



NETWORKERS 2004

ADVANCED IP MULTICAST

tkramer@cisco.com

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

1

Agenda

Cisco.com

- **MBGP (routing)**
- **MSDP (source discovery)**
- **MBGP/MSDP Examples**
- **SSM (Source Specific Multicast)**
- **Security**
- **Source Redundancy**

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

2

MULTIPROTOCOL BGP (MBGP)



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

3

MBGP Overview

Cisco.com

- **MBGP: Multiprotocol BGP**
 - Defined in RFC 2283 (extensions to BGP)
 - Can carry different types of routes
 - IPv4 Unicast IPv6 Unicast
 - IPv4 Multicast IPv6 Multicast
 - May be carried in same BGP session
 - Does not propagate multicast state info
 - Still need PIM to build Distribution Trees
 - Same path selection and validation rules
 - AS-Path, LocalPref, MED, ...

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

4

•MBGP Overview

•Multiprotocol BGP (MBGP) is defined in RFC 2283. This RFC defines extensions to the existing BGP protocol to allow it to carry more than just IPv4 route prefixes. Examples of some of the new types of routing information include (but are not limited to):

- IPv4 prefixes for Unicast routing
- IPv4 prefixes for Multicast RPF checking
- IPv6 prefixes for Unicast routing

•A common misconception is that MBGP is a replacement for PIM. This is incorrect. MBGP does not propagate *any* multicast state information nor does it build any sort of multicast distribution trees. MBGP *can* distribute unicast prefixes that can be used for the multicast RPF check.

•Because MBGP is an extension to the existing BGP protocol, the same basic rules apply to path selection, path validation, etc.

MBGP Overview

Cisco.com

- **Separate BGP tables maintained**
 - Unicast BGP Table (U-Table)
 - Multicast BGP Table (M-Table)
 - BGP 'nlri' keyword specifies which BGP Table
 - Allows different unicast/multicast topologies or policies
- **Unicast BGP Table (U-Table)**
 - Contains unicast prefixes for unicast forwarding
 - Populated with BGP unicast NLRI
- **Multicast BGP Table (M-Table)**
 - Contains unicast prefixes for RPF checking
 - Populated with BGP multicast NLRI

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

5

•Routing Information Bases

- Previously, BGP only maintained a single BGP Table for IPv4 unicast prefixes. In the case of MBGP, separate BGP Tables must be maintained for each type of routing information being exchanged. This implies that a separate Unicast BGP Table (U-Table) and a separate Multicast BGP Table (M-Table) can be maintained by MBGP.

- A new "nlri" keyword was added to the Cisco IOS command structure to differentiate between the U-Table and the M-Table. (Note: This keyword will soon be depreciated in order to generalize MBGP for other protocols such as IPv6. Consult you latest IOS Documentation for the correct syntax.)

- Unicast BGP Table (U-Table)

- This table contains the unicast prefixes that was previously used by BGP for IPv4 unicast traffic forwarding.

- Multicast BGP Table (M-Table)

- This new table contains the same *type* of unicast prefixes contained in the U-Table except that the prefixes stored in the M-Table are used to RPF check arriving multicast traffic.

MBGP Update Message

Cisco.com

- **Address Family Information (AFI)**
 - **Identifies Address Type (see RFC1700)**
 - **AFI = 1 (IPv4)**
 - **AFI = 2 (IPv6)**
- **Sub-Address Family Information (Sub-AFI)**
 - **Sub category for AFI Field**
 - **Address Family Information (AFI) = 1 (IPv4)**
 - **Sub-AFI = 1 (NLRI is used for unicast)**
 - **Sub-AFI = 2 (NLRI is used for multicast RPF check)**

RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

6

- **Address Family Information (AFI)**
 - This field is based on the address families defined in RFC1700.
 - **AFI = 1 (IPv4)**
 - **AFI = 2 (IPv6)**
- **Sub-Address Family Information (Sub-AFI)**
 - This field contains further information regarding the type of routing information being exchanged in the NLRI field. The following are the current definitions for Sub-AFI codes associated with the IPv4 Address Family:
 - **Sub-AFI = 1 (NLRI is used for unicast routing)**
 - **Sub-AFI = 2 (NLRI is used for multicast RPF check)**
 - **Sub-AFI = 3 (NLRI is used for both unicast routing and multicast RPF checking)**

PIM RPF Calculation Details

Cisco.com



Global Command: `ip multicast longest-match`

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

7

MBGP—Capability Negotiation

Cisco.com

- **Keyword on neighbor command**

```
neighbor <foo> remote-as <asn> nlri multicast unicast
```

- Configures router to negotiate either or both NLRI
- If neighbor configures both or subset, common NLRI is used in both directions
- If there is no match, notification is sent and peering doesn't come up


RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

8

MBGP — Capability Negotiation

Cisco.com




•MBGP Capability Negotiation

- In this example, the router on the left is configured to peer with the router on the right. The command:
 - neighbor 192.168.100.2 remote-as 321 nlri unicast multicast
- instructs the router on the left to attempt to negotiate both unicast and multicast NLRI exchange.

MBGP — Capability Negotiation

Cisco.com



•MBGP Capability Negotiation


•In this example, the router on the right is configured to peer with the router on the left. The command:

```
neighbor 192.168.100.1 remote-as 123 nlri unicast
multicast
```

instructs the router on the right to attempt to negotiate *both* unicast and multicast NLRI exchange.

MBGP — Capability Negotiation

Cisco.com




•MBGP Capability Negotiation

- In this example, the two routers can be seen exchanging Capabilities in the router debug message output. In this case, both unicast and multicast NLRI has been successfully negotiated. Therefore, both unicast and multicast NLRI will be exchanged in this single MBGP peer session

MBGP—NLRI Information

Cisco.com



- Storage of arriving NLRI information depends on AFI/SAFI fields in the Update message
 - Unicast BGP Table only (AFI=1/SAFI=1 or old style NLRI)

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.


12

•Receiving Update Messages

- Storage of arriving NLRI information depends on the AFI/Sub-AFI fields in the MP_REACH_NLRI attribute.
- In the above example, an Update message is received that contains an MP_REACH_NLRI with an AFI/Sub-AFI = 1/1 which indicates that the information is for the Unicast RIB. As a result, the information is processed by the MBGP input process and written into the U-RIB as shown.

MBGP—NLRI Information

Cisco.com



- Storage of arriving NLRI information depends on AFI/SAFI fields in the Update message
 - Unicast BGP Table only (AFI=1/SAFI=1 or old style NLRI)
 - **Multicast BGP Table only (AFI=1/SAFI=2)**

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

13


•Receiving Update Messages

•In the above example, an Update message is received that contains an MP_REACH_NLRI with an AFI/Sub-AFI = 1/2 which indicates that the information is for the Multicast RIB. As a result, the information is processed by the MBGP input process and written into the M-RIB as shown.

MBGP—NLRI Information

Cisco.com

Congruent Topologies



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

14

MBGP—NLRI Information

Cisco.com

Incongruent Topologies



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

15

MBGP Syntax Change

Cisco.com

NLRI Syntax

```
router bgp 5
 network 171.69.214.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 171.69.214.38 remote-as 2 nlri unicast
 neighbor 171.69.214.50 remote-as 2 nlri multicast
```

Address-Family Syntax

```
router bgp 5
 no bgp default ipv4-unicast
 neighbor 171.69.214.38 remote-as 2
 neighbor 171.69.214.50 remote-as 2
 !
 address-family ipv4 unicast
 neighbor 171.69.214.38 activate
 network 171.69.214.0 mask 255.255.255.0
 exit-address-family
 !
 address-family ipv4 multicast
 neighbor 171.69.214.50 activate
 network 171.69.214.0 mask 255.255.255.0
 exit-address-family
```

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

16

MBGP—Summary

Cisco.com

- **Solves part of inter-domain problem**
 - Can exchange multicast routing information
 - Uses standard BGP configuration knobs
 - Permits separate unicast and multicast topologies if desired
- **Still must use PIM to:**
 - Build distribution trees
 - Actually forward multicast traffic
 - PIM-SM recommended

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

17

MULTICAST SOURCE DISCOVERY PROTOCOL (MSDP)



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

18

MSDP Overview

MSDP Example

MSDP Peers 
Source Active Messages 




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

MSDP Overview

MSDP Example

MSDP Peers



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

20

MSDP Overview

MSDP Example

MSDP Peers ———
Multicast Traffic ———



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

21

MSDP Overview

MSDP Example

MSDP Peers ———
Multicast Traffic ———



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

22

MSDP Overview

MSDP Example

MSDP Peers ———
Multicast Traffic ———



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

23

- **MSDP Source Active (SA) Messages**
 - **Used to advertise active Sources in a domain**
 - **Carry 1st multicast packet from source**
 - Hack for Bursty Sources (ala SDR)
 - **SA Message Contents:**
 - IP Address of Originator (RP address)
 - Number of (S, G)'s pairs being advertised
 - List of active (S, G)'s in the domain
 - Encapsulated Multicast packet

Receiving SA Messages Forwarding of SA messages ← “RPF-check”

Cisco.com

- **RPF Check Rules depend on peering**
 - Rule 1: Sending MSDP peer = i(m)BGP peer
 - Rule 2: Sending MSDP peer = e(m)BGP peer
 - Rule 3: Sending MSDP peer != (m)BGP peer
- **Exceptions:**
 - RPF check is skipped when:
 - Sending MSDP peer = Originating RP
 - Sending MSDP peer = Mesh-Group peer
 - Sending MSDP peer = only MSDP peer
 - (i.e. the ‘default-peer’ or the only ‘msdp-peer’ configured.)

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

25

•Receiving SA Messages

- RPF Check rules depend on the BGP peering between the MSDP peers.
 - Rule 1: Applied when the sending MSDP peer is also an i(m)BGP peer.
 - Rule 2: Applied when the sending MSDP peer is also an e(m)BGP peer.
 - Rule 3: Applied when the sending MSDP peer is not an (m)BGP peer.
- RPF Checks are not done in the following cases:
 - If the sending MSDP peer is the only MSDP Peer. This would be the case if a single ‘msdp-peer’ command is configured or if only the ‘default-peer’ command is used.
 - If the sending MSDP peer is a Mesh-Group peer.
 - If the sending MSDP peer address is the RP address contained in the SA message.

RPF Check Rule 1

Cisco.com

- **When MSDP peer = i(m)BGP peer**
 - Find “Best Path” to RP in BGP Tables
 - Search MRIB first then URIB
 - If no path to Originating RP found, RPF Fails
 - Note “BGP peer” that advertised path
 - (i.e. IP Address of BGP peer that sent us this path)
 - **Warning:**
 - This is not the same as the Next-hop of the path!!!
 - i(m)BGP peers normally do not set Next-hop = Self.
 - This is also not necessarily the same as the Router-ID!
 - **Rule 1 Test Condition:**
 - **MSDP Peer address = BGP peer address?**
 - If Yes, RPF Succeeds

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.


26

•RPF Check Rule 1

- Applied when the sending MSDP peer is also an i(m)BGP peer.
- Search the BGP MRIB for the “best-path” to the RP that originated the SA message. If a path is not found in the MRIB, search the URIB. If a path is still not found, the RPF check fails.
- Determine the address of the “BGP Neighbor” for this path. (This is the address of the BGP neighbor the sent us this path in a BGP Update message.)
 - Be careful not to assume the “BGP Neighbor” address is the same as the Next-Hop address in the path. Since i(m)BGP peers do not update the Next-Hop attribute of a path, it is usually the case that the Next-Hop address is not the same as the address of the BGP peer that sent us the path.
 - The “BGP Neighbor” address is also not necessarily the same as the BGP “Router-Id” of the peer that sent us the path.
- Rule 1 Test:
 - If the IP address of the sending MSDP peer is the same as the BGP Neighbor address (i.e. the address of the BGP peer that sent us the path), then the RPF check succeeds; otherwise it fails.

Rule1: MSDP peer = i(m)BGP peer

Cisco.com




•Rule 1 Example 1

- In this example, router A receives an SA message originated by router G from router E which is an i(m)BGP peer.
- Applying Rule 1, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - This best path was received from i(m)BGP peer, 172.16.3.1
 - The sending MSDP peer address is also 172.16.3.1
 - Therefore the RPF check Rule 1 succeeds.

Rule1: MSDP peer = i(m)BGP peer

Cisco.com




•Rule 1 Example 2

- In this example, router A receives the same SA message (originated by router G) from router D which is an i(m)BGP peer.
- Applying Rule 1, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - This best path was received from i(m)BGP peer, 172.16.3.1
 - The sending MSDP peer address is 172.16.4.1
 - Therefore RPF check Rule 1 fails.

Rule1: MSDP peer = i(m)BGP peer

Cisco.com




•Rule 1 Common Mistake 1

- The most common mistake is for the MSDP and (m)BGP peering sessions to the same router to use different IP addresses. In this example,
 - The MSDP peer address to router C is 172.16.20.1
 - The (m)BGP peer address to router C is 172.16.3.1
- Router A receives the SA message (originated by router G) from router E which is an i(m)BGP peer.
- Applying Rule 1, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - This best path was received from router C which is the i(m)BGP peer with an IP address of 172.16.3.1
 - However, the sending MSDP peer (also router C) address is 172.16.20.1
 - Therefore RPF check Rule 1 fails. (Even though the SA message arrived via the correct path.)

Rule1: MSDP peer = i(m)BGP peer

Cisco.com



•Rule 1 Common Mistake 2

•The other common mistake is to failure of the MSDP topology to follow the i(m)BGP peering topology. This can happen when Route Reflectors are used. In this example,

- The MSDP peer address of router E is 172.16.3.1
- The i(m)BGP peer router is the Route Reflector “RR” whose peer address is 172.16.1.1
- Router A receives an SA message (originated by router G) from router E which is the MSDP peer.
- Applying Rule 1, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - This best path was received from the Route Reflector which is the i(m)BGP peer with an IP address of 172.16.1.1
 - However, the sending MSDP peer is router E whose address is 172.16.20.1
 - Therefore RPF check Rule 1 fails.

RPF Check Rule 2

Cisco.com

- **When MSDP peer = e(m)BGP peer**
 - Find (m)BGP “Best Path” to RP
 - **Search MRIB first then URIB**
 - If no path to Originating RP found, RPF Fails
 - **Rule 2 Test Condition:**
 - **First AS in path to the RP = MSDP peer?**
 - If Yes, RPF Succeeds

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.


31

•RPF Check Rule 2

- Applied when the sending MSDP peer is also an e(m)BGP peer.
- Search the BGP MRIB for the “best-path” to the RP that originated the SA message. If a path is not found in the MRIB, search the URIB. If a path is still not found, the RPF check fails.
- Rule 2 Test:
 - If the first AS in the ‘best-path’ to the RP is the same as the AS of the e(m)BGP peer (which is also the sending MSDP peer), then the RPF check succeeds; otherwise it fails.

Rule2: MSDP peer = e(m)BGP peer

Cisco.com




•Rule 2 Example 1

- In this example, router A receives an SA message originated by router G via router E which is an e(m)BGP peer.
- Applying Rule 2, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - The first-hop AS in the best path to the originating RP is AS3.
 - The origin AS of the sending MSDP peer (172.16.3.1) is also AS3. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
 - Therefore the RPF check Rule 2 succeeds.

Rule2: MSDP peer = e(m)BGP peer

Cisco.com



•Rule 2 Example 2

- In this example, router A receives the same SA message (originated by router G) via router D which is also an e(m)BGP peer.
- Applying Rule 2, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - The first-hop AS in the best path to the originating RP is AS3.
 - The origin AS of the sending MSDP peer (172.16.4.1) is not AS3, it is AS1. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
 - Therefore the RPF check Rule 2 fails.

RPF Check Rule 3

Cisco.com

- **When MSDP peer != (m)BGP peer**
 - Find (m)BGP “Best Path” to RP
 - Search MRIB first then URIB
 - If no path to Originating RP found, RPF Fails
 - Find (m)BGP “Best Path” to MSDP peer
 - Search MRIB first then URIB
 - If no path to sending MSDP Peer found, RPF Fails
 - Note AS of sending MSDP Peer
 - Origin AS (last AS) in AS-PATH to MSDP Peer
 - Rule 3 Test Condition:
 - First AS in path to RP = Sending MSDP Peer AS ?
 - If Yes, RPF Succeeds

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.


34

•RPF Check Rule 3

- Applied when the sending MSDP peer is not an (m)BGP peer at all.
- Search the BGP MRIB for the “best-path” to the RP that originated the SA message. If a path is not found in the MRIB, search the URIB. If a path is still not found, the RPF check fails.
- Search the BGP MRIB for the “best-path” to the MSDP peer that sent us the SA message. If a path is not found in the MRIB, search the URIB. If a path is still not found, the RPF check fails.
- Note the AS of MSDP peer that sent us the SA. (This is the “origin AS” which is the last AS in the AS-PATH to the MSDP peer.)
- Rule 3 Test:
 - If the first AS in the ‘best-path’ to the RP is the same as the AS of the sending MSDP peer, then the RPF check succeeds; otherwise it fails.

Rule3: MSDP peer != BGP peer

Cisco.com



•Rule 3 Example 1


•In this example, router A receives an SA message originated by router G via router E which is neither an i(m)BGP peer nor an e(m)BGP peer.

•Applying Rule 3, the following occurs:

- The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
- The first-hop AS in the best path to the originating RP is AS3.
- The origin AS of the sending MSDP peer (172.16.3.1) is also AS3. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
- Therefore the RPF check Rule 3 succeeds.

Rule3: MSDP peer != BGP peer

Cisco.com




•Rule 3 Example 2

- In this example, router A receives the same SA message (originated by router G) via router D which is neither an i(m)BGP peer nor an e(m)BGP peer.
- Applying Rule 3, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - The first-hop AS in the best path to the originating RP is AS3.
 - The origin AS of the sending MSDP peer (172.16.4.1) is AS1. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
 - Therefore the RPF check Rule 3 fails.

Rule3: MSDP peer != BGP peer

Cisco.com



- **Configure peers**

```
ip msdp peer <ip-address> [connect-source <i/f>]
```

- **Configure default peer**


```
ip msdp default-peer <ip-address> [prefix-list acl]
```

- **Mesh groups**

```
ip msdp mesh-group <name> <ip-address>
```

MSDP Mesh-Group Example

Cisco.com



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

39

•MSDP Mesh-Group Example.

- In the above example, routers R1, R2 and R3 are all configured as members of the same MSDP mesh-group. In addition, router R1 is also MSDP peering with router R4 and router R3 is MSDP peering with router R5. Neither R4 nor R5 are members of the MSDP mesh-group.

- Assume router R4 originates an SA message for a source in its local PIM-SM domain. This message is sent to route R2 as shown in the drawing above.

- When router R2 receives this SA message, it must perform an RPF check on the message because it was received from an MSDP peer that is not a member of the mesh-group.


- In this case the RPF check is successful and router R2 floods the SA message (received from a non-mesh-group member) to all other members of the mesh-group.

- When routers R1 and R3 receive the SA message from mesh-group member R2, they do not have to perform an RPF check on the arriving message nor do they flood the SA message to each other since they are both members of the mesh-group. (They know that the other members of the mesh-group will have received a copy directly from R2 and therefore they do not have to forward the SA message to each other. This is why a full mesh between mesh-group members is required.)

Intradomain MSDP peering scenarios

Cisco.com

Hierarchical mesh groups can help reduce MSDP peerings



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

40

MSDP Configuration

Cisco.com

- **RFC 3618**
- **Filtering**
 - Can filter SA in/out, groups, with acls or route-maps
- **For configuration commands see:**
 - <ftp://ftpeng.cisco.com/ipmulticast/Multicast-Commands>
- **For MSDP BCP (Best Current Practice) Draft:**
 - [draft-ietf-mboned-msdp-deploy-06.txt](#)

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

41

- **New IOS command**

```
ip msdp rpf rfc3618
```

- **MSDP SA RPF check using IGP**
- **Accept SA's from BGP NEXT HOP**
- **Accept SA's from closest peer along the best path to the originating RP**
- **“show ip msdp rpf”**
- **12.0(27)S**

MSDP RPF check using IGP

Cisco.com

- **When MSDP peer = IGP peer (No BGP)**

Find best IGP route to RP

Search URIB

If route to Originating RP found and:

**If IGP next hop (or advertiser) address for RP is the
MSDP peer and in UP state, then that is the RPF
peer.**

If route not found: Fall through to the next rule.


RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

43

IGP Rule: MSDP peer = IGP peer (Next hop)

Cisco.com




•Rule 2 Example 1

- In this example, router A receives an SA message originated by router G via router E which is an e(m)BGP peer.
- Applying Rule 2, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - The first-hop AS in the best path to the originating RP is AS3.
 - The origin AS of the sending MSDP peer (172.16.3.1) is also AS3. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
 - Therefore the RPF check Rule 2 succeeds.

IGP Rule: MSDP peer = IGP peer (Advertiser)

Cisco.com




•Rule 2 Example 1

- In this example, router A receives an SA message originated by router G via router E which is an e(m)BGP peer.
- Applying Rule 2, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - The first-hop AS in the best path to the originating RP is AS3.
 - The origin AS of the sending MSDP peer (172.16.3.1) is also AS3. (This is determined by locating the best-path to the MSDP peer and then finding the last AS in the AS-Path list.)
 - Therefore the RPF check Rule 2 succeeds.

SA's accepted from Next Hop

Cisco.com




•Rule 1 Common Mistake 2

•The other common mistake is to failure of the MSDP topology to follow the i(m)BGP peering topology. This can happen when Route Reflectors are used. In this example,

- The MSDP peer address of router E is 172.16.3.1
- The i(m)BGP peer router is the Route Reflector “RR” whose peer address is 172.16.1.1
- Router A receives an SA message (originated by router G) from router E which is the MSDP peer.
- Applying Rule 1, the following occurs:
 - The best path in the BGP M-RIB for 172.16.6.1 (the originating RP) is located.
 - This best path was received from the Route Reflector which is the i(m)BGP peer with an IP address of 172.16.1.1
 - However, the sending MSDP peer is router E whose address is 172.16.20.1
 - Therefore RPF check Rule 1 fails.

Accept SA along RPF path

Cisco.com



Existing Rule: If first AS in best path to the RP != MSDP peer
RPF Fails

New code: Choose peer in CLOSEST AS along best AS path to the RP.
Loosens rule a bit.

RPF Succeeds.

BGP Peer ———
MSDP Peer - - -
SA Message →

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

47

New MSDP RPF command

Cisco.com

```
Router-A# show ip msdp rpf 2.1.1.1
RPF peer information for Router-B (2.1.1.1)
  RPF peer: Router-C (3.1.1.1)
  RPF route/mask: 2.1.1.0/24
  RPF rule: Peer is IGP next hop of best route
  RPF type: unicast (ospf 1)
```

MBGP/MSDP EXAMPLES



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

49

MSDP Application—Anycast RP

Cisco.com

- **RFC 3446 Anycast RP mechanism using PIM and MSDP**
- **Within a domain, deploy more than one RP for the same group range**
- **Give each RP the same IP address assignment**
- **Sources and receivers use closest RP**
- **Used intra-domain to provide redundancy and RP load sharing**

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

50

MSDP Application—Anycast RP

Cisco.com

- **Sources from one RP are made known to other RPs using MSDP**
- **When an RP goes down, sources and receivers are taken to new RP via unicast routing**
 - **Fast convergence**


RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

51

Anycast RP—Overview

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

52

Anycast RP—Overview

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

53

Anycast RP Configuration

Cisco.com



•Anycast RP Example

- In this example, two Anycast RP's are configured with the same IP address, 10.1.1.1, using Loopback 0.
- Each are connected via MSDP using their Loopback 1 addresses, 10.0.0.1 and 10.0.0.2.
 - (Yes, you must use some other address in the 'ip msdp peer' commands than 10.0.0.1.)

MBGP/MSDP Single Homed Customer

BackUp - Slides



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.


55

Single-Homed, Customer RP, Non-MBGP

Cisco.com

Tail-site Customer

Transit AS109



Note: Access-list 111 = Recommended SA Filter

RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

56


Single-Homed, Customer RP, Non-MBGP

Cisco.com

Tail-site Customer



Transit AS109



Note: Access-list 111 = Recommended SA Filter

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

57


Single-Homed, Customer RP, Non-MBGP

Cisco.com

MSDP RPF Check

Tail-site Customer

Transit AS109



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

58


Single-Homed, Customer RP, Non-MBGP

Cisco.com

Multicast RPF Check

Tail-site Customer

Transit AS109




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

59

Single-Homed, Customer RP, MBGP

Cisco.com



Note: Access-list 111 = Recommended SA Filter


RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

60

Single-Homed, Customer RP, MBGP

Cisco.com



```
int pos0/0
ip pim sparse-mode
ip pim bsr-border
ip multicast boundary 1

ip msdp sa-filter out 4.4.4.4 111
ip msdp sa-filter in 4.4.4.4 111
```

RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

61

Single-Homed, Customer RP, MBGP

Cisco.com

MSDP RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

62

Single-Homed, Customer RP, MBGP

Cisco.com

Multicast RPF Check




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

63

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

64

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

65

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast


Cisco.com



Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com

MSDP RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

67

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com

Multicast RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

68

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com

Multicast RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

69

Dual-Homed, Customer RP, MBGP Incongruent Multicast—Unicast

Cisco.com

Multicast RPF Check




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

70

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

71

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com




RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

72

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast


Cisco.com



Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com

MSDP RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

74

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com

Multicast RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

75

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com

Multicast RPF Check



RST-1701
9817_05_2004_X


© 2004 Cisco Systems, Inc. All rights reserved.

76

Dual-Homed, Customer RP, MBGP Congruent Multicast—Unicast

Cisco.com

Multicast RPF Check



RSF-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

77

GLOP—Static Allocation of 233/8

Cisco.com

- **Temporary allocation of 233/8**
 - RFC 2770
- **Statically assigned by mapping AS number into middle octets**
 - <http://gigapop.uoregon.edu/glop/index.html>
- **Provides each AS with /24 addresses to use while waiting another solution**

RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

78

GLOP—Static Allocation of 233/8

Cisco.com

- The hexadecimal value of 5662 is 161E. 16 hex equals 22 decimal and 1E hex equals 30 decimal. We get 233.22.30.0/24.
- The lazy (smart?) way to calculate your GLOP address space is by entering it here and it will calculate it for you:

<http://www.ogig.net/glop/>


RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

79

Wonderful Bedtime Stories

Cisco.com



RST-1701
9817_05_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

80

CISCO SYSTEMS

