim send-rp-announce <intfc> scope <ttl> [group-list acl]`
  - **'Deny' in group-list has variable meaning**
    - **Before 12.0(1.1) Deny = "I'm not C-RP for this group-range"**
    - **After 12.0(1.1) Deny = "Force group-range to always be DM"**

Module 4 7

- **Auto-RP Candidate RP's (C-RP's)**
  - Multicast RP-Announcement messages to the Cisco-Announce (224.0.1.39) group. These messages "announce" this router as being a Candidate for selection as RP and are sent every 60 seconds by default.
  - RP-Announce messages contain:
    - Group Range (default is all multicast groups or 224.0.0.0/4)
    - The Candidate's IP address
    - A holdtime which is used to detect when the C-RP has failed. This holdtime is 3 times the announcement interval or 3x60 = 180 seconds = 3 minutes
  - C-RP's are configured using the (rather obtuse) command:

    ip pim send-rp-announce <intfc> scope <ttl> [group-list <acl>]
    - The <intfc> specifies which IP address is used as the source address in the RP-Announce messages that are sent out all multicast interfaces on the router.
    - The <ttl> value controls the TTL of the RP-Announce message.
    - The optional 'group-list' permits a group range other than the default to be assigned.
    - This command may be configured more than once on a router so that the router will function as C-RP for multiple group ranges.
  - Note: A 'deny' in the 'group-list' access-list has a different meaning beginning with IOS release 12.0(1.1).
    - Before 12.0.(1.1): Deny means "I'm not the RP for this group range."
    - After 12.0.(1.1): Deny means "Force this group range to always work in Dense mode. Note: Only a single C-RP needs to "deny" this group range to force this to happen. In other words, the 'deny' overrides any other router's "permit" advertisement.

Auto-RP Fundamentals

Cisco.com

- **Mapping agents**
  - **Receive RP-Announcements**
    - **Stored in Group-to-RP Mapping Cache with holdtimes**
    - **Elects highest C-RP IP address as RP for group range**
  - **Multicast RP-Discovery messages**
    - **Sent to Cisco-Discovery (224.0.1.40) group**
    - **Sent every 60 seconds or when changes detected**
  - **RP-Discovery messages contain:**
    - **Elected RP's from MA's Group-to-RP Mapping Cache**
  - **Configured via global config command**

    `ip pim send-rp-discovery [<interface>] scope <ttl>`

    - **Source address of packets set by '<interface>' (12.0)**
      - **If not specified, source address = output interface address**
      - **Results in the appearance of multiple MA's. (one/interface)**

Module 4 8

- **Auto-RP Mapping Agents (MA's)**
  - Mapping Agents join the RP-Announce group (224.0.1.39) in order to receive RP Announcements sent by all Candidate RP's.
  - When they receive an Announcement they:
    - Save the Announcement in the Group-to-RP mapping cache
    - Select the C-RP with the highest IP address as RP for the group range
    - The holdtimes are used to timeout an entry in the cache if a C-RP fails and is no longer sending periodic C-RP announcements.
  - Mapping Agents periodically send the elected RP's from their Group-to-RP mapping cache to all routers in the network via RP Discovery messages.
    - RP Discovery messages are multicast to the Auto-RP Discovery group 224.0.1.40.
    - They are sent every 60 seconds or when a change to the information in the Group-to-Mapping takes place.
  - MA's are configured using the (rather obtuse) command:

    ip pim send-rp-discovery[<intfc>] scope <ttl>

    - The optional <intfc> specifies which IP address is used as the source address in the RP-Discovery messages that are sent out all multicast interfaces on the router. (A Loopback interface is normally specified here.) If this interface is not specified, the source address of each multicast interface on the router is used.

      Note: The reason that this is an optional clause is strictly to be backwards compatible with IOS releases prior to 12.0 that did not allow the interface to be specified. In practice, an interface should always be specified.
    - The <ttl> value controls the TTL of the RP-Discovery message.

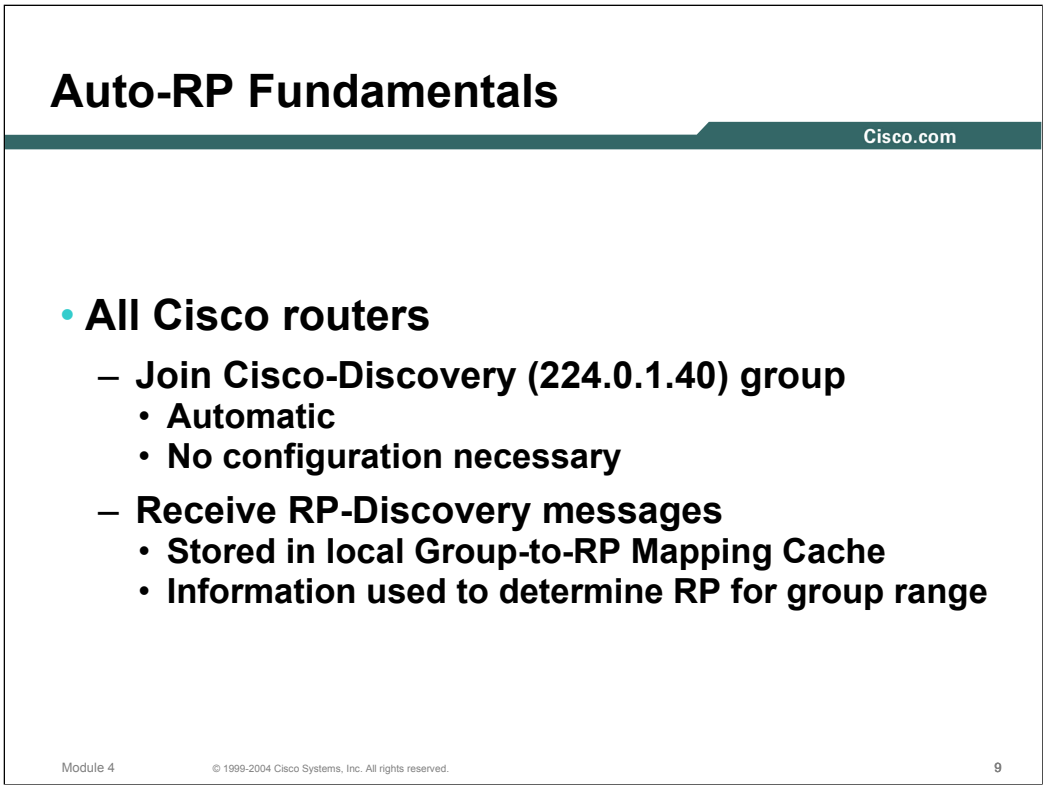

- **All Cisco Routers**
  - Automatically join the Cisco-Discovery (224.0.1.40) group in order to receive Group-to-RP mapping information being multicast by the Mapping Agents in the network.
    - *No configuration is necessary!*
  - Group-to-RP mapping information contained in the RP-Discovery messages is stored in the router's local Group-to-RP mapping cache. This information is used by the router to map a Group address to the IP address of the active RP for the group.

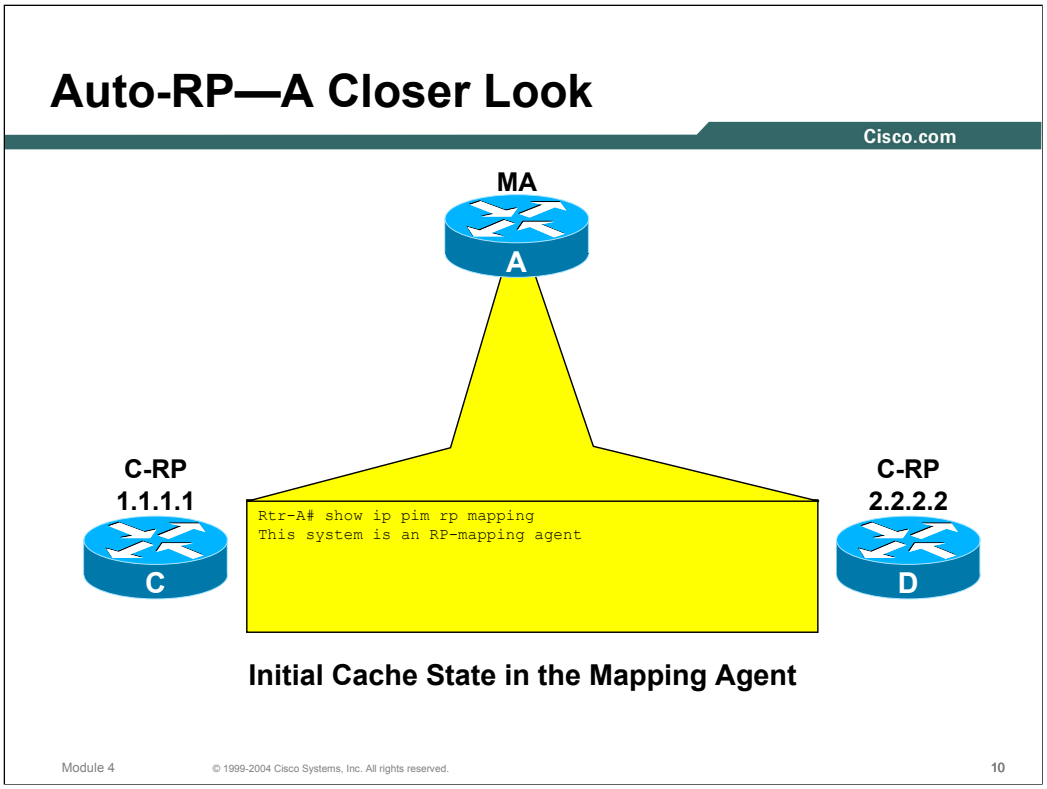

- **Auto-RP Up Close**

  This is the same example that was presented in the previous slides. However, in this case, we will examine the process in more detail at each step.

  – Step 1
    - At time zero, the Group-to-RP mapping caches in the Mapping Agents are empty since no RP-Announcements have been received.
    - The output of the 'show ip pim rp mapping' command shows that router A is a Mapping Agent and that the Group-to-RP mapping cache is empty.

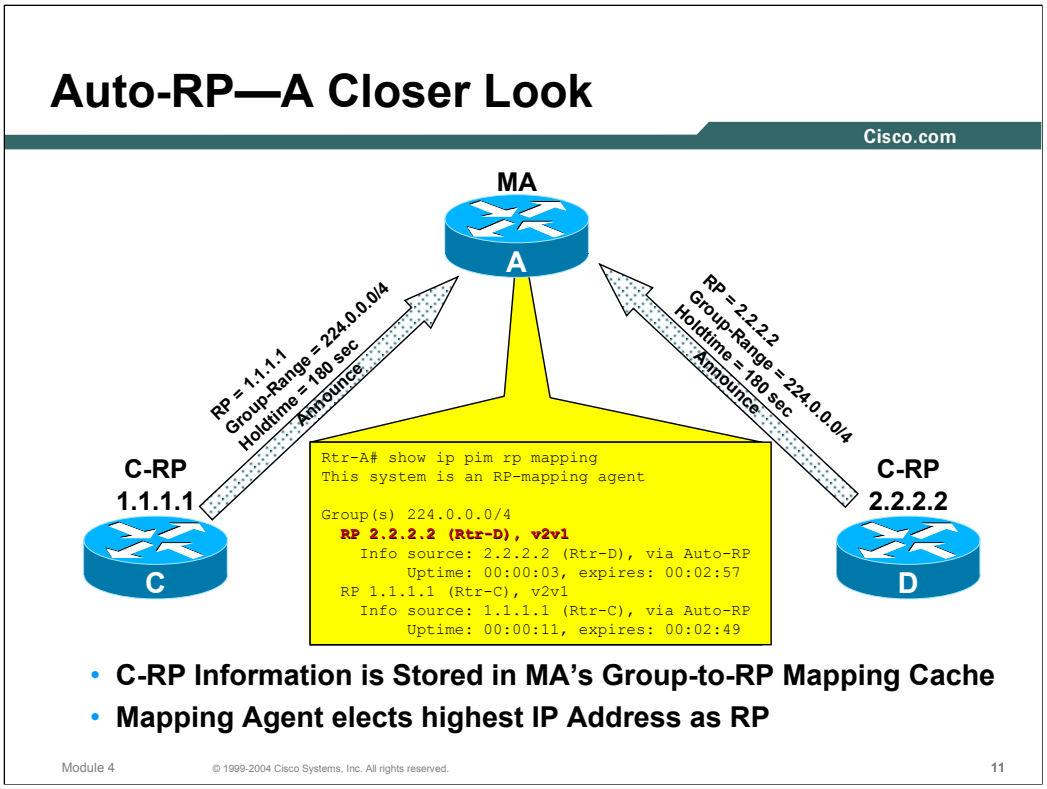

- **Auto-RP Up Close**
  - Step 2
    - Routers C and D begin sending their RP Announce messages advertising themselves as a candidate to be RP for all multicast groups. (Note the group range, the IP address of the C-RP and the holdtime in the message.)
  - Step 3
    - The Mapping Agent (router A) receives these RP Announcements and stores this information in its Group-to-RP mapping cache.
    - The output of the 'show ip pim rp mapping' command on the Mapping Agent (router A) now shows both router C and D as candidates for group range 224.0.0.0/4 (i.e. all multicast groups with the exception of the Auto-RP groups).
    - The Mapping Agent then elects the C-RP with the highest IP address as the active RP for the group range.

Auto-RP—A Closer Look

Cisco.com

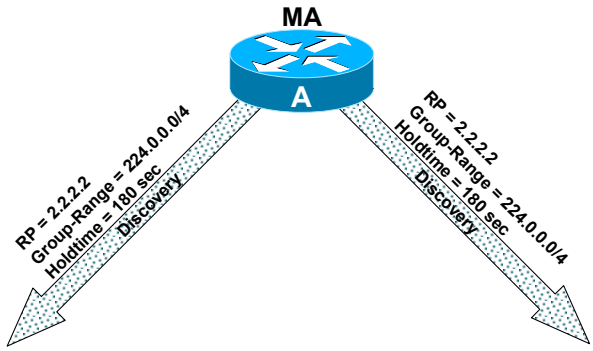

- **Mapping Agent advertises elected RP via Discovery messages**

Module 4 12

- **Auto-RP Up Close**
  - Step 4
    - The Mapping Agent begins advertising the results of the RP election to the rest of the network via Auto-RP Discovery messages.

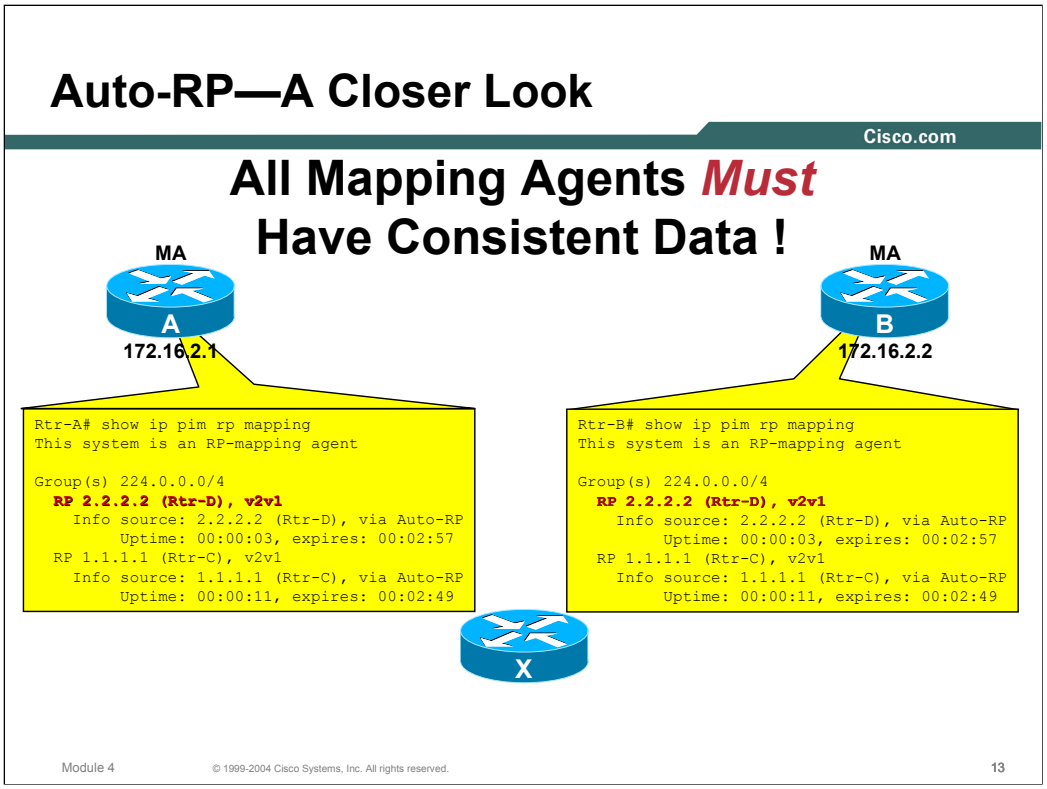

- **Auto-RP Up Close**

  It is critical that all Mapping Agents in the PIM-SM domain have identical information in their Group-to-RP mapping caches. Note that in our example network, they do.

  If the information in the mapping caches are *not* identical, it can cause the routers in the network to flip-flop between two different RPs.

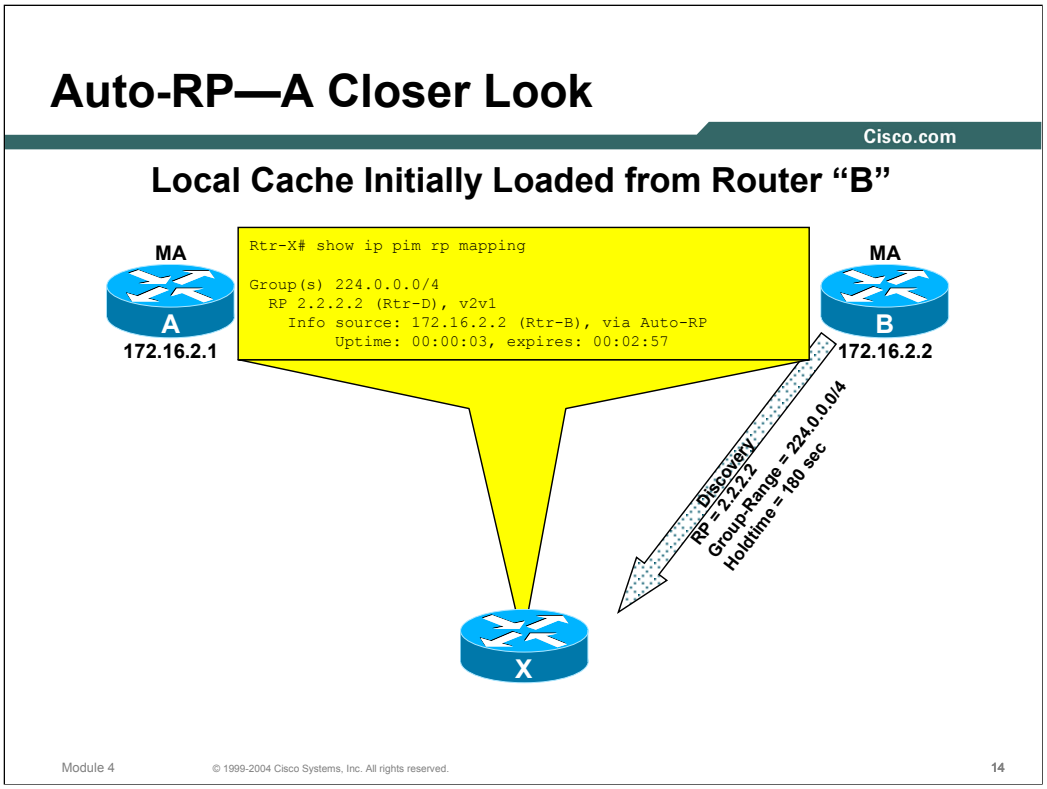

- **Auto-RP Up Close**
  - Step 6
    - Assume that router B is the first MA to send its RP Discovery message containing its Group-to-RP mapping cache contents.
  - Step 7
    - The routers in the network (router X in this example) all receive this RP Discovery message and install the information in their local Group-to-RP mapping cache.
    - The output of the 'show ip pim rp mapping' command shows that router D is currently selected as the RP for group range 224.0.0.0/4 (i.e. all multicast groups with the exception of the Auto-RP groups) and that this information was most recently received from router B.

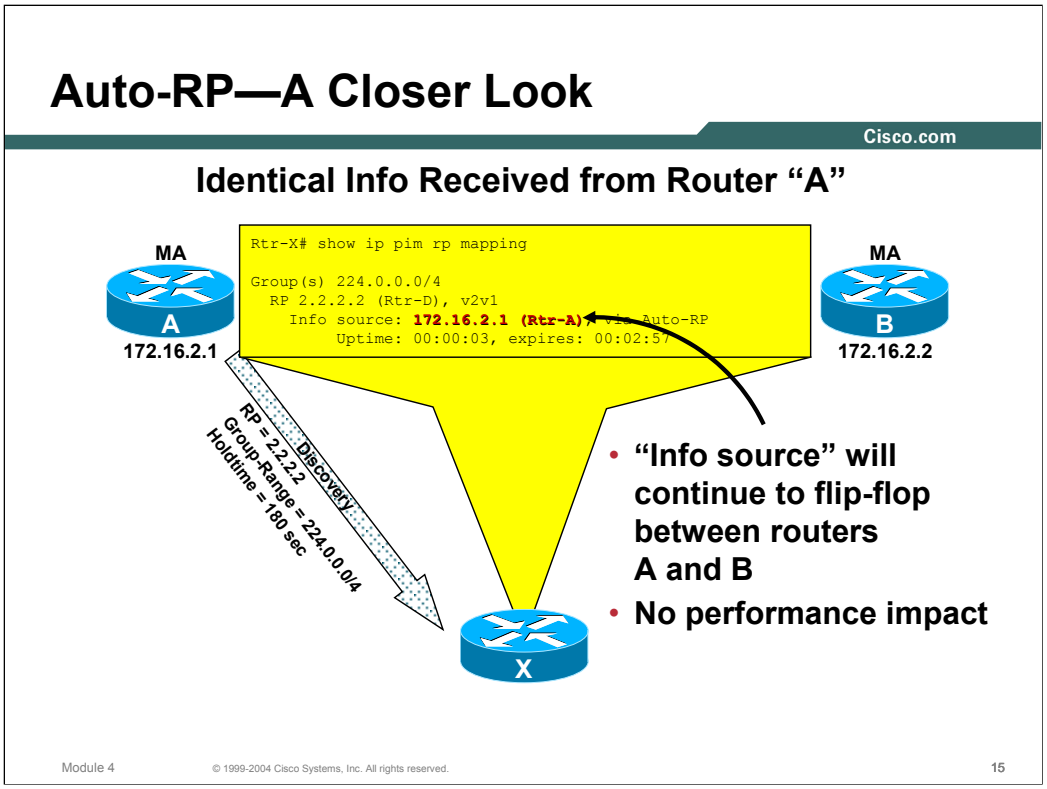

- **Auto-RP Up Close**
  - Step 8
    - Next, router A sends an RP Discovery message containing its Group-to-RP mapping cache contents.
  - Step 9
    - The routers in the network (router X in this example) all receive this RP Discovery message and update the information in their local Group-to-RP mapping cache. Since both Mapping Agents are sending identical information, the only thing that will change in the local Group-to-RP mapping cache is the "source" of the information.
    - The output of the 'show ip pim rp mapping' command shows that router D is still selected as the RP for group range 224.0.0.0/4 (i.e. all multicast groups with the exception of the Auto-RP groups). However, the data reflects that this information was most recently received from router A.
    - The flip-flop of the information source in the routers' local Group-to-RP mapping cache has little or no performance impact on the router.

PIMv2 BSR

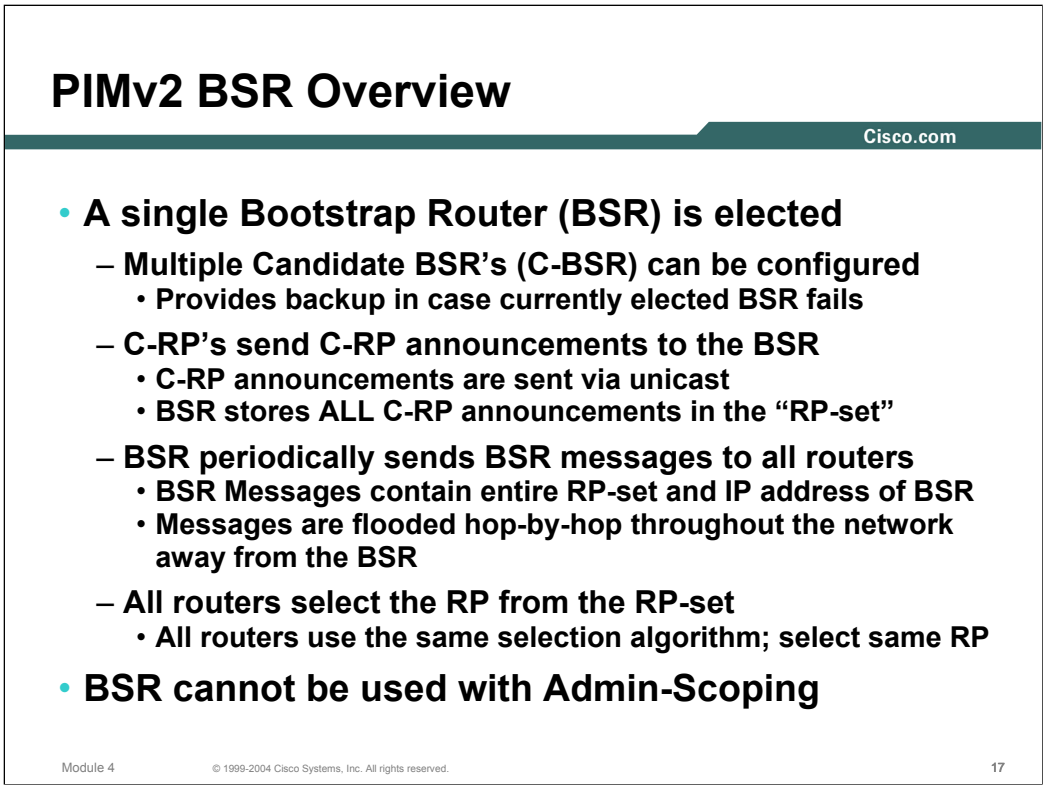

- **BSR Overview**
  - Bootstrap Router (BSR)
    - A single router is elected as the BSR from a collection of Candidate BSR's.
    - If the current BSR fails, a new election is triggered.
    - The election mechanism is pre-emptive based on C-BSR priority.
  - Candidate RP's (C-RP's)
    - Send C-RP announcements directly to the BSR via unicast. (Note: C-RP's learn the IP address of the BSR via periodic BSR messages.)
    - The BSR stores the complete collection of all received C-RP announcements in a database called the "RP-set".
  - The BSR periodically sends out BSR messages to all routers in the network to let them know the BSR is still alive.
  - BSR messages are flooded hop-by-hop throughout the network.
    - Multicast to the "All-PIM Routers" group (224.0.0.13) with a TTL of 1.
  - BSR messages also contain:
    - The complete "RP-set" consisting of *all* C-RP announcements.
    - The IP Address of the BSR so that C-RP's know where to send their announcements.
  - All routers receive the BSR messages being flooded throughout the network.
    - Select the active RP for each group range using a common hash algorithm that is run against the RP-set. This results in all routers in the network selecting the same RP for a given group-range.
  - ***BSR cannot be used with Admin-Scoping!***
    - Admin scoping was not considered when BSR was designed. One problem is that C-RP announcements that are unicast to the BSR cross multicast boundaries. There are several other problems as well.

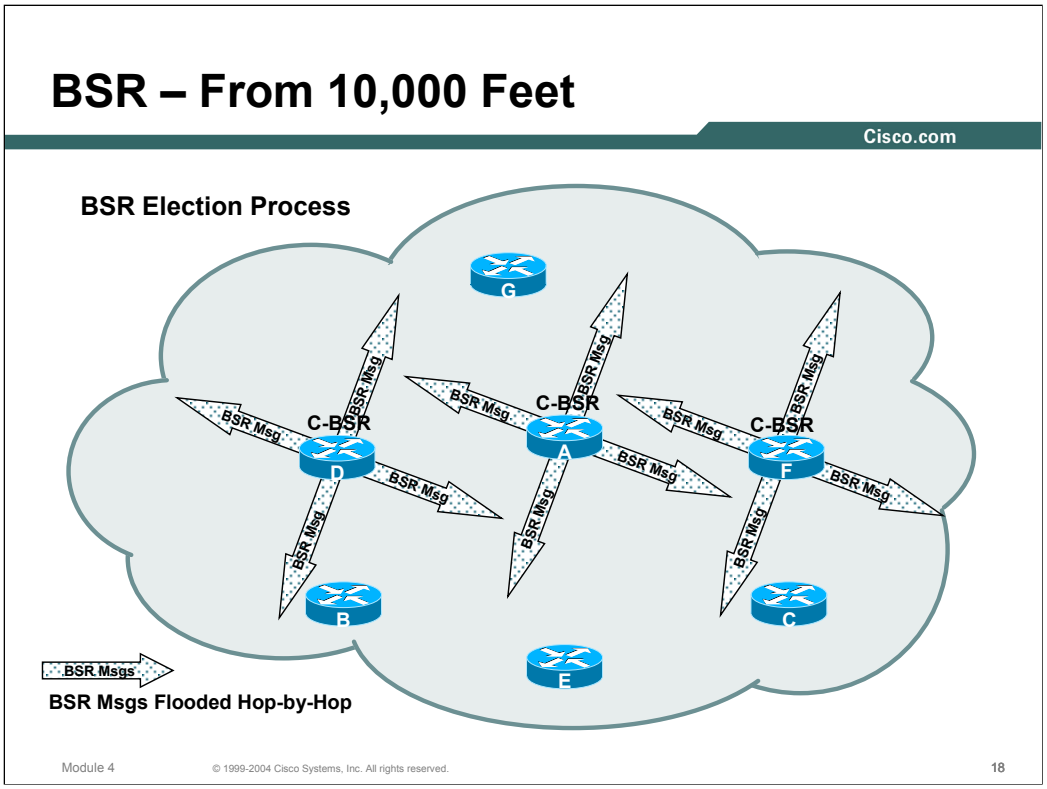

- **BSR Overview**
  - The network administrator configures one or more routers in the network to serve as Candidate BSR's (C-BSR).
    - At network startup, all Candidate BSR's participate in the BSR election process by sending a PIM BSR message containing its BSR priority out all interfaces. These BSR messages are flooded hop-by-hop throughout the entire network.

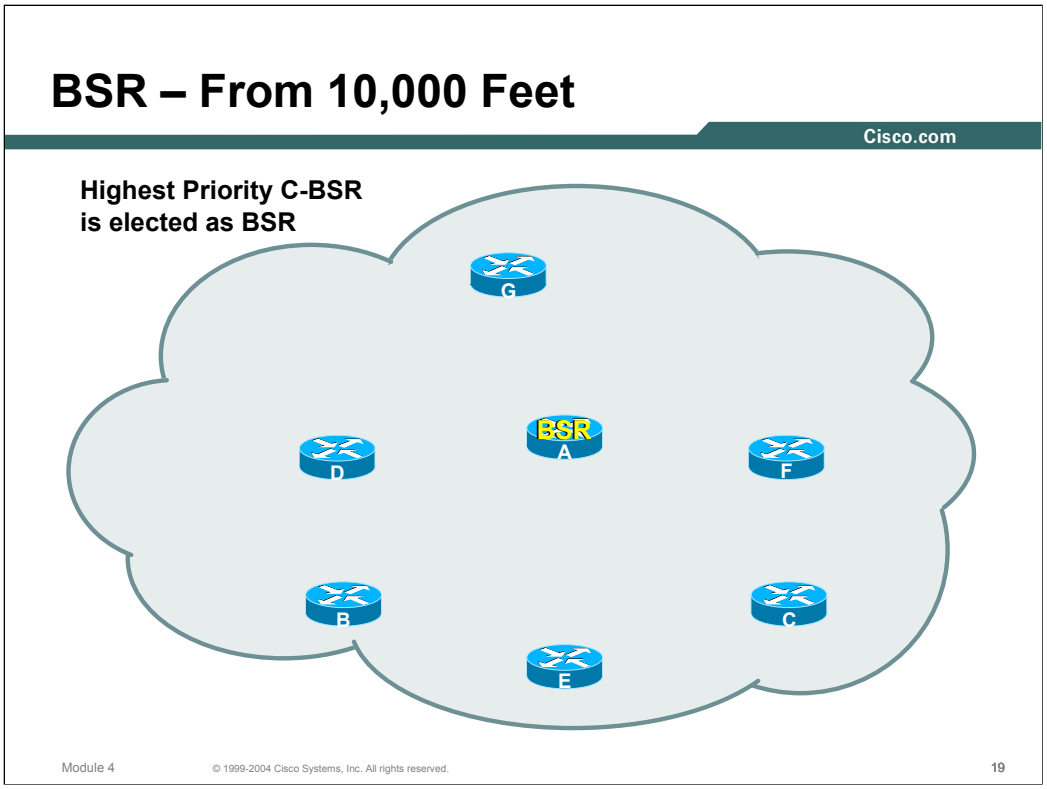

- **BSR Overview**
  - At the end of the “BSR-Election-Interval”, the BSR with the highest BSR priority is elected as the active BSR (Bootstrap Router).
    - (Note: The BSR election process is similar in nature to the Root-Bridge election mechanism in the Spanning-Tree protocol.)

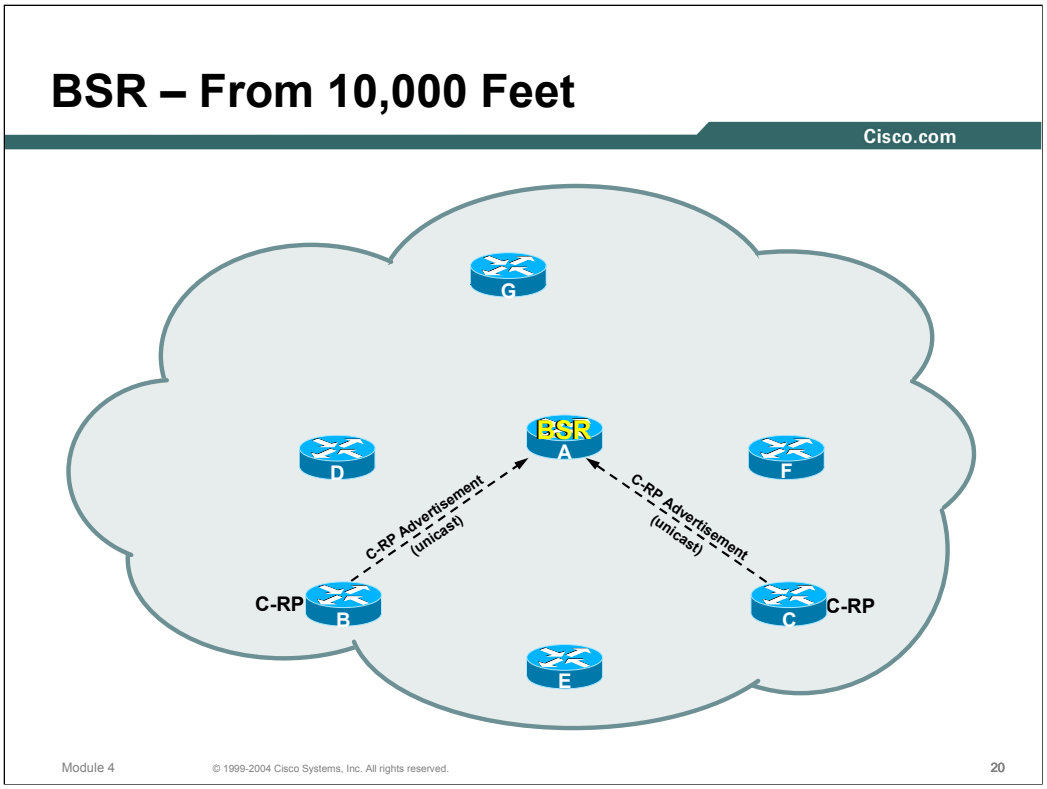

- **BSR Overview**
  - All routers in the network participate in the BSR Election process by forwarding BSR messages to their downstream neighbors. As a result, at the end of the BSR Election Interval, all routers in the network (including the routers that have been configured as Candidate RPs) know which C-BSR has been elected as the currently active BSR.
  - Since the Candidate RP's know the IP address of the currently active BSR, they can unicast their C-RP Announcement messages directly to the active BSR.

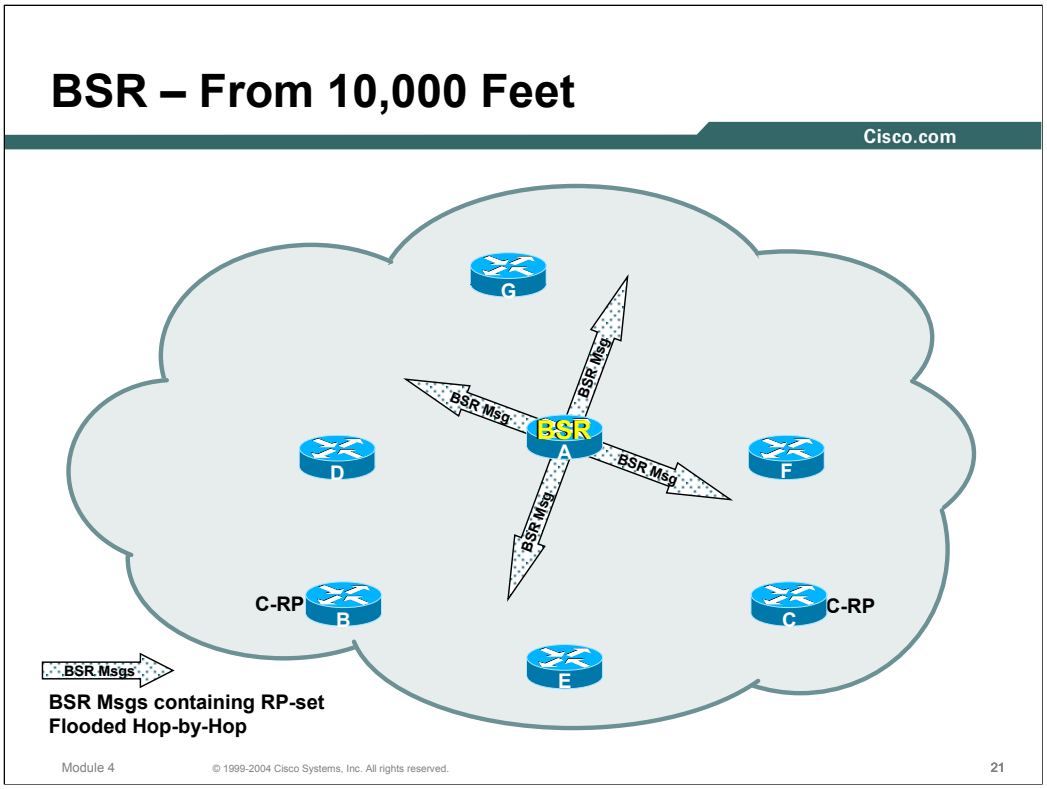

- **BSR Overview**
  - The active BSR stores all incoming C-RP Announcements in its Group-to-RP mapping cache. The BSR then sends the entire list of C-RP's from its Group-to-RP mapping cache in periodic BSR messages which are flooded hop-by-hop throughout the entire network. As each router receives a copy of these BSR messages, it updates the information in it's local Group-to-RP mapping cache so it knows the IP address of all C-RP's in the network.
  - However, unlike Auto-RP where the Mapping Agent "elects" the active RP for a group range and announces the election results to the network, the BSR does not "elect" the active RP for a group. Instead, it leaves this task to each individual router in the network.
  - Each router in the network will using a well-known hashing algorithm to elect the currently active RP for a particular group range. Since each router is running the same algorithm against the same list of C-RP's, then they will all elect the same RP for a particular group range.

PIMv2 BSR Fundamentals

Cisco.com

- **Candidate RPs**
  - **Unicast PIMv2 C-RP messages to BSR**
    - **Learns IP address of BSR from BSR messages**
    - **Sent every rp-announce-interval (default: 60 sec)**
  - **C-RP messages contain:**
    - **Group Range (default = 224.0.0.0/4)**
    - **Candidate's RP address**
    - **Holdtime = 3 x <rp-announce-interval>**
  - **Configured via global config command**

```
ip pim rp-candidate <intfc> [group-list acl]
```

Module 4 22

- **BSR Candidate RP's (C-RP)**
  - C-RP Messages
    - Sent periodically (default: 60sec) directly to the BSR via unicast.
    - Messages contain the Group-range, C-RP address and a holdtime.
    - The IP address of the current BSR is learned from the periodic BSR messages that are received by all routers in the network.
  - C-RP's are configured using the following command:

```
ip pim rp-candidate <intfc> [group-list <acl>]
```

    - The <intfc> parameter dictates the IP Address that is advertised in the C-RP message. In most cases, a Loopback interface is used.
    - The optional 'group-list' access-list can be used to specify a group-range other than the default of 224.0.0.0/4 (i.e. all multicast groups)
    - This command may be configured more than once on a router so that the router will function as C-RP for multiple group ranges.

PIMv2 BSR Fundamentals

Cisco.com

- **Bootstrap router (BSR)**
  - **Receive C-RP messages**
    - **Accepts and stores ALL C-RP messages**
    - **Stored in Group-to-RP Mapping Cache w/holdtimes**
  - **Originates BSR messages**
    - **Multicast to All-PIM-Routers (224.0.0.13) group**
      - **(Sent with a TTL = 1)**
    - **Sent out all interfaces. Propagate hop-by-hop**
    - **Sent every 60 seconds or when changes detected**
  - **BSR messages contain:**
    - **Contents of BSR's Group-to-RP Mapping Cache**
    - **IP Address of active BSR**

Module 4 23

- **Bootstrap Router**
  - The primary purpose of the Bootstrap router is to collect all C-RP announcements in to a database called the RP-set and to periodically send the RP-set out to all other routers in the network inside of BSR messages.
  - BSR Messages
    - Sent periodically (default: 60 secs) by the BSR out all multicast interfaces.
    - BSR messages are multicast to the All-PIM-Routers (224.0.0.13) group with a TTL of 1. These messages are received by all PIM neighbors who retransmit them (again with a TTL of 1) out all interfaces except the one in which the messages was received. (An RPF check is done to insure the BSR message came in on the correct interface in the direction of the BSR.)
    - BSR messages contain the RP-set and the IP address of the currently active BSR. (This is how C-RP's know where to unicast their C-RP messages.

PIMv2 BSR Fundamentals

Cisco.com

- **Candidate bootstrap router (C-BSR)**
  - **C-BSR with highest priority elected BSR**
    - **C-BSR IP address used as tie-breaker**
      - » **(Highest IP address wins)**
    - **The active BSR may be preempted**
      - » **New router w/higher BSR priority forces new election**
  - **Configured via global config command**

```
ip pim bsr-candidate <intfc> <hash-length> [priority <pri>]
```

    - **<intfc>**
      - » **Determines IP address**
    - **<hash-length>**
      - » **Sets RP selection hash mask length**
    - **<pri>**
      - » **Sets the C-BSR priority (default = 0)**

Module 4 24

- **Candidate Bootstrap Routers (C-BSRs)**
  - C-BSR's participate in the BSR election mechanism.
    - The C-BSR with the highest priority is elected as the BSR.
    - The highest IP address of the C-BSR's is used as a tie-breaker.
    - The election mechanism is preemptive. If a new C-BSR with a higher priority comes up, it triggers a new election.
  - C-BSR's are configured using the following command:
    - ip pim bsr-candidate <intfc> <hash-length> [priority <pri>]
    - The <intfc> parameter is used to specify the BSR's IP address which is forwarded in BSR messages. (This is where C-RP's will send their messages if the C-BSR is elected as BSR.)
    - The <hash-length> parameter specifies the number of bits in the hash. This can be used to control RP load balancing across a group range where different RP's are selected for different groups within a group range whose size is defined by the hash-length in bits.
    - The optional <pri> value permits the C-BSR to be configured with a priority other than the default of zero.

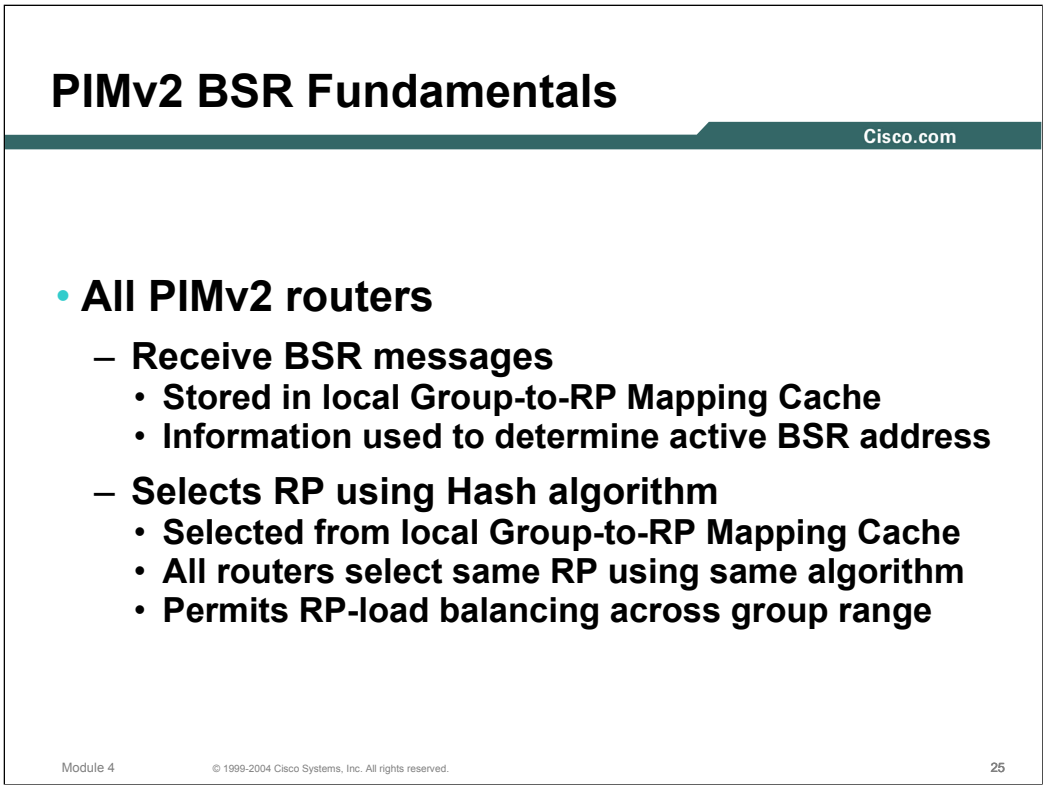

- **All PIMv2 Routers**
  - Accept BSR messages based on the rules described in the previous pages. When a BSR message is accepted:
    - The RP-Set in the BSR message is stored in the local Group-to-RP mapping cache.
    - The BSR message is forwarded out all other interfaces (except the one in which it was received) on the router.
  - Selects RP using a Hash Algorithm
    - The RP for a group is selected from the set of C-RP's (stored in the Group-to-RP mapping cache) that have advertised their candidacy for a matching group-range.
    - The same hashing algorithm is used by all routers to select the RP from the set of C-RP's in the RP-set. Since all routers run the same algorithm on the same RP-set (received from the BSR), all routers will select the same RP for a given group.
    - The hashing algorithm permits multiple C-RP's to load balance the duties of RP across a range of groups. Only one C-RP will be selected as RP for any *single* group in the group range. However, the hash algorithm may select other C-RP's as RP for another group *within the group range.*
    - For example, given a BSR hash length of 30 bits being used on IPv4 group addresses, this results in a remainder of 2 bits of an IPv4 address or 4 group addresses that a C-RP will serve as RP. In this scenario, if C-RP routers A and B both advertise their candidacy for group-range 224.1.1.0/24 and the hash algorithm selects router A as RP for 224.1.1.0, the hash length of 28 bits will also cause router A to be selected as RP for groups 224.1.1.1, 224.1.1.2 and 224.1.1.3 (i.e. a contiguous group range of of 4 addresses.) If the hash algorithm selects router B as RP for group 224.1.1.4, it will also select router B for groups 224.1.1.5, 224.1.1.6 and 224.1.1.7.

Static RP's

Static RP's

Cisco.com

- **Hard-coded RP address**
  - **When used, must be configured on every router**
  - **All routers must have the same RP address**
  - **RP fail-over not possible**
    - **Exception: If Anycast RPs are used. (See MSDP module.)**
- **Command**

```
ip pim rp-address <address> [group-list <acl>] [override]
```

  - **Optional group list specifies group range**
    - **Default: Range = 224.0.0.0/4** ***(Includes Auto-RP Groups!!!!)***
  - **Override keyword "overrides" Auto-RP information**
    - **Default: Auto-RP learned info takes precedence**

Module 4 27

- **Hard-code RP Addresses**
  - Requires every router in the network to be manually configured with the IP address of a *single* RP.
  - If this RP fails, there is no way for routers to fail-over to a standby RP.
    - The exception to this rule is if "Anycast-RP's" are in use. This requires MSDP to be running between each RP in the network.
- **Command**

```
ip pim rp-address <address> [group-list <acl>] [override]
```

  - The 'group-list' allows a group range to be specified.
    - The default is ALL multicast groups or 224.0.0.0/4
    - DANGER, WILL ROBINSON!!!

      The default range includes the Auto-RP groups (224.0.1.39 and 224.0.1.40) which will cause this router to attempt to operate these groups in Sparse mode. This is normally not desirable and can often lead to problems where some routers in the network are trying to run these groups in Dense mode (which is the normal method) while others are trying to use Sparse mode. This will result in some routers in the network being starved of Auto-RP information. This in turn, can result in members of some groups to not receive multicast traffic.
  - The 'override' keyword permits the statically defined RP address to take precedence over Auto-RP learned Group-to-RP mapping information.
    - The default is that Auto-RP learned information has precedence.

Anycast-RP's

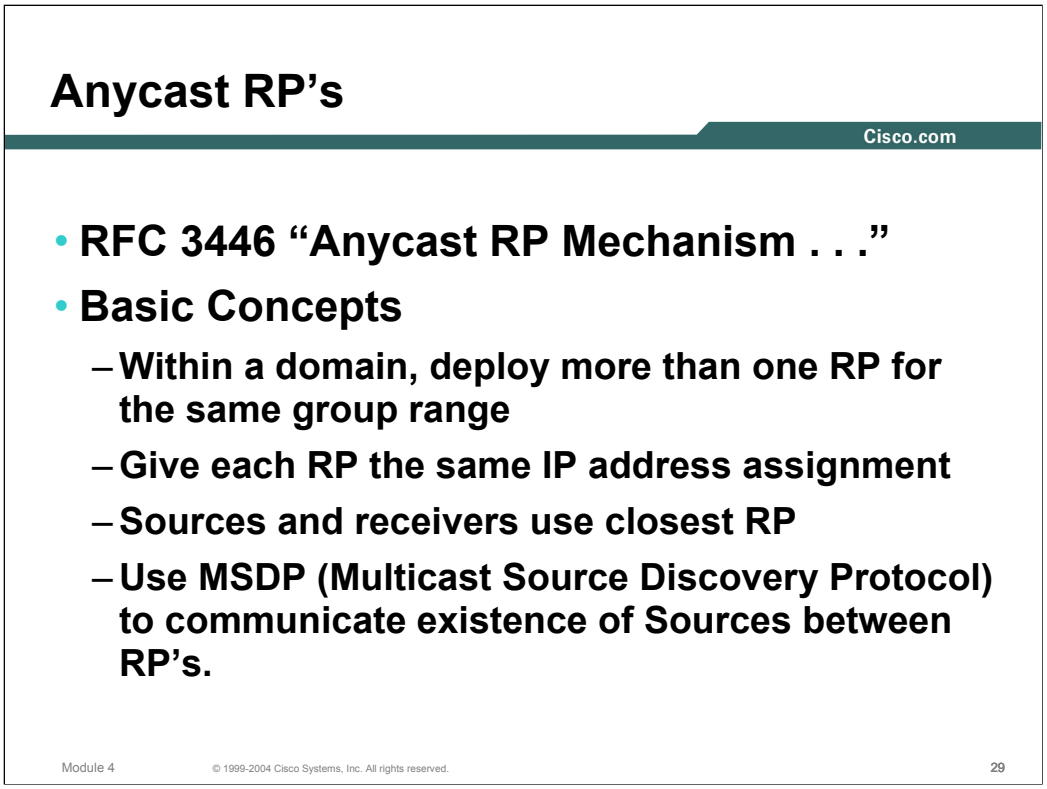

- **Anycast RP's**
  - RFC 3446, "Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", describes the methodology for configuring Anycast RP's.
  - Basic Concept
    - Within a PIM-SM domain, more than one RP's are deployed for a given group range. Each of these RP's are assigned the same IP address which is in turn, advertised as a host route by each RP. This results in the other routers in the PIM-SM domain calculating a route to the nearest RP based on the unicast routing metrics. The net result is that Sources and Receivers in one portion of the PIM-SM domain will be Register and Joined to the closest RP.
    - Normally, the above technique would result in a partitioning of the PIM-SM domain into multiple smaller PIM-SM domains centered around the Anycast RP's. However, by using the Multicast Source Discovery Protocol (MSDP) which was originally designed to interconnect PIM-SM domains together, the partitioning can be healed.

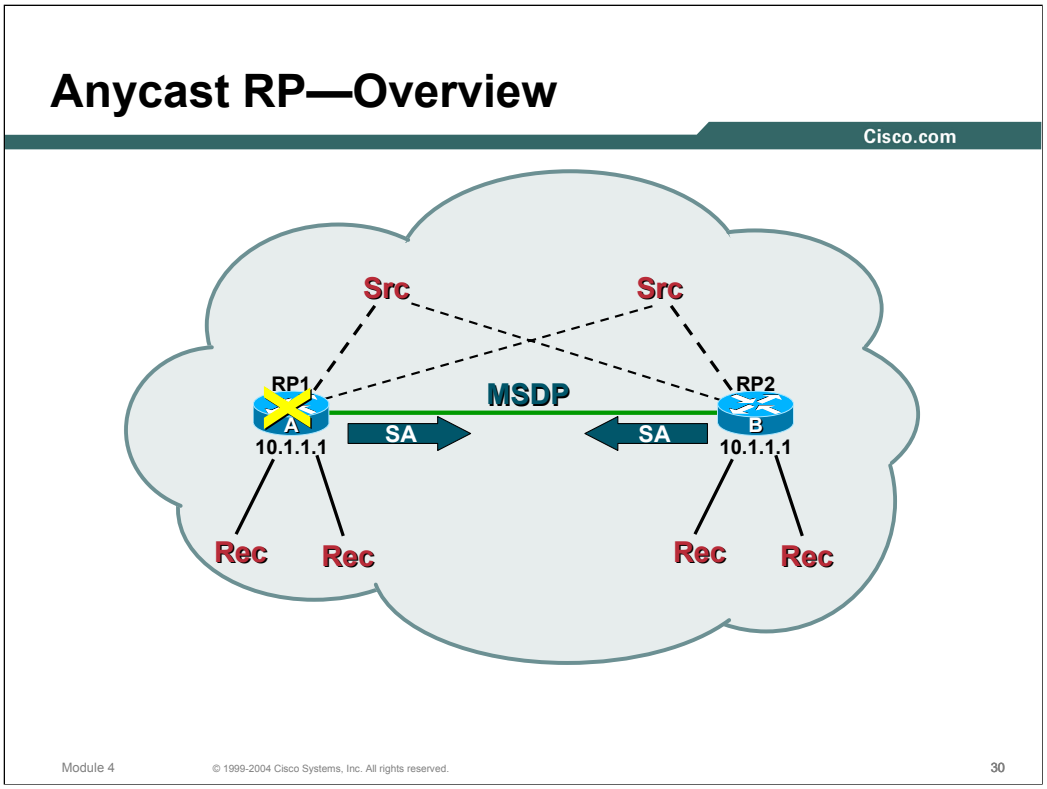

- **Anycast RP – Overview**
  - The Last-Hop routers that are directly connected to receivers will join the Shared Tree to the nearest RP.
  - The First-Hop routers directly connected to sources will register the sources to the nearest RP.
  - Whenever an RP learns of a new source in it's portion of the network, it sends an MSDP "Source Active" (SA) message to the other RP(s) in the Anycast RP cluster to notify it (them) of the new source.
  - When an SA message is received, an RP with active receivers for the group will send an (S,G) Join toward the source to "pull" the flow of source traffic to its portion of the network and will forward this traffic down its Shared Tree. At this point the normal PIM-SM mechanisms take over and (typically) the Last-Hop routers in this portion of the network will Join the SPT to the source and bypass the local Anycast RP.

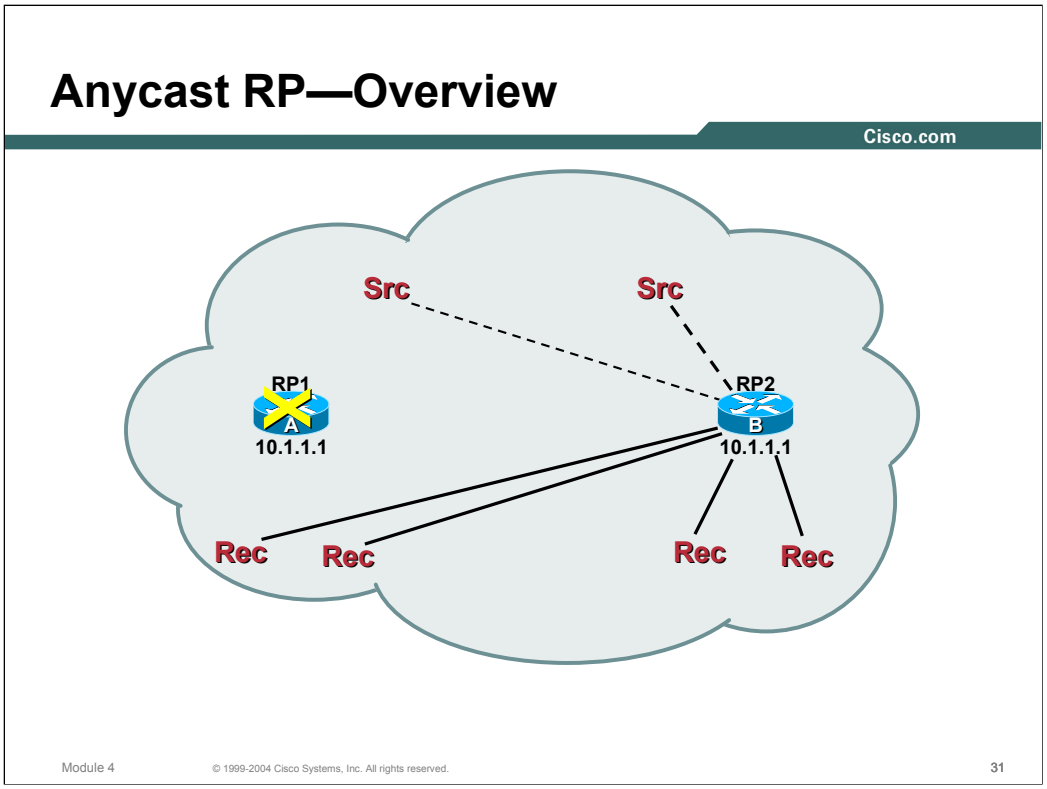

- **Anycast RP – Overview**
  - If an Anycast RP fails, it will cease advertising its address as a host route to the network. Within a few seconds, the unicast routing in the network will reconverge and only see the routes to the remaining Anycast RP(s) in the network.
  - Once the unicast routing has converged, the PIM protocol will reconverge by calculating new RPF information for every mroute entry in its multicast forwarding table. Where there has been a change in the RPF information for the Shared Tree and SPT's, new PIM Join/Prune control messages are sent out to rebuild the trees using the remaining RP.

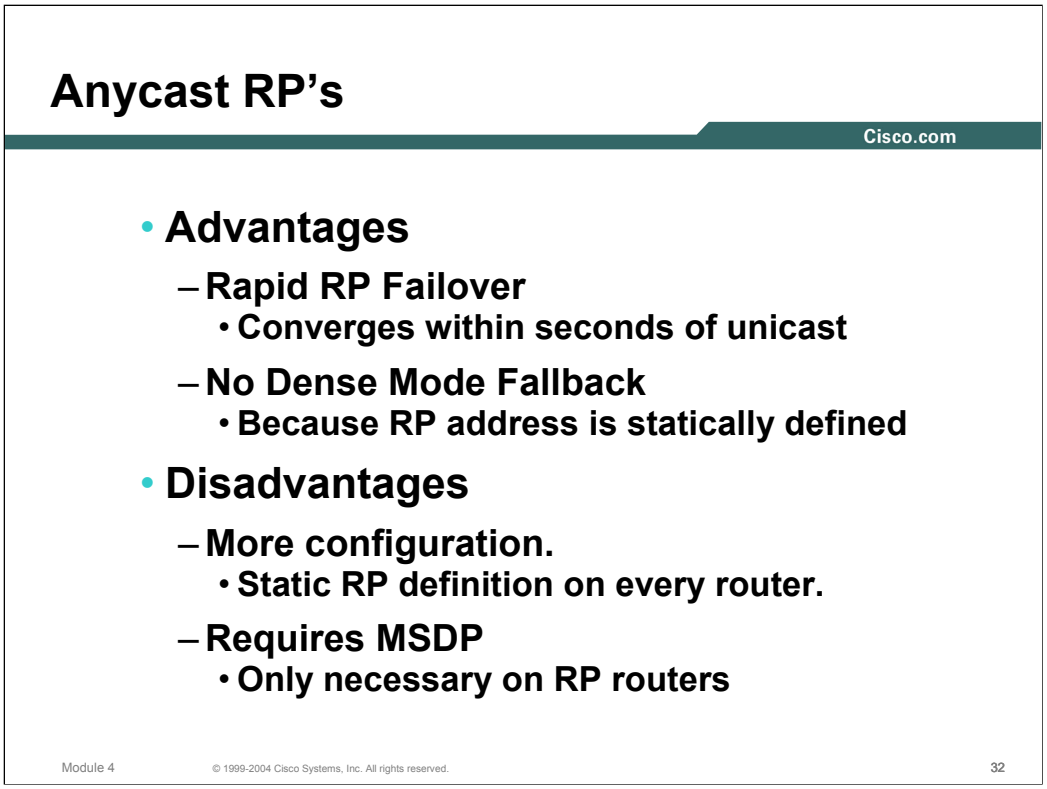

- **Anycast RP's**
  - Advantages
    - Anycast RP's have the fastest reconvergence time of all RP methods. In general, the PIM-SM network will reconverge on the remaining Anycast RP(s) within seconds of unicast routing reconverging.
    - The Anycast RP mechanism relies on statically configuring a single common RP address for a group range on every router in the network. As a result, it is not possible to loose RP information and fallback into Dense mode operation.
  - Disadvantages
    - Anycast RP requires more configuration than Auto-RP or BSR as it requires the static configuration of the Anycast RP address on every router in the network.
    - The use of MSDP is required to "heal" the network which is naturally partitioned by the use of Anycast RP. Not only is this an extra configuration step, it brings another protocol into use which increases network complexity.

Anycast RP Configuration

Cisco.com

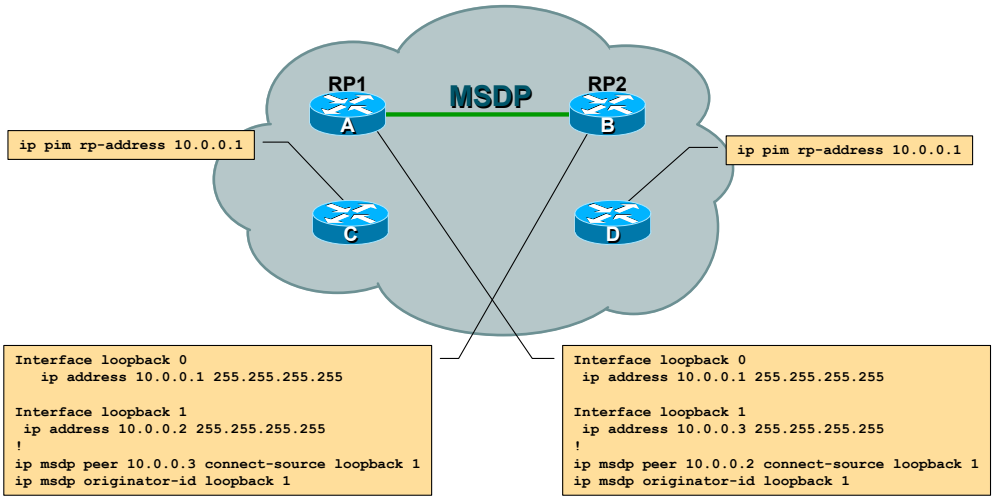

- **Anycast RP Example**
  - In this example, two Anycast RP's are configured with the same IP address, 10.0.0.1, using Loopback 0.
  - Each are connected via MSDP using their Loopback 1 addresses, 10.0.0.2 and 10.0.0.3.
    - (Yes, you must use some other address in the 'ip msdp peer' commands than 10.0.0.1.)

Tuning RP Operations

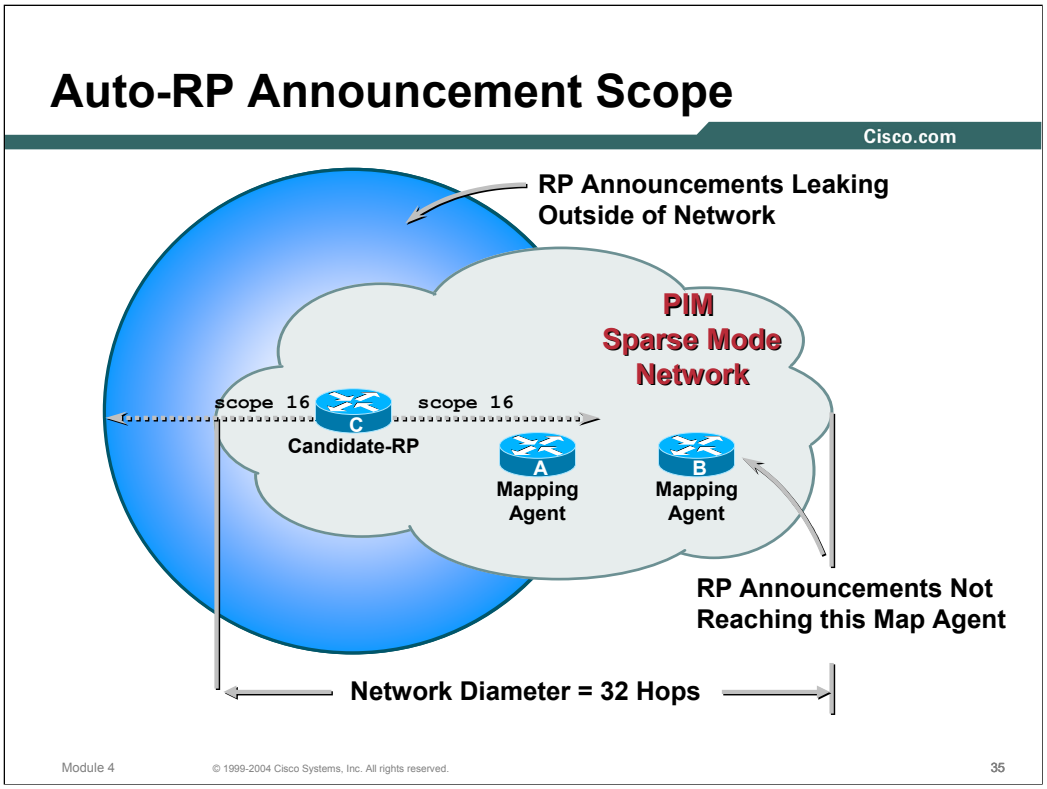

- **Auto-RP Announcement Scope**
  - Care must be taken in the selection of the TTL scope of RP Announcement messages that are sent by C-RPs to insure that the messages reach all Mapping Agents in the network.
- **Example**

  In the diagram above, an arbitrary scope of 16 was used in the 'ip pim send-rp-announce' command on the C-RP router. However, the maximum diameter of the network is greater than 16 hops and in this case one Mapping Agent is further away than 16 hops. As a result, this Mapping Agent does not receive the RP Announcement messages from the C-RP. This can cause the two Mapping Agents to have different information in their Group-to-RP mapping caches. If this occurs, each Mapping Agent will advertise a different router as the RP for a group which will have disastrous results.

  Notice also, that the C-RP is fewer than 16 hops way from the edge of the network. This can result in RP Announcement messages leaking into adjacent networks and causing Auto-RP problems in those networks.

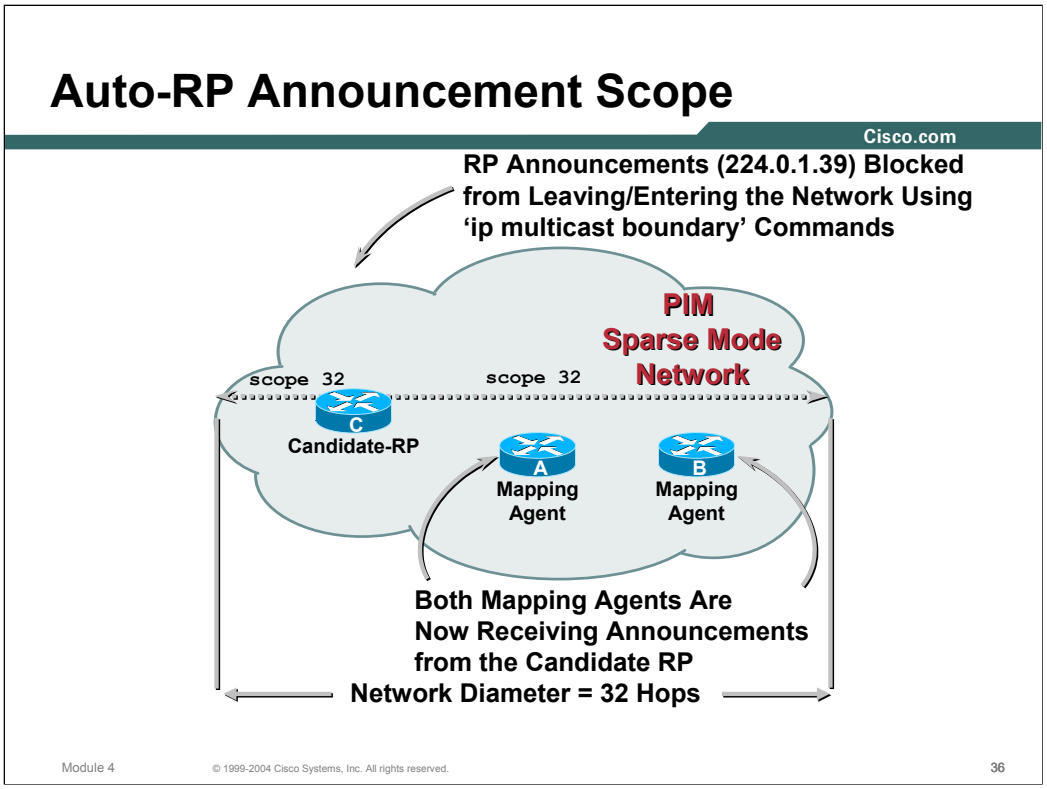

- **Auto-RP Announcement Scope**
  - The best way to avoid the problems on the preceding page is to use a sufficiently large enough scope so that the RP Announcement messages reach all Mapping Agents in the network.
- **Example**
  - In the above diagram, the maximum network diameter is 32. Therefore by setting the scope to 32 or greater, we are assured that the RP Announcements will reach both Mapping Agents shown in the example network.
  - In order to prevent RP Announcement messages from leaking into adjacent networks, a multicast boundary is defined for the Cisco-Announce (224.0.1.39) multicast group on all border routers in the network. This not only stops RP Announcement messages from leaking out, *it more importantly, stops any from leaking in from adjacent networks.*

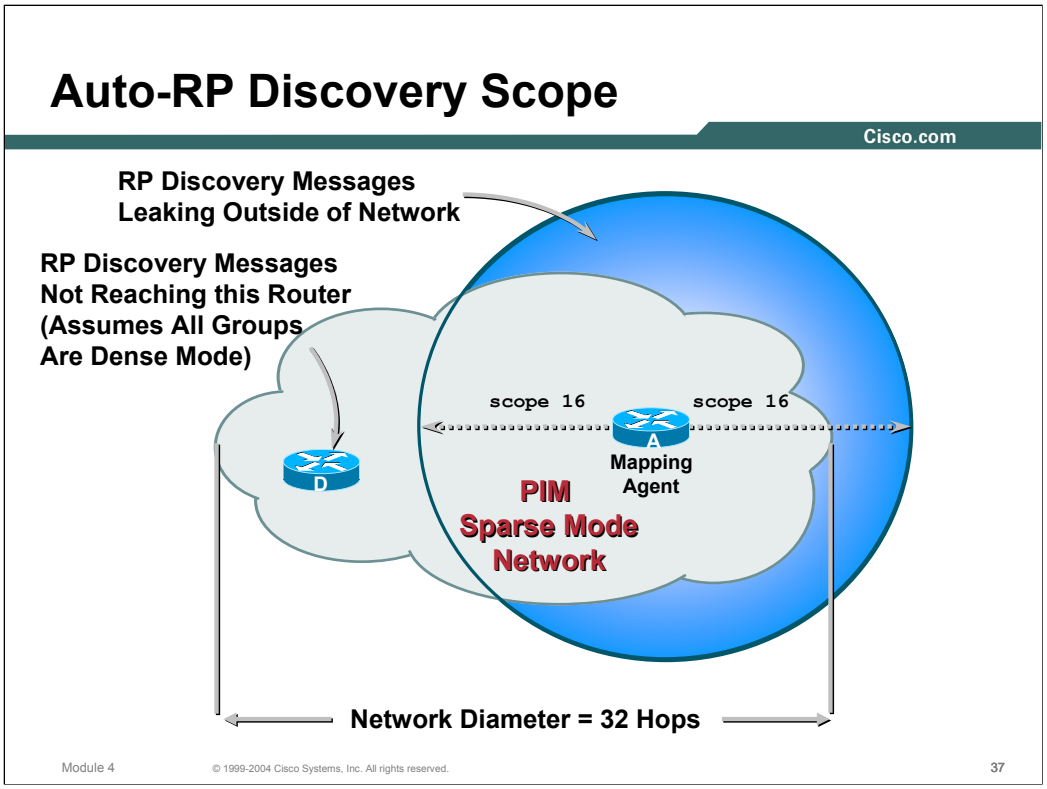

- **Auto-RP Discovery Scope**
  - Care must be taken in the selection of the TTL scope of RP Discovery messages that are sent by Mapping Agents to insure that the messages reach all routers in the network.
- **Example**

  In the diagram above, an arbitrary scope of 16 was used in the 'ip pim send-rp-discovery' command on the Mapping Agent. However, the maximum diameter of the network is greater than 16 hops and in this case, at least one router is further away than 16 hops. As a result, this router does not receive the RP Discovery messages from the MA. This can result in the router having no Group-to-RP mapping information. If this occurs, the router will attempt operate in Dense mode for all multicast groups while other routers in the network are working in Sparse mode.

  Notice also, that the MA is fewer than 16 hops way from the edge of the network. This can result in RP Discovery messages leaking into adjacent networks and causing Auto-RP problems in those networks.

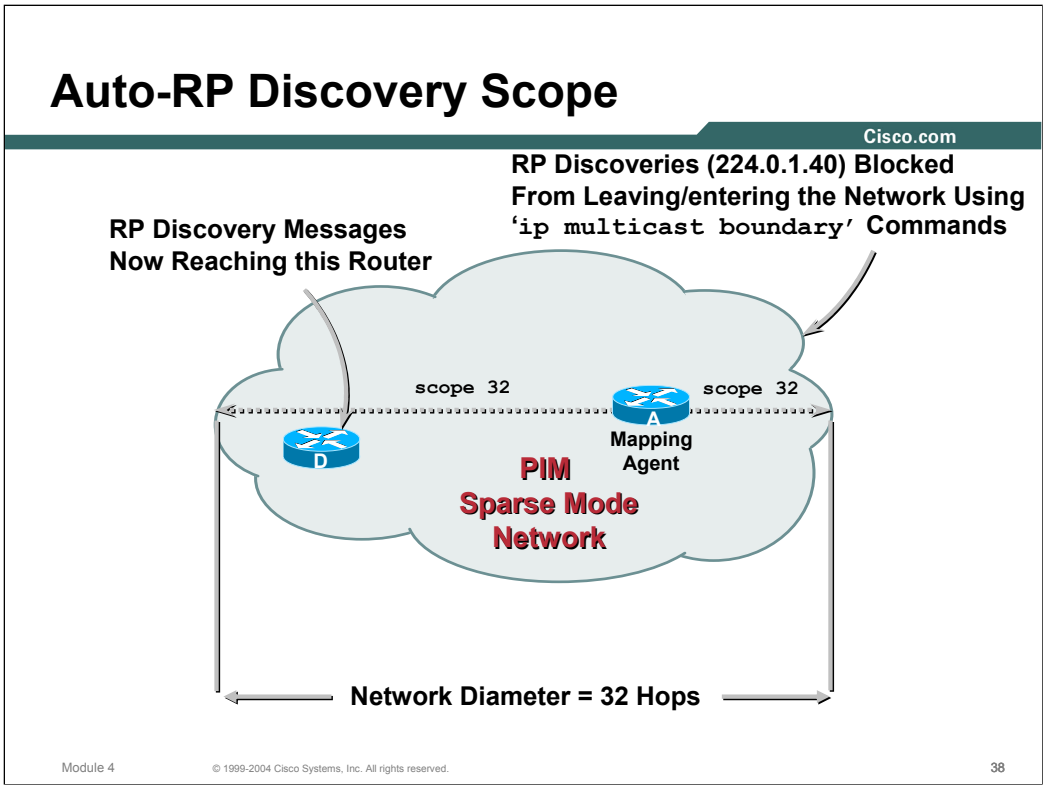

- **Auto-RP Discovery Scope**
  - The best way to avoid the problems on the preceding page is to use a sufficiently large enough scope so that the RP Discovery messages reach all routers in the network.
- **Example**
  - In the above diagram, the maximum network diameter is 32. Therefore by setting the scope to 32 or greater, we are assured that the RP Discovery messages will reach the farthest router in the network.
  - In order to prevent RP Discovery messages from leaking into adjacent networks, a multicast boundary is defined for the Cisco-Discovery (224.0.1.40) multicast group on all border routers in the network. This not only stops RP Discovery messages from leaking out, *it more importantly, stops any from leaking in from adjacent networks.*

Constraining Auto-RP Messages

Cisco.com

Need to Block Auto-RP Discovery (224.0.1.40) and Announcement (224.0.1.39) Messages from Entering/Leaving the Network

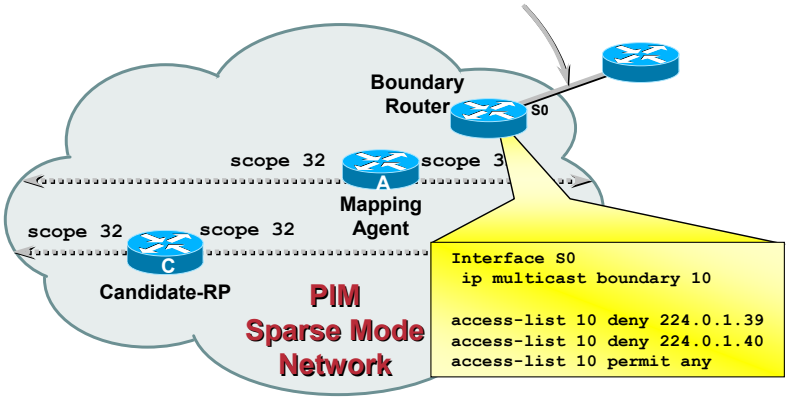

Module 4 39

- **Constraining Auto-RP Messages**
  - This example shows how to configure the multicast boundary on a border router so that Auto-RP messages do not leak into or out of the network.
    - On the border interface (in this case, Serial0) the 'ip multicast boundary' command is used.
    - The access list associated with the 'ip multicast boundary' command is as follows:

```
access-list 10 deny 224.0.1.39
access-list 10 deny 224.0.1.40
access-list 10 permit any
```

      The above access list stops the flow of multicast traffic for the two Auto-RP groups (224.0.1.39 and 224.0.1.40) while allowing all other multicast traffic to enter or exit via interface Serial 0.

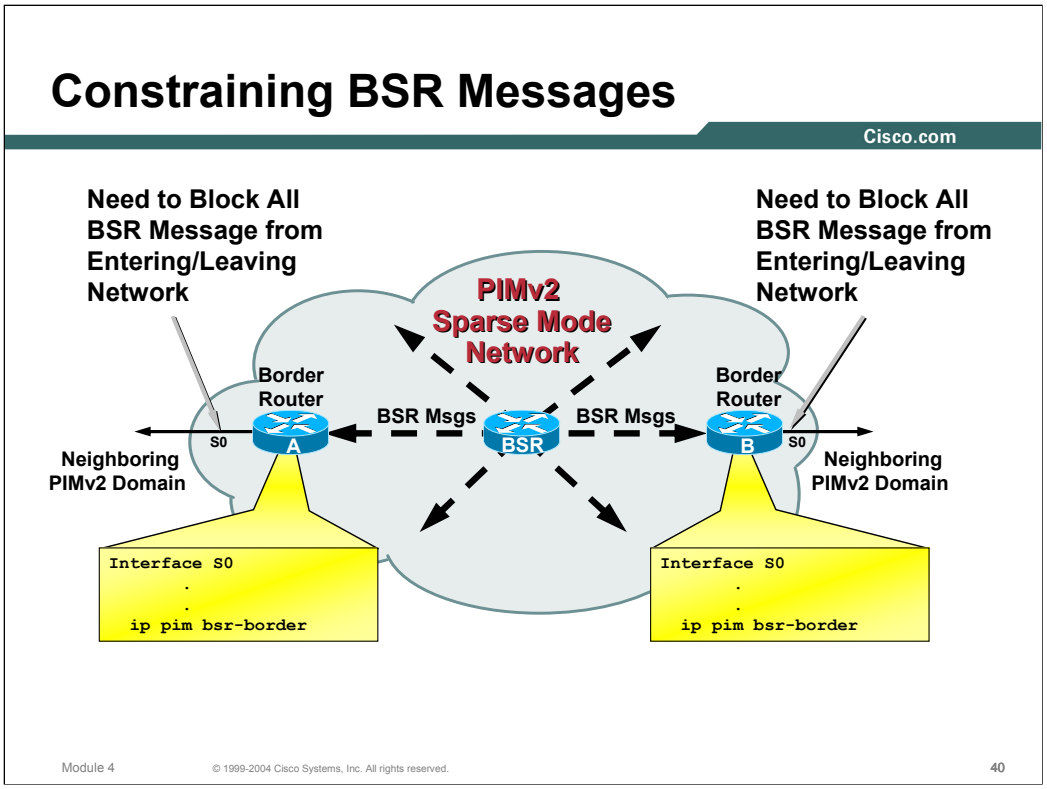

- **Constraining BSR Messages**
  - Like Auto-RP, allowing BSR messages to leak into or out of a network can cause problems both in the local network and in adjacent networks.
  - In order to block BSR messages from entering or exiting on a given interface, the ‘ip pim bsr-border’ interface command can be used.

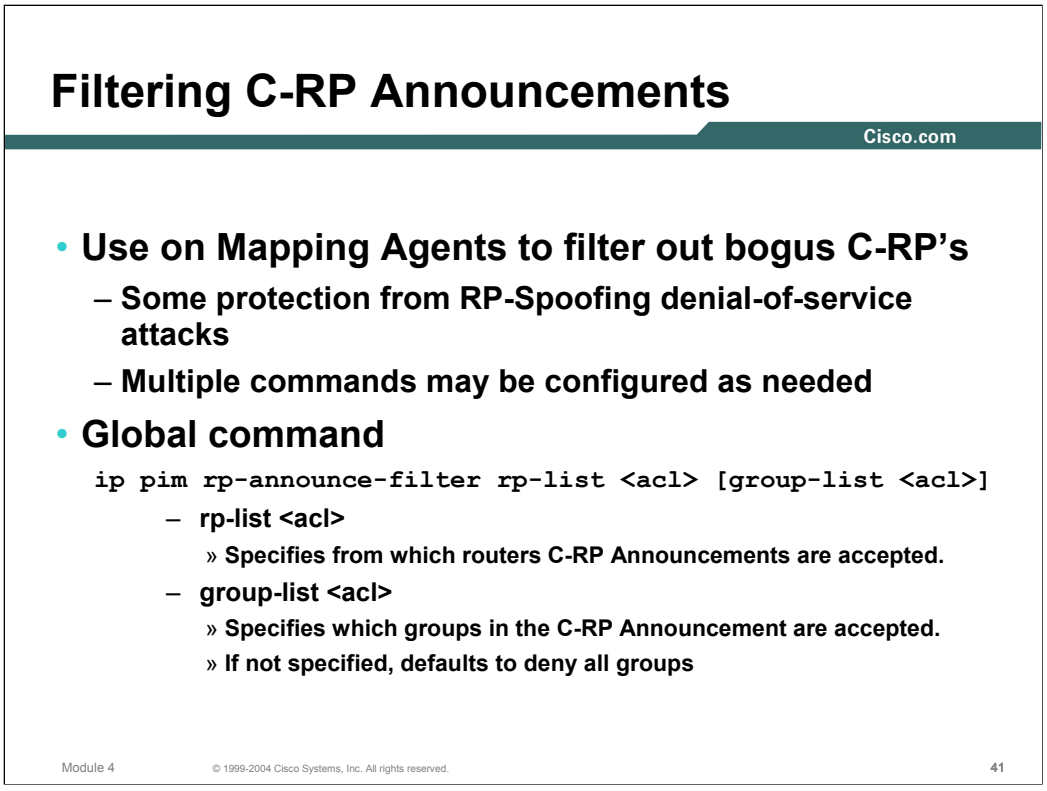

- **Filtering RP Announcements**

  Network Administrators may wish to configure Mapping Agents so that they will only accept C-RP Announcements from well-known routers in the network. This will prevent C-RP Announcements from bogus routers from being accepted and potentially being selected as the RP.

- **Global Command**

  `ip pim rp-announce-filter rp-list <acl> [group-list <acl>]`

  – The rp-list <acl> specifies the IP address(es) from which C-RP announcements will be accepted.

  – The option group-list <acl> specifies the group range(s) that are acceptable for the routers in the rp-list. If not specified, the default group-list <acl> is

    `deny all`

  – Multiple instances of this command may be configured.

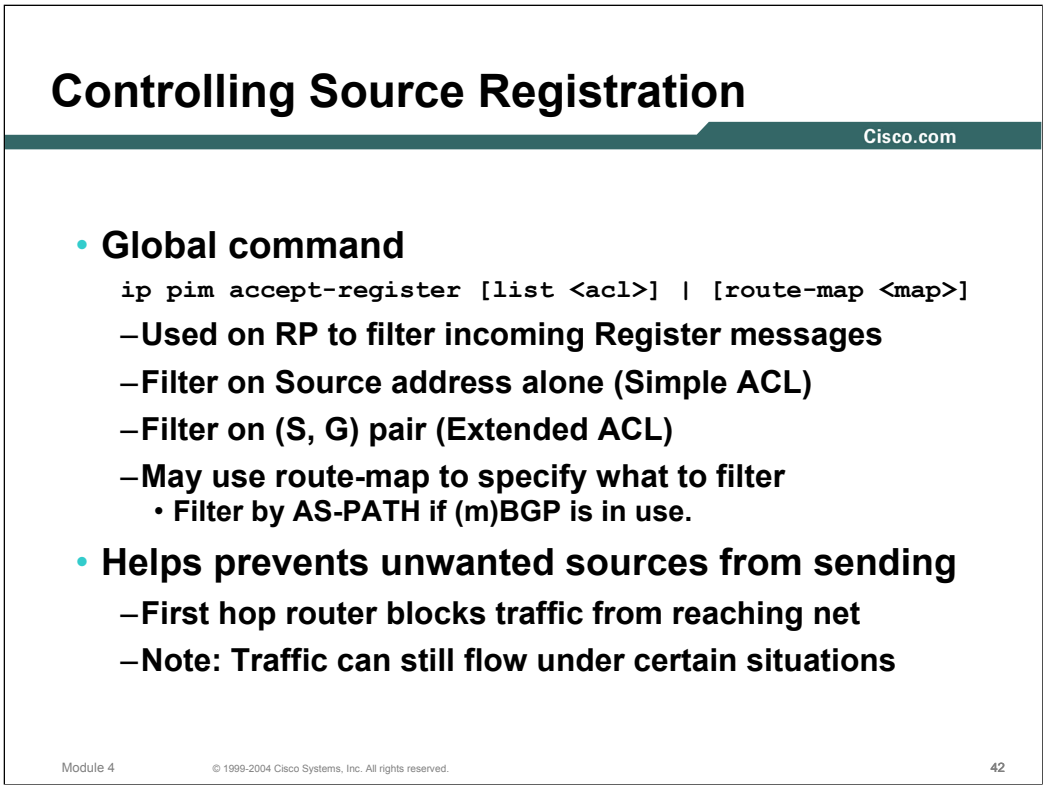

- **Controlling Source Registration**

  In some cases, it may be desirable to control which hosts in the network can actually source traffic to a group. While there is currently no way to prevent a bogus source from transmitting traffic on its local segment, we can prevent it from being registered to the RP. This will, in most cases, prevent this traffic from going past the first-hop router and reaching other hosts in the network.

  A new IOS command, 'ip pim accept-register' was introduced which when configured on an RP, controls which (S, G) Register messages will be accepted and which will be rejected.

- **Global Command (IOS 12.0(6) or later)**

  ip pim accept-register [list <acl>] | [route-map <map>]

  – If the "list <acl>" is specified, the <acl> can either be a simple access list to control which hosts may send to any groups or an extended access list that specifies both source and group address combinations that are permitted or denied from sending.

  – If the "route-map <map>" is specified, then only matching (S, G) traffic will be accepted. (Note: This permits other matching criteria to be considered such as AS-PATH.)

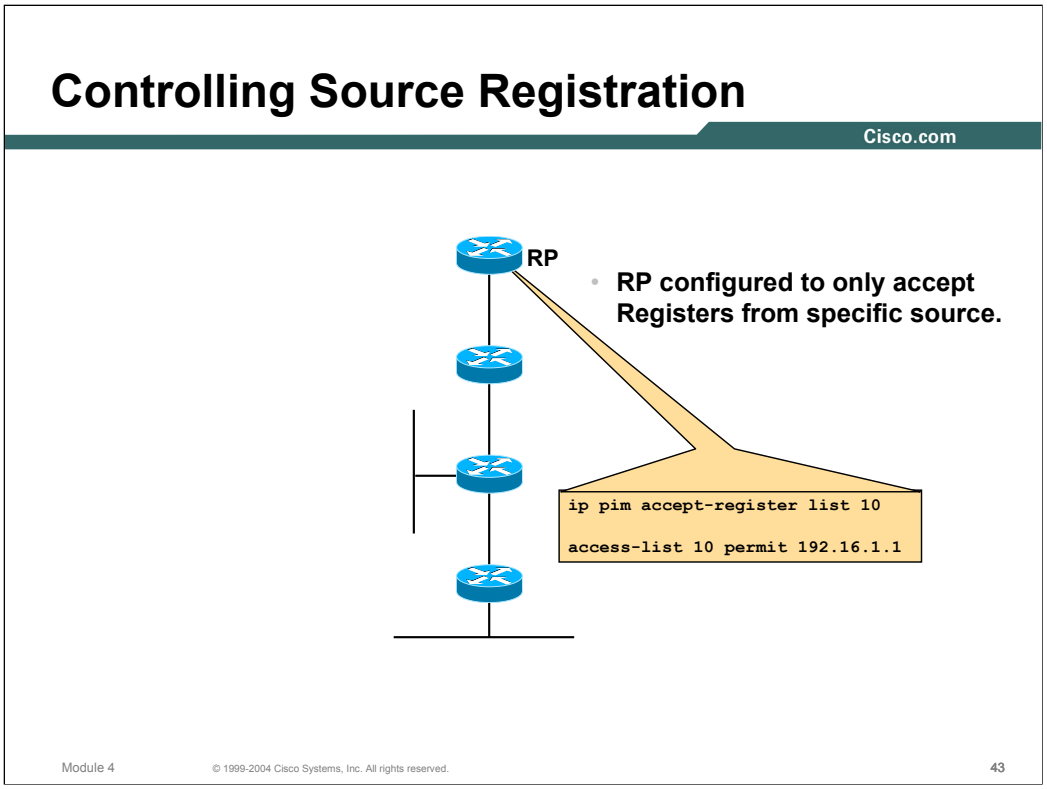

- **Controlling Source Registration – Example**
  - In this example, the RP is configured with an **ip pim accept-register** command along with an ACL to permit PIM Registers to be accept only for source 192.16.1.1.

Controlling Source Registration

Cisco.com

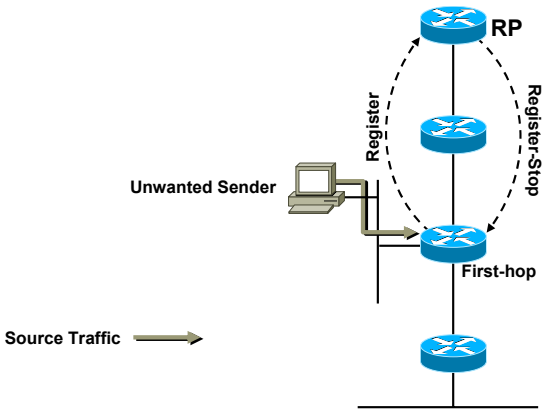

- **Unwanted source traffic hits first-hop router.**
- **First-hop router creates (S,G) state and sends Register.**
- **RP rejects Register, sends back a Register-Stop.**

Module 4 44

- **Controlling Source Registration – Example**
  - If an unwanted source begins transmitting, the first-hop router creates (S,G) state and sends a PIM **Register** message for this source to the RP.
  - The RP rejects this Register as a result of the **ip pim accept-register** command and sends back an immediate PIM **Register-Stop** message to the first-hop router.

Controlling Source Registration

Cisco.com

Weaknesses in 'accept-register' usage.

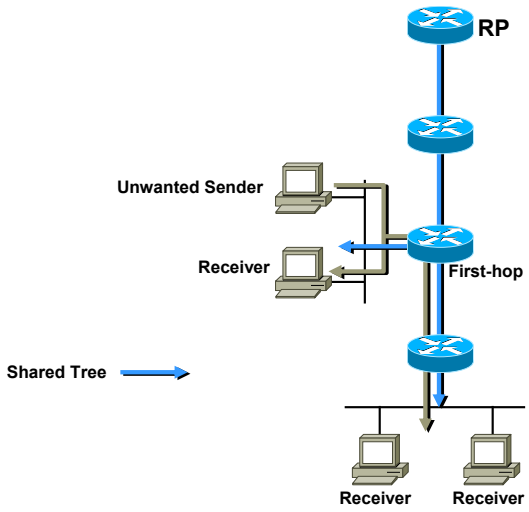

- **Traffic will flow on local subnet where source resides.**
- **Traffic will flow from first-hop router down any branches of the Shared Tree.**
  - **Results when (*,G) OIL is copied to (S,G) OIL at first-hop router.**
  - **Causes (S,G) traffic to flow down all interfaces in (*,G) OIL of first-hop router.**

Module 4 45

- **Controlling Source Registration – Example**
  - Weaknesses in the '**accept-register**' mechanism.
    - Source traffic will still flow on the local subnet where the source resides. Any hosts that have joined this multicast group will get the unwanted traffic regardless if the **accept-register** command has been configured on the RP or not.
    - Unwanted source traffic can still flow from the first-hop router *down* any branches of the Shared Tree as shown in the example above. This is occurs when the interfaces in the (*,G) outgoing interface list (OIL) are copied into the OIL of the newly created (S,G) entry.


Cisco Systems