NETWORKERS 2004

Deploying IP Multicast

Session RST-2701

RST-2701
9799_05_2004_X

1

Networkers Multicast Sessions

Cisco.com

- Breakout Sessions
  - RST 1701 – Introduction to IP Multicast
  - RST 2701 – Deploying IP Multicast
  - RST 2702 – Deploying IP Multicast VPN's
  - RST 4701 – Advanced IP Multicast
- Techtorials
  - RST 2T07 – Enterprise IP Multicast
- Multicast BoF

RST-2701
9799_05_2004_X

2

Agenda

Geekometer

Cisco.com

- **Basic Multicast Engineering**
  - PIM Configuration Steps
  - Which Mode: Sparse or Dense?
  - Basic RP Engineering
- **Advanced Multicast Engineering**
  - PIM Protocol Extensions
  - Combining Auto-RP and Anycast RP
  - Multicast Group Control
  - Using Admin. Scoped Zones

Basic Multicast Engineering

Basic Multicast Configuration – PIM Configuration Steps

RST-2701
9799_05_2004_X

PIM Configuration Steps

Cisco.com

- Enable Multicast Routing on ***every*** router
- Configure ***every*** interface for PIM
- Configure the RP
  - Using Auto-RP or BSR
    - Configure certain routers as Candidate RP(s)
    - All other routers automatically learn elected RP
  - Anycast/Static RP addressing
    - RP address must be configured on every router
    - Note: Anycast RP requires MSDP

RST-2701
9799_05_2004_X

Configure PIM on Every Interface

Cisco.com

Classic Partial Multicast Cloud Mistake #1

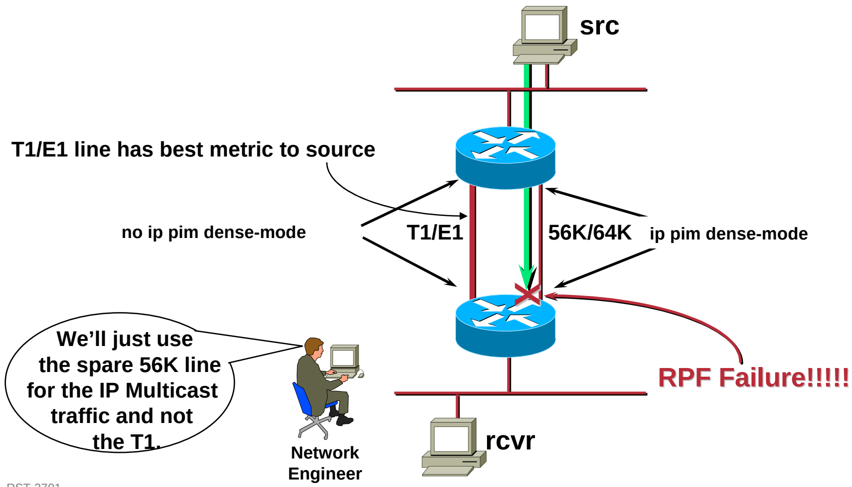

RST-2701
9799_05_2004_X

7

Configure PIM on Every Router

Cisco.com

Classic Partial Multicast Cloud Mistake #2

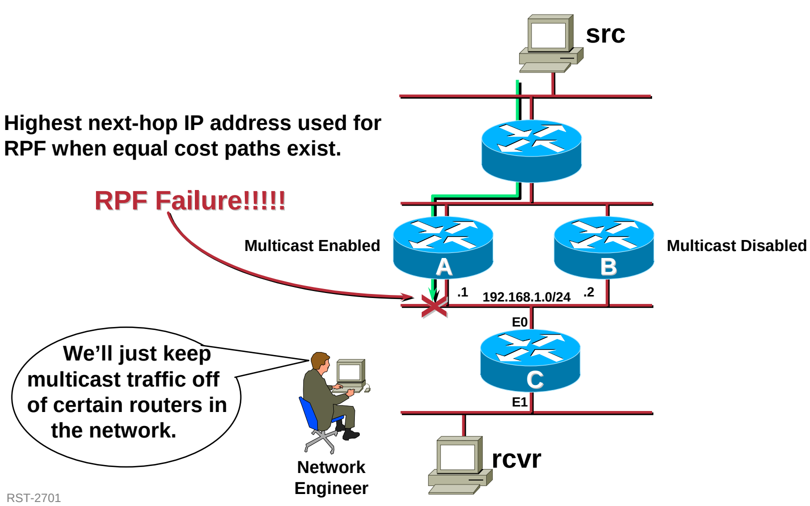

RST-2701
9799_05_2004_X

8

Which Mode: Sparse or Dense?

Which Mode—Sparse or Dense

Cisco.com

- **Dense mode**
  - **Flood and Prune behavior very inefficient**
    - **Can cause problems in certain network topologies**
  - **Creates (S, G) state in EVERY router**
    - **Even when there are no receivers for the traffic**
  - **Complex Assert mechanism**
  - **Mixed control and data planes**
    - **Results in (S, G) state in every router in the network**
    - **Can result in non-deterministic topological behavior**
      Read: It can black-hole traffic and/or melt down your network!
  - **Primarily usage:**
    - **Testing a router's performance in the lab**

Which Mode—Sparse or Dense

Cisco.com

- Sparse mode
  - Must configure a Rendezvous Point (RP)
  - Very efficient
    - Uses Explicit Join model
    - Traffic only flows to where it's needed
  - Separated control and data planes
    - Router state only created along flow paths
    - Deterministic topological behavior
  - Scales well
    - Works for both sparsely or densely populated networks

RST-2701
9799_05_2004_X 11

CONCLUSION

"Sparse mode Good! Dense mode Bad!"

Source: "The Caveman's Guide to IP Multicast", ©2000, R. Davis

RST-2701 12

Group Mode vs. Interface Mode

Cisco.com

- Group & Interface mode are independent.
  - Interface Mode
    - Determines how the interface operates when sending/receiving multicast traffic.
  - Group Mode
    - Determines whether the group is Sparse or Dense.

RST-2701
9799_05_2004_X

13

Group Mode

Cisco.com

- Group mode is controlled by local RP info
  - Local RP Information
    - Stored in the Group-to-RP Mapping Cache
    - May be statically configured or learned via Auto-RP or BSR
  - If RP info exists, Group = Sparse
  - If RP info does not exist, Group = Dense
  - Mode Changes are automatic.
    - i.e. if RP info is lost, Group falls back to Dense.

RST-2701
9799_05_2004_X

14

Configuring Interface

Cisco.com

- Interface Mode Configuration Commands
  - Enables multicast forwarding on the interface.
  - Controls the *interface's* mode of operation.

`ip pim dense-mode`

- Interface mode is set to Dense mode operation.

`ip pim sparse-mode`

- Interface mode is set to Sparse mode operation.

`ip pim sparse-dense-mode`

- Interface mode is determined by the Group mode.
  - If Group is Dense, interface operates in Dense mode.
  - If Group is Sparse, interface operates in Sparse mode.

RST-2701
9799_05_2004_X

15

Basic RP Engineering –
RP Configuration Methods

RST-2701
9799_05_2004_X

16

RP Configuration Methods

Cisco.com

- Static
- Auto-RP
- BSR
- Anycast-RP’s

RST-2701
9799_05_2004_X

17

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. (More on that later.)
  - Group can never fall back into Dense mode.

RST-2701
9799_05_2004_X

18

Auto-RP Overview

Cisco.com

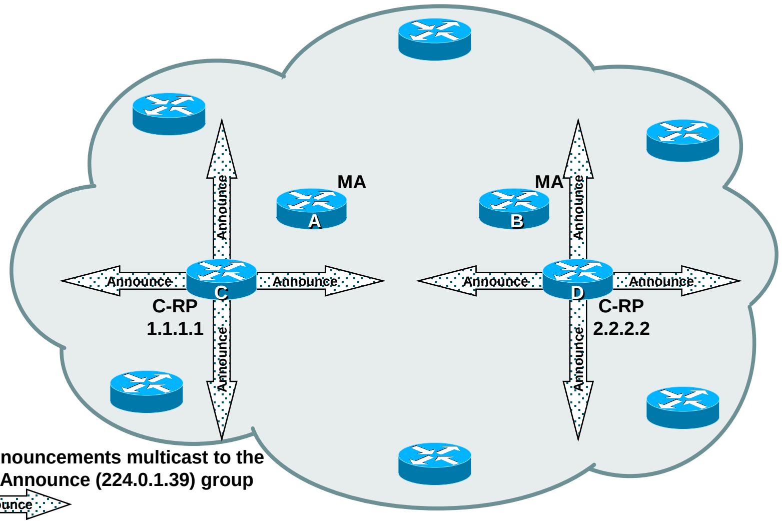

RP-Announcements multicast to the Cisco Announce (224.0.1.39) group

Announce

RST-2701
9799_05_2004_X

19

Auto-RP Overview

Cisco.com

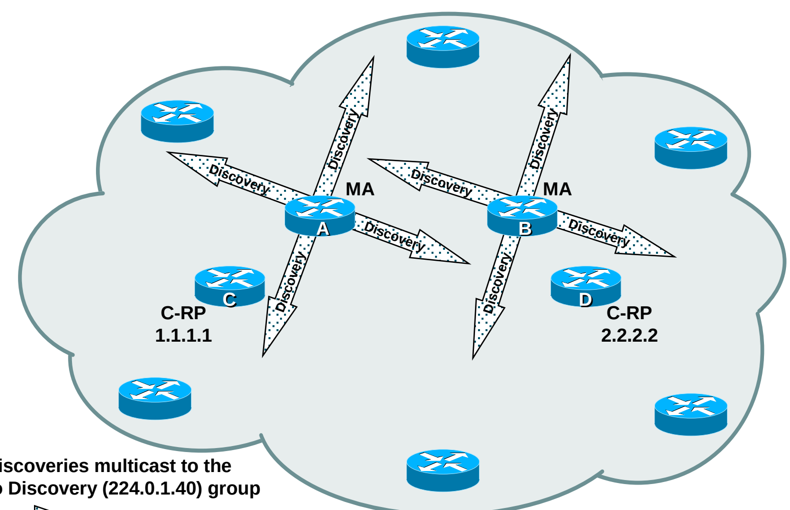

RP-Discoveries multicast to the Cisco Discovery (224.0.1.40) group

Discovery

RST-2701
9799_05_2004_X

20

BSR Overview

Cisco.com

BSR Election Process

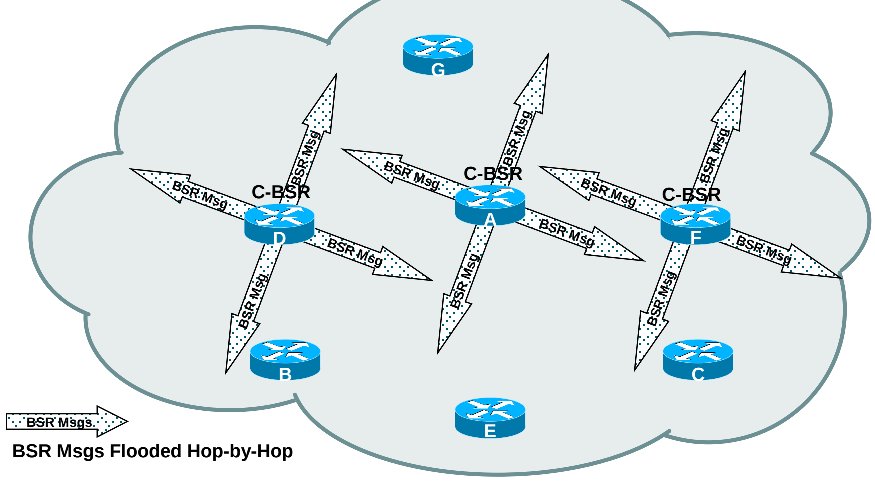

BSR Msgs Flooded Hop-by-Hop

RST-2701
9799_05_2004_X

21

BSR Overview

Cisco.com

Highest Priority C-BSR
is elected as BSR

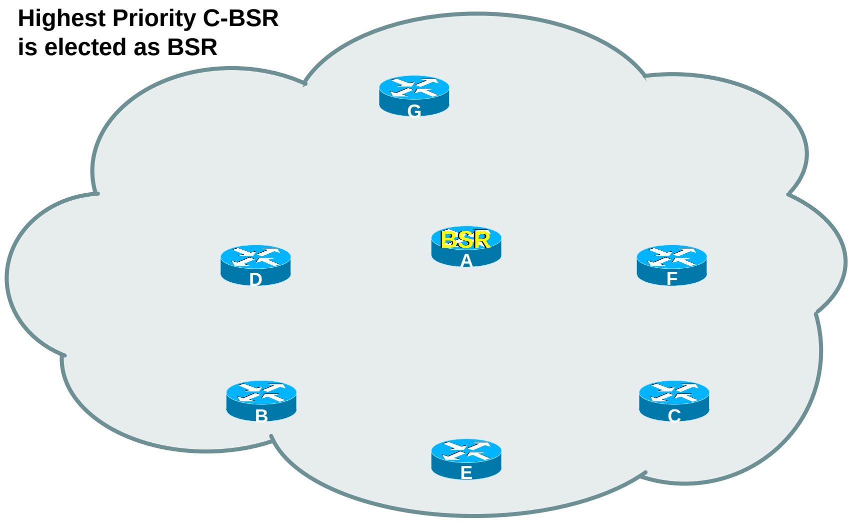

RST-2701
9799_05_2004_X

22

BSR Overview

Cisco.com

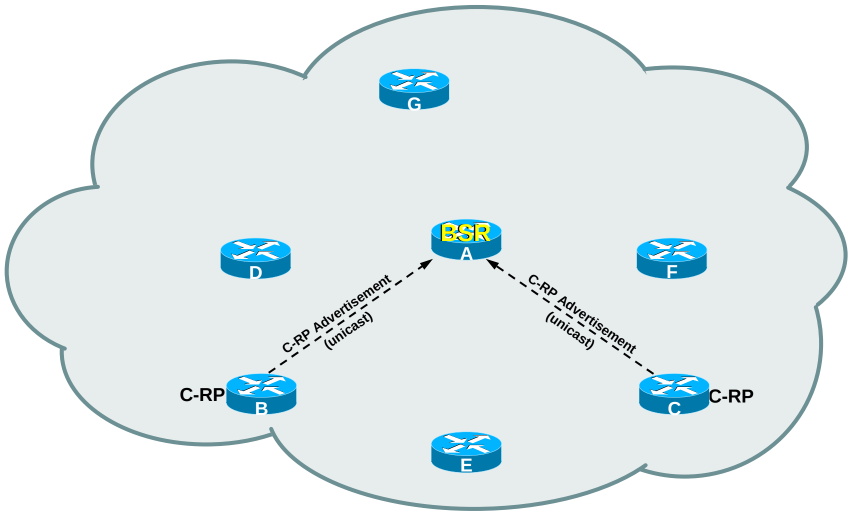

RST-2701
9799_05_2004_X

23

BSR Overview

Cisco.com

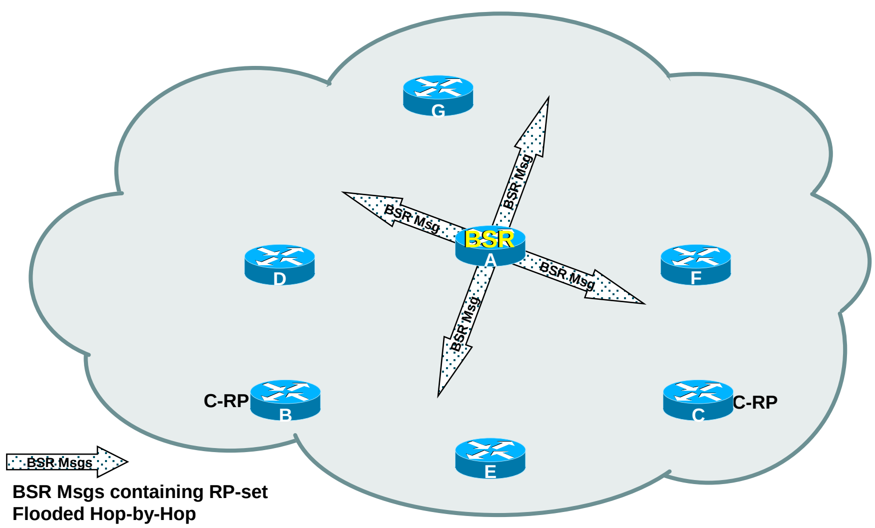

BSR Msgs containing RP-set Flooded Hop-by-Hop

RST-2701
9799_05_2004_X

24

Anycast RP—Overview

Cisco.com

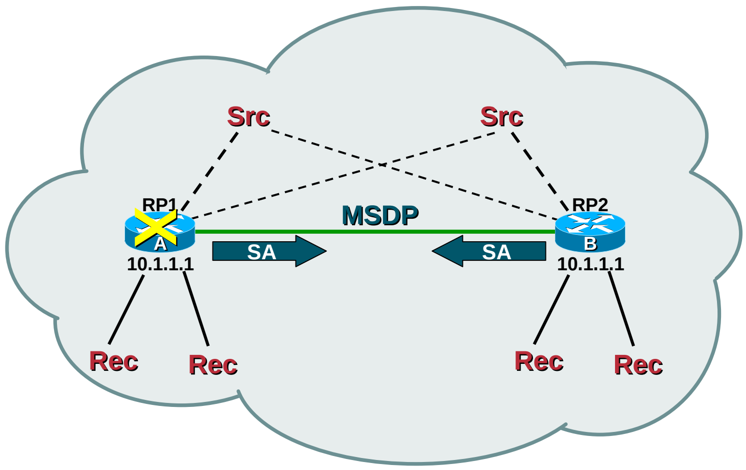

RST-2701
9799_05_2004_X

25

Anycast RP—Overview

Cisco.com

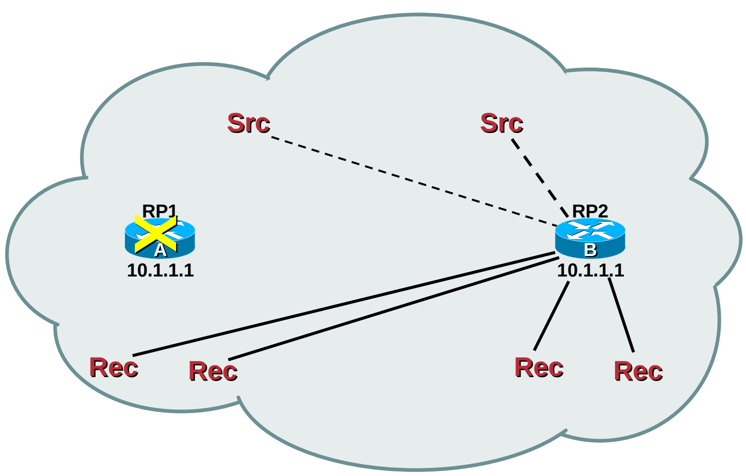

RST-2701
9799_05_2004_X

26

Basic RP Engineering –
General RP Recommendations

General RP Recommendations

- Use Anycast RP's:
  – When network must connect to Internet or
  – When rapid RP failover is critical
- Pros
  – Fastest RP Convergence method
  – Required when connecting to Internet
- Cons
  – Requires more configuration
  – Requires use of MSDP between RP's

General RP Recommendations

Cisco.com

- Use Auto-RP
  - When minimum configuration is desired and/or
  - When maximum flexibility is desired
- Pros
  - Most flexible method
  - Easiest to maintain
- Cons
  - Increased RP Failover times vs Anycast
  - Special care needed to avoid DM Fallback
    - Some methods greatly increase configuration

RST-2701
9799_05_2004_X

29

General RP Recommendations

Cisco.com

- Use BSR:
  - When Static/Anycast RP's cannot be used and
  - When maximum interoperability is needed
- Pros
  - Interoperates with all Vendors
- Cons
  - Increased RP Failover times vs Anycast
  - Special care needed to avoid DM Fallback
    - Some methods greatly increase configuration
  - Does not support Admin. Scoping

RST-2701
9799_05_2004_X

30

Basic RP Engineering –
Avoiding Dense Mode Fallback

RST-2701
9799_05_2004_X

31

Dense Mode Fallback

Cisco.com

- Caused by loss of local RP information in older IOS releases.
  - Entry in Group-to-RP mapping cache times out.
- Can happen when:
  - All C-RP’s fail.
  - Auto-RP/BSR mechanism fails.
    - Generally a result of network congestion.
- Group is switched over to Dense mode.
  - Dense mode state is created in the network.
  - Dense mode flooding begins if interfaces configured as **ip pim sparse-dense-mode**.

RST-2701
9799_05_2004_X

32

Dense Mode Fallback

Cisco.com

Avoiding Dense Mode Fallback

To always guarantee Sparse mode operation (and avoid falling back to Dense mode), make sure that every router *always* knows of an RP for every group.

RST-2701
9799_05_2004_X

33

Avoiding DM Fallback – Old Workaround

Cisco.com

- Define an “RP-of-last-resort”
  - Configure as a Static RP on every router
    - Will only be used if all Candidate-RP’s fail
    - Can be a dummy address or local Loopback
      - Recommendation: Use local Loopback on each router
  - *MUST use ACL to avoid breaking Auto-RP!*

```
ip pim rp-address <RP-of-last-resort> 10
access-list 10 deny 224.0.1.39
access-list 10 deny 224.0.1.40
access-list 10 permit any
```

RST-2701
9799_05_2004_X

34

Avoiding DM Flooding

Cisco.com

- New IOS global command
  `ip pim autorp-listener`
- Added support for Auto-RP Environments
  - Modifies interface behavior
    - Interface always uses DM for Auto-RP groups
    - Permits use of ip pim sparse-mode interfaces and Auto-RP.
  - Prevents DM Flooding
    - When ip pim sparse-mode used on interfaces.
  - Does not prevent DM Fallback!
- Available 12.3(4)T, 12.2(28)S

RST-2701
9799_05_2004_X 35

Avoiding DM Flooding

Cisco.com

- Deploying ip pim autorp-listener
  - Must be configured on every router.
  - Use RP-of-last-resort on older IOS versions until upgraded
    - Assign local Loopback as RP-of-last-resort on each router.
    - Example

```
ip pim rp-address <local_loopback> 10
access-list 10 deny 224.0.1.39
access-list 10 deny 224.0.1.40
access-list 10 permit any
```

RST-2701
9799_05_2004_X 36

Avoiding DM *Fallback*

Cisco.com

- **New IOS global command**

```
no ip pim dm-fallback
```

- **Totally prevents DM Fallback!!**
  - **No DM Flooding since all state remains in SM**
- **Default RP Address = 0.0.0.0 [nonexistent]**
  - **Used if all RP’s fail.**
    - **Results in loss of Shared Tree.**
    - **All SPT’s remain active.**
- **Available 12.3(4)T, 12.2(28)S**

Advanced Multicast Engineering

PIM Protocol Extensions

RST-2701
9799_05_2004_X

Source Specific Multicast

RST-2701
9799_05_2004_X

0863_04F9_c2.scr

Barriers to Multicast Deployment

Cisco.com

- Global Multicast Address Allocation
  - Dynamic Address Allocation
    - No adequate dynamic address allocation methods exist
    - SDR – Doesn't scale
    - MASC – Long ways off!
  - Static Address Allocation (GLOP)
    - Based on AS number.
    - Insufficient address space for large Content Providers.
- Multicast Content "Jammers"
  - Undesirable sources on a multicast group.
    - "Capt. Midnight" sources bogus data/noise to group.
    - Can cause DoS attack by congesting low speed links.

RST-2701
9799_05_2004_X 41

Source Specific Multicast (SSM)

Cisco.com

- Uses Source Trees only.
- Assumes One-to-Many model.
  - Most Internet multicast fits this model.
  - IP/TV also fits this model.
- Hosts responsible for source discovery.
  - Typically via some out-of-band mechanism.
    - Web page, Content Server, etc.
  - Eliminates need for RP and Shared Trees.
  - Eliminates need for MSDP.

RST-2701
9799_05_2004_X 42

SSM Overview

Cisco.com

- Hosts join a specific source within a group.
  - Content identified by specific (S,G) instead of (*,G).
  - Hosts responsible for learning (S,G) information.
- Last-hop router sends (S,G) join toward source
  - Shared Tree is never Joined or used.
  - Eliminates possibility of content Jammers.
  - Only specified (S,G) flow is delivered to host.
- Simplifies address allocation.
  - Dissimilar content sources can use same group without fear of interfering with each other.

RTS-2701
9799_05_2004_X 43

SSM Example

Cisco.com

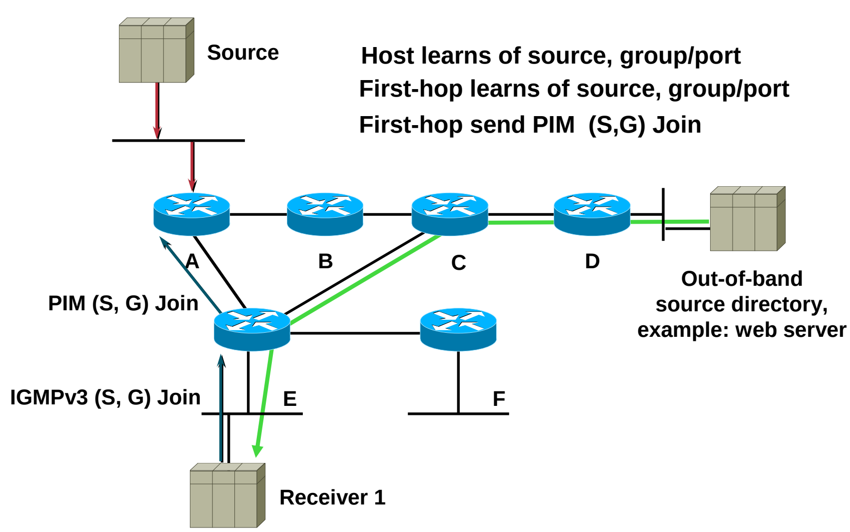

RTS-2701
9799_05_2004_X 44

SSM Example

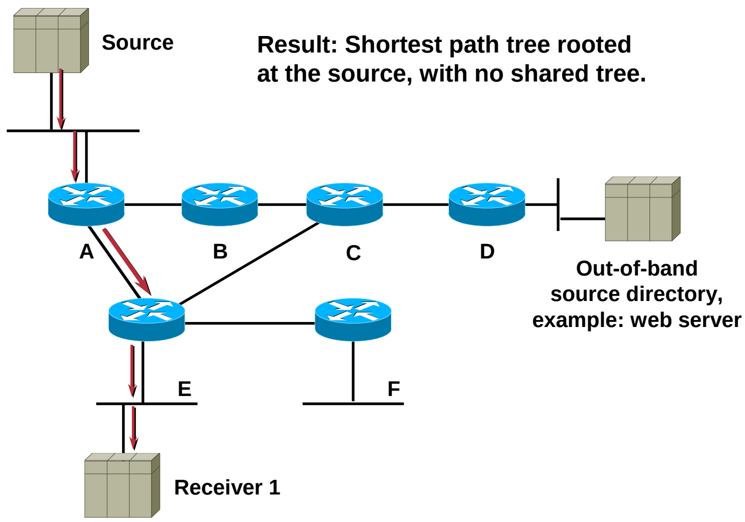

SSM Configuration

- **Global command**

  `ip pim ssm {default | <acl>}`

  - Defines SSM address range
    - Default range = 232.0.0.0/8
    - Use ACL for other ranges
  - Prevents Shared Tree Creation
    - (*, G) Joins never sent or processed
    - PIM Registers never sent or processed
  - Available in IOS versions
    - 12.1(5)T, 12.2, 12.0(15)S, 12.1(8)E

SSM – Summary

Cisco.com

- Uses Source Trees only.
  - Hosts are responsible for source & group discovery.
  - Hosts must signal router which (S,G) to join.
- Solves multicast address allocation problems.
  - Flows differentiated by *both* source and group.
  - Content providers can use same group ranges.
    - Since each (S,G) flow is unique.
- Helps prevent certain DoS attacks
  - “Bogus” source traffic:
    - Can’t consume network bandwidth.
    - Not received by host application.

RST-2701
9799_05_2004_X 47

So where is SSM?

Cisco.com

- Dependant on IGMPv3
  - Microsoft supports IGMPv3 in Windows XP
- Workarounds
  - IGMPv3 lite
    - API/Library/DLL
    - Used by Cisco IP/TV 3.2 and later.
  - URL RenDezvous (URD)
    - Redirect from Web page with specific information intercepted by Router
  - Static Source Mapping
    - Router maps IGMPv2 Joins in SSM range to well-known sources via DNS or static configuration

RST-2701
9799_05_2004_X 48

IGMP v3lite

Cisco.com

- Source side:
  - No application changes required!
- Receiver side:
  - Application must use IGMPv3 API:
  - IGMP v3lite Library Component
    - Provides the IP SSM subset of IGMPv3 API
      - Applications must still filter out unwanted traffic.
  - IGMP v3lite Daemon Component
    - Sends special (S,G) Join to local router via UDP port 465

RST-2701
9799_05_2004_X 49

URD

Cisco.com

- A content provider builds a web page that contains URD links.
  - List of sources willing to provide multicast content
- The user (receiver) clicks on one of the links
- Web Server sends back an HTTP redirect containing source and group info to TCP port 465
- Host sends the redirect via TCP port 465
- Local router intercepts TCP port 465 traffic
  - Uses source/group information in the redirect to identify the requested SSM flow.

RST-2701
9799_05_2004_X 50

SSM Mapping

Cisco.com

- Allows only for one source per Group
- Router maps group to a single source
  - Uses either DNS or static internal database
    - DNS method allows content providers to provide the mapping
    - DNS Method independent from network operators

RST-2701
9799_05_2004_X

51

SSM Mapping – DNS Example

Cisco.com

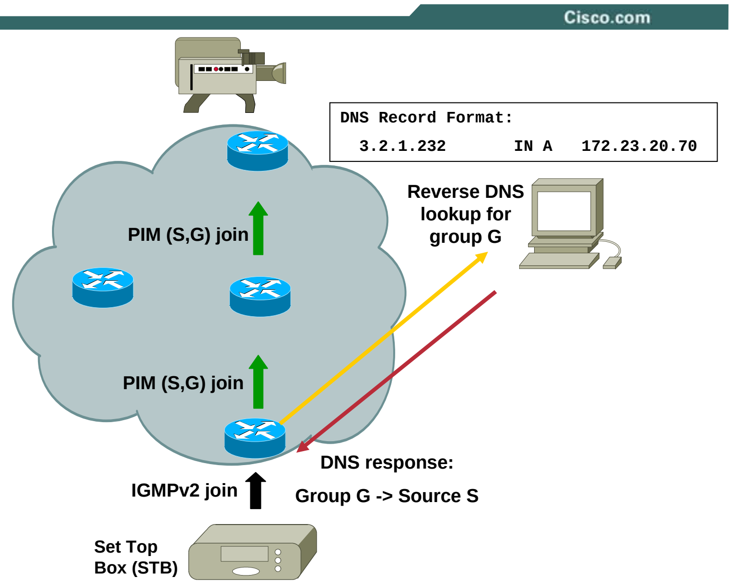

RST-2701
9799_05_2004_X

52

0863_04F9_c2.scr

SSM Mapping Configuration

Cisco.com

Enabling SSM mapping on the router

```
ip igmp ssm-map enable
```

For static mapping:

```
ip igmp ssm-map static <acl-1> <source-1 IP address>
ip igmp ssm-map static <acl-2> <source-2 IP address>
```

For DNS mapping (existing commands):

```
ip domain-server <ip address>
ip domain-name <domain.com>
```

To disable DNS mapping

```
no ip igmp ssm-map query dns
```

```
DNS Record Format:   3.2.1.232       IN A   172.23.20.70
```


Bidirectional (Bidir) PIM

Multicast Application Categories

- One-to-Many Applications
  - Video, TV, Radio, Concerts, Stock Ticker, etc.
- Few-to-Few Applications
  - Small (<10 member) Video/Audio Conferences
- Few-to-Many Applications
  - TIBCO RV Servers (Publishing)
- Many-to-Many Applications
  - Stock Trading Floors, Gaming
- Many-to-Few Applications
  - TIBCO RV Clients (Subscriptions)

Multicast Application Categories
PIM-SM (S, G) State

- One-to-Many Applications
  - Single (S,G) entry
- Few-to-Few Applications
  - Few (<10 typical) (S,G) entries
- Few-to-Many Applications
  - Few (<10 typical) (S,G) entries
- Many-to-Many Applications
  - Unlimited (S,G) entries
- Many-to-Few Applications
  - Unlimited (S,G) entries

Multicast State Maintenance

Cisco.com

- CPU load factors
  - Must send/receive Registers
  - Must send periodic Joins/Prunes
  - Must perform RPF recalculations
    - Watch the total number of mroute table entries
    - Unicast route table size impacts RPF recalculation
- Memory load factors
  - (*, G) entry ~ 380 bytes + OIL size
  - (S, G) entry ~ 220 bytes + OIL size
  - Outgoing interface list (OIL) size
    - Each oil entry ~ 150 bytes

RST-2701
9799_05_2004_X

57

Many-to-Any State Problem

Cisco.com

- Creates huge amounts of (S,G) state
  - State maintenance workloads skyrocket
    - High OIL fanouts make the problem worse
  - Router performance begins to suffer
- Using Shared-Trees only
  - Provides some (S,G) state reduction
    - Results in (S,G) state only along SPT to RP
    - Frequently still too much (S,G) state
    - Need a solution that only uses (*,G) state

RST-2701
9799_05_2004_X

58

Bidirectional (Bidir) PIM

Cisco.com

- Idea:
  - Use the same tree for traffic from sources towards RP and from RP to receivers
- Benefits:
  - Less state in routers
    - Only (*, G) state is used
    - Source traffic follows the Shared Tree
      - Flows up the Shared Tree to reach the RP.
      - Flows down the Shared Tree to reach all other receivers.

RST-2701
9799_05_2004_X

59

Bidirectional (Bidir) PIM

Cisco.com

- Bidirectional Shared-Trees
  - Violates current (*,G) RPF rules
    - Traffic often accepted on ***outgoing*** interfaces.
    - Care must be taken to avoid multicast loops
  - Requires a Designated Forwarder (DF)
    - Responsible for forwarding traffic up Shared Tree
      - DF's will accept data on the interfaces in their OIL.
      - Then send it out all other interfaces. (Including the IIF.)

RST-2701
9799_05_2004_X

60

Bidirectional (Bidir) PIM

Cisco.com

- Designated Forwarders (DF)
  - On each link the router with the best path to the RP is elected to be the DF
    - Note: Designated Routers (DR) are not used for bidir groups
  - The DF is responsible for forwarding traffic upstream towards the RP
  - No special treatment is required for local sources

RST-2701
9799_05_2004_X

61

Bidirectional PIM — Example

Cisco.com

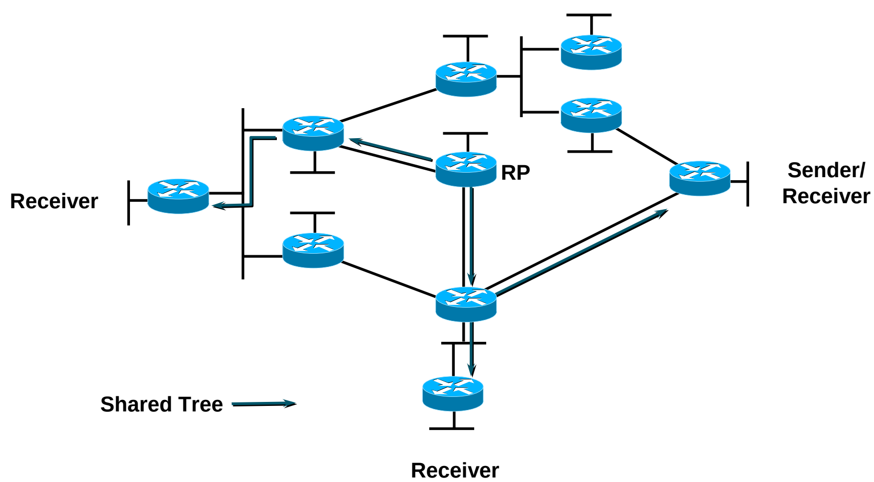

RST-2701
9799_05_2004_X

62

Bidirectional PIM — Example

Cisco.com

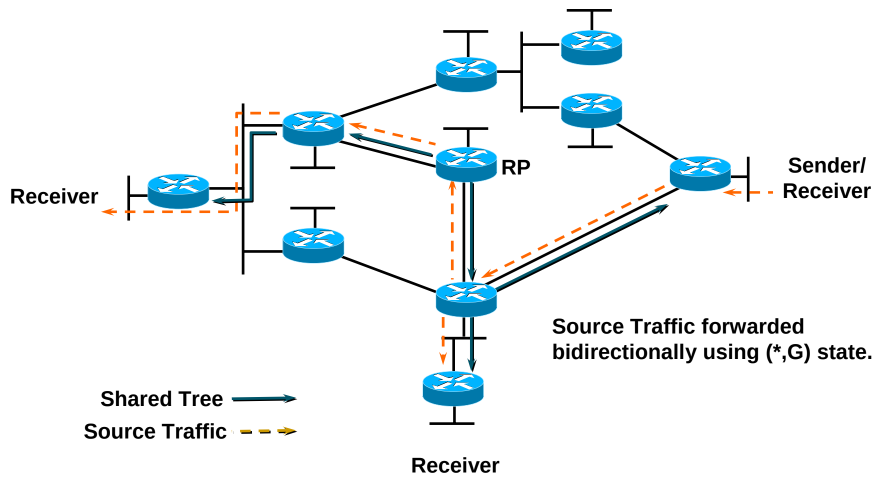

RST-2701
9799_05_2004_X

63

Configuring Bidir PIM
(Auto-RP Example)

Cisco.com

- **Define Candidate RP and groups / modes it is willing to serve**

```
ip pim send-rp-announce Loopback0 scope 10 group-list 45 bidir
ip pim send-rp-announce Loopback1 scope 10 group-list 46
! Two loopbacks needed due to a nature of ACLs (permit, deny)
ip pim send-rp-discovery scope 10

access-list 45 permit 224.0.0.0 0.255.255.255
access-list 45 permit 227.0.0.0 0.255.255.255
! 224/8 and 227/8 will be PIM Bidir groups
access-list 45 deny 225.0.0.0 0.255.255.255
! 225/8 will be a PIM Dense Mode group

access-list 46 permit 226.0.0.0 0.255.255.255
! 226/8 will be a PIM Sparse Mode group
```

RST-2701
9799_05_2004_X

64

Bidir PIM – Phantom RP

Cisco.com

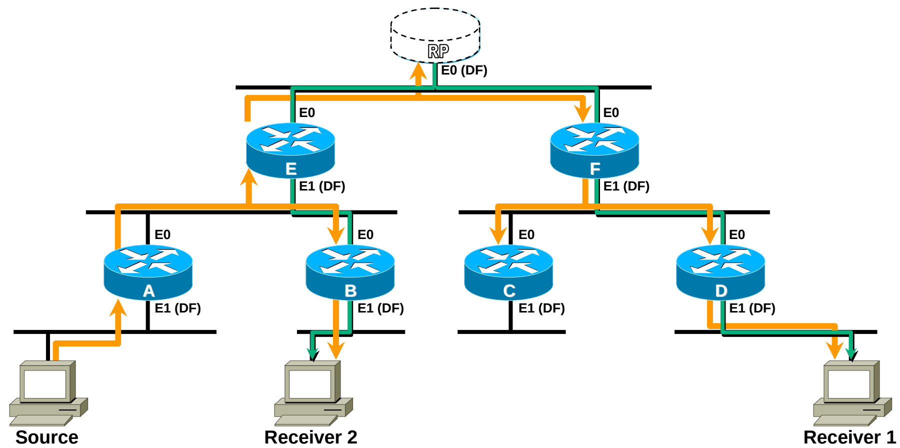

Question: Does a Bidir RP even have to physically exist?
Answer: No. It can just be a phantom address.

RST-2701
9799_05_2004_X

65

Bidir PIM – Phantom RP

Cisco.com

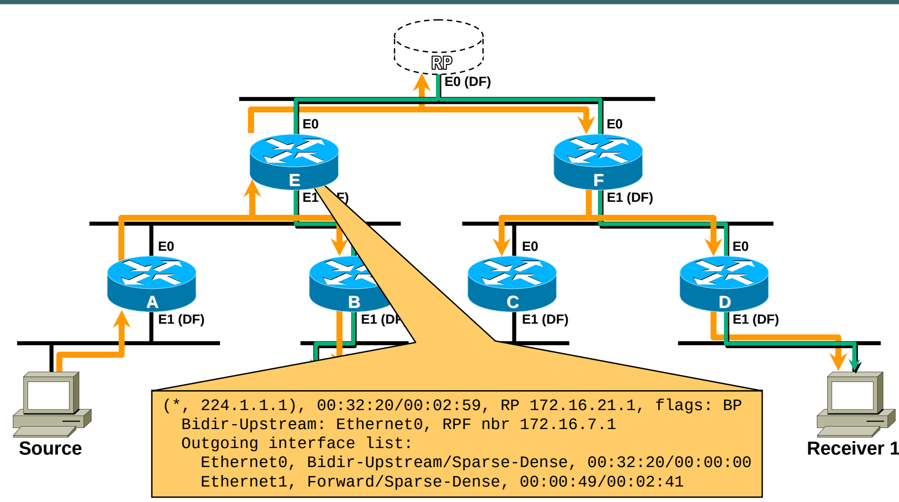


```
(*, 224.1.1.1), 00:32:20/00:02:59, RP 172.16.21.1, flags: BP
  Bidir-Upstream: Ethernet0, RPF nbr 172.16.7.1
  Outgoing interface list:
    Ethernet0, Bidir-Upstream/Sparse-Dense, 00:32:20/00:00:00
    Ethernet1, Forward/Sparse-Dense, 00:00:49/00:02:41
```

Router “E” forwards traffic onto core LAN segment.

RST-2701
9799_05_2004_X

66

Bidir PIM – Phantom RP

Cisco.com

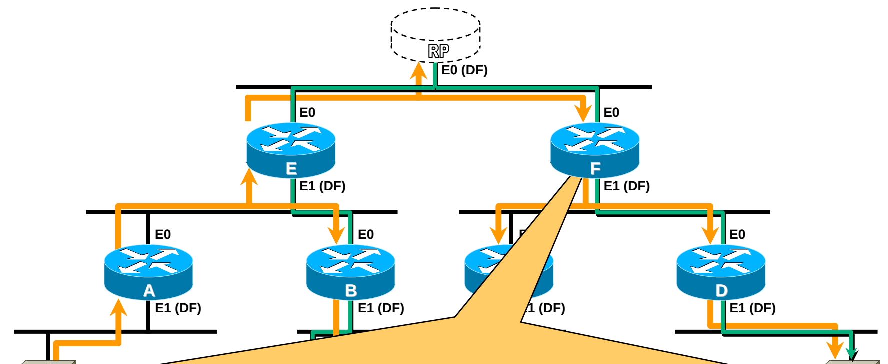


```
(*, 224.1.1.1), 00:00:49/00:02:41, RP 172.16.21.1, flags: B
  Bidir-Upstream: Ethernet0, RPF nbr 172.16.1.1
  Outgoing interface list:
    Ethernet0, Bidir-Upstream/Sparse-Dense, 00:00:49/00:00:00
    Ethernet1, Forward/Sparse-Dense, 00:00:49/00:02:41
```

Router "F" forwards traffic on down the Shared Tree ala normal PIM-SM.

RP doesn't even have to physically exist.

Phantom RP on Point-to-Point Core

Cisco.com

Static Route Method

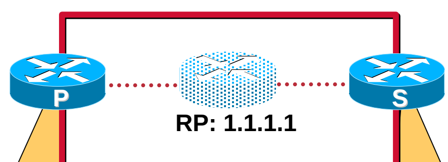


```
ip multicast-routing

interface Loopback0
 ip address 11.0.0.1 255.255.255.255
 ip pim sparse-mode

router ospf 11
 redistribute static subnets

ip route 1.1.1.1 255.255.255.255 Loopback0

ip pim bidir-enable
ip pim rp-address 1.1.1.1 bidir
```

```
ip multicast-routing

interface Loopback0
 ip address 11.0.0.2 255.255.255.255
 ip pim sparse-mode

router ospf 11
 redistribute static subnets

ip route 1.1.1.0 255.255.255.254 Loopback0

ip pim bidir-enable
ip pim rp-address 1.1.1.1 bidir
```

Phantom RP on Point-to-Point Core

Cisco.com

Netmask Method

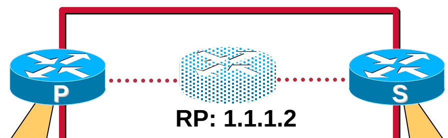


```
ip multicast-routing
!
interface Loopback0
 ip address 1.1.1.1 255.255.255.252
 ip pim sparse-mode
 ip ospf network point-to-point
!
router ospf 11
 network 1.1.1.0 0.0.0.3 area 0
 network 10.1.1.0 0.0.0.255 area 0
 network 10.1.2.0 0.0.0.255 area 0
!
ip pim bidir-enable
ip pim rp-address 1.1.1.1
ip pim rp-address 1.1.1.2 bidir
```

```
ip multicast-routing
!
interface Loopback0
 ip address 1.1.1.1 255.255.255.248
 ip pim sparse-mode
 ip ospf network point-to-point
!
router ospf 11
 network 1.1.1.0 0.0.0.7 area 0
 network 10.1.1.0 0.0.0.255 area 0
 network 10.1.2.0 0.0.0.255 area 0
!
ip pim bidir-enable
ip pim rp-address 1.1.1.1
ip pim rp-address 1.1.1.2 bidir
```

RST-2701
9799_05_2004_X

69

Bidir PIM—Summary

Cisco.com

- Drastically reduces network mroute state
  - Eliminates ***ALL*** (S,G) state in the network
    - SPT's between sources to RP eliminated
    - Source traffic flows both up and down Shared Tree
  - Allows Many-to-Any applications to scale
    - Permits virtually an unlimited number of sources

RST-2701
9799_05_2004_X

70

Multicast Group Control

Controlling Receivers

Cisco.com

IGMP Access-Group Approach

```
interface VLAN100
 ip igmp access-group IPMC-ACL

ip access-list standard IPMC-ACL
 permit 239.192.244.1
 deny any
```

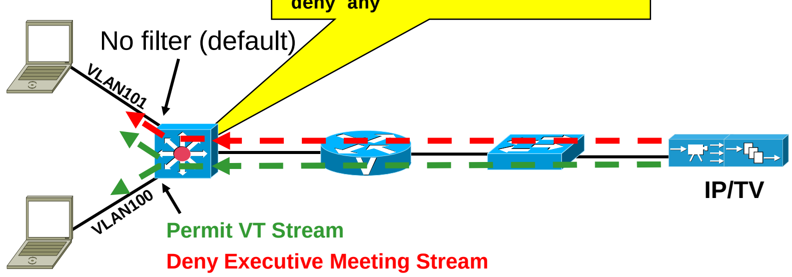

This is micro-management of IP Multicast traffic!!!

Controlling Source Registration

- Global command

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

  – Used on RP to filter incoming Register messages

  – Filter on Source address alone (Simple ACL)

  – Filter on (S, G) pair (Extended ACL)

  – May use route-map to specify what to filter

    - Filter by AS-PATH if (m)BGP is in use.

- Helps prevents unwanted sources from sending

  – First hop router blocks traffic from reaching net

  – Note: Traffic can still flow under certain situations

Controlling Source Registration

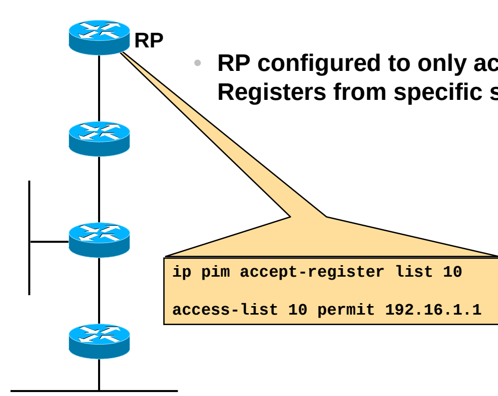

- RP configured to only accept Registers from specific source.

Controlling Source Registration

Cisco.com

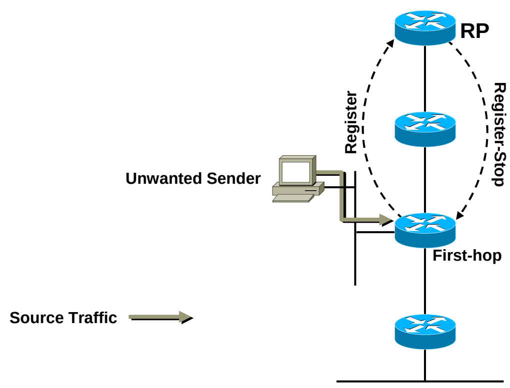

- Unwanted source traffic hits first-hop router.
- First-hop router creates (S,G) state and sends Register.
- RP rejects Register, sends back a Register-Stop.

RST-2701
9799_05_2004_X

75

Controlling Source Registration

Cisco.com

Weaknesses in ‘accept-register’ usage.

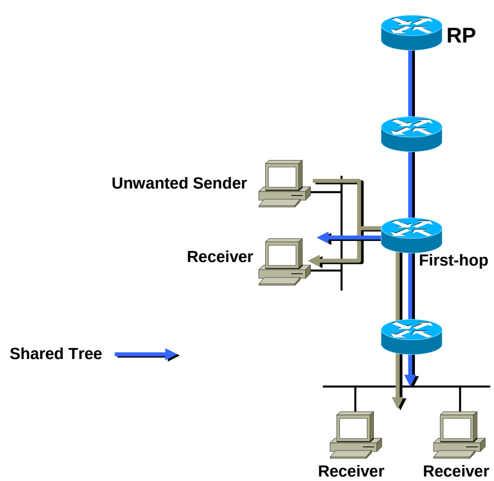

- Traffic will flow on local subnet where source resides.
- Traffic will flow from first-hop router down any branches of the Shared Tree.
  - Results when (*,G) OIL is copied to (S,G) OIL at first-hop router.
  - Causes (S,G) traffic to flow down all interfaces in (*,G) OIL of first-hop router.
  - Fundamental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

RST-2701
9799_05_2004_X

76

Disabling Entire Group Ranges

Cisco.com

- **Accept-Register Method**

```
ip pim accept-register group-list 10
access-list 10 deny 224.2.0.0 0.0.255.255
access-list 10 permit any
```

- **Pros**
  - Only configured on RP(s)
- **Cons**
  - Shared Trees and (*,G) state still created.
    - Results in unwanted (*,G) PIM Control Traffic.
  - Source traffic can still flow.

    (See previous section on Accept-Register)

Disabling Entire Group Ranges

Cisco.com

- **Garbage Can RP Method**
  - Concept:
    - Separate RP for "disabled" groups
      - Could be non-existant router
    - Blackholes all Registers and Joins
  - Implementation:
    - Define separate RP for disabled groups
      - Use Auto-RP, BSR or Static RP definition
    - Disable RP functionality on Garbage Can RP
      - Use 'accept-rp' command on GC RP to "deny" it from serving as RP for the disabled group range.

Disabling Entire Group Ranges

Cisco.com

- Garbage Can RP Method
  - Pros:
    - Few if any.
  - Cons:
    - Periodic Registers still sent to GC RP
    - Periodic Joins still sent to GC RP
    - Has same source issues as Accept-Register
      - Source traffic can still flow under certain conditions.
    - Adds *significant* complexity to network

RST-2701
9799_05_2004_X

79

Disabling Entire Group Ranges

Cisco.com

- Local Loopback RP Method
  - Concept:
    - Only Auto-RP-learned groups are authorized.
    - All other groups are considered *unauthorized*.
  - Implementation:
    - Define local Loopback as RP for unauthorized groups on each router.

```
ip pim rp-address <local_loopback> 10
access-list 10 permit 224.2.0.0 0.0.255.255
```

Note: The permit clause defines the unauthorized group.

RST-2701
9799_05_2004_X

80

Disabling Entire Group Ranges

Cisco.com

- Local Loopback RP Method
  - Operation:
    - Each router serves as RP for unauthorized groups.
      - Collapses PIM-SM domain of unauthorized groups down to the local router.
    - Unauthorized group traffic cannot flow beyond local router.

RST-2701
9799_05_2004_X

81

Disabling Entire Group Ranges

Cisco.com

- Local Loopback RP Method
  - Pros:
    - No PIM control traffic sent.
      - Local router is RP so no Registers/Joins are sent.
    - No additional workload on local router.
      - First-hop routers always have to create state anyway.
    - Can also serve as RP-of-last-resort
      - Solving DM Fallback problem at the same time.
  - Cons:
    - Must be configured on every router.
    - Local sources can still send to local receivers.

RST-2701
9799_05_2004_X

82

Disabling Entire Group Ranges

Cisco.com

- New no ip pim dm-fallback command
  - Groups with no known RP default to an RP address of 0.0.0.0.
    - Effectively disables multicast for these groups.
    - New sources are not Registered.
    - New receivers are not Joined.
- Available 12.3(4)T, 12.2(28)S.

RST-2701
9799_05_2004_X

83

Disabling Entire Group Ranges

Cisco.com

- Recommendations
  - Use no ip pim dm-fallback command
    - Available 12.3(4)T, 12.2(28)S
  - Use Local Loopback RP Method
    - *Effectively* disables unauthorized group traffic.
    - Can also serve as RP-of-last-resort

```
ip pim rp-address <local_loopback> 10
access-list 10 deny 224.0.1.39
access-list 10 deny 224.0.1.40
access-list 10 permit any
```

RST-2701
9799_05_2004_X

84

Combining Anycast RP & Auto-RP

Combining Auto-RP and Anycast-RP

- Anycast-RP and Auto-RP may be combined.
  - Provides advantages of both methods
    - Rapid RP failover of Anycast RP
    - No DM Fallback
    - Configuration flexibility of Auto-RP
    - Ability to effectively disable undesired groups

Combining Auto-RP and Anycast-RP

Cisco.com

Configuration Steps

1. Enable Auto-RP
   – Newer IOS images
     – Use ip pim autorp listener global command and configure ip pim sparse-mode on all interfaces.
   – Older IOS images
     – Configure ip pim sparse-dense-mode on all interfaces.
2. Configure Auto-RP Mapping Agents

```
ip pim send-rp-discovery interface Loopback0 scope 32
```

RST-2701
9799_05_2004_X

Combining Auto-RP and Anycast-RP

Cisco.com

Configuration Steps

3. Block DM Fallback
   – Newer IOS images
     – Use no ip pim dm-fallback
   – Older IOS images
     – Configure RP-of-last-Resort

```
ip pim rp-address <local_loopback> 10
access-list 10 deny 224.0.1.39
access-list 10 deny 224.0.1.40
access-list 10 permit any
```

4. Configure Anycast RP’s for desired group range.
5. Configure Anycast RP’s as Auto-RP C-RP’s

```
ip pim send-rp-discovery Loopback0 scope 32 group-list 10
```

   – Loopback0 = Anycast RP Address
     – Anycast-RP’s will announce Anycast-RP address via Auto-RP

RST-2701
9799_05_2004_X

0863_04F9_c2.scr

Example Auto-RP and Anycast-RP

Cisco.com

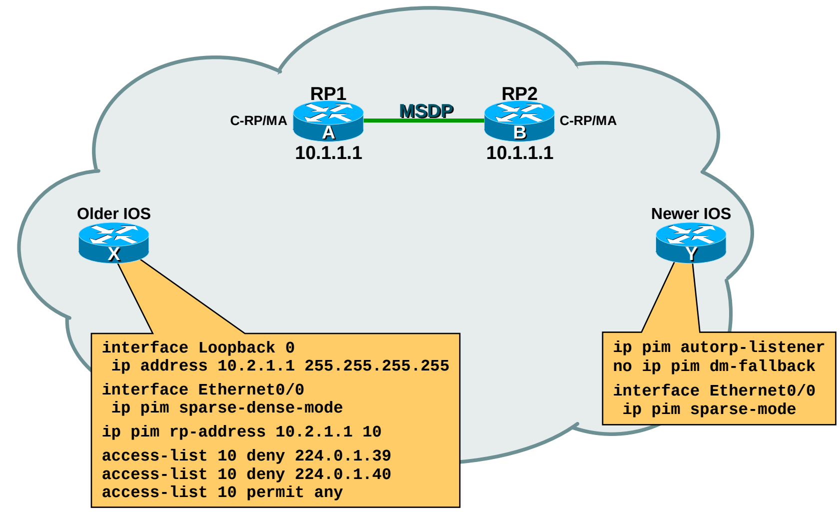

RST-2701
9799_05_2004_X

89

Example Auto-RP and Anycast-RP

Cisco.com

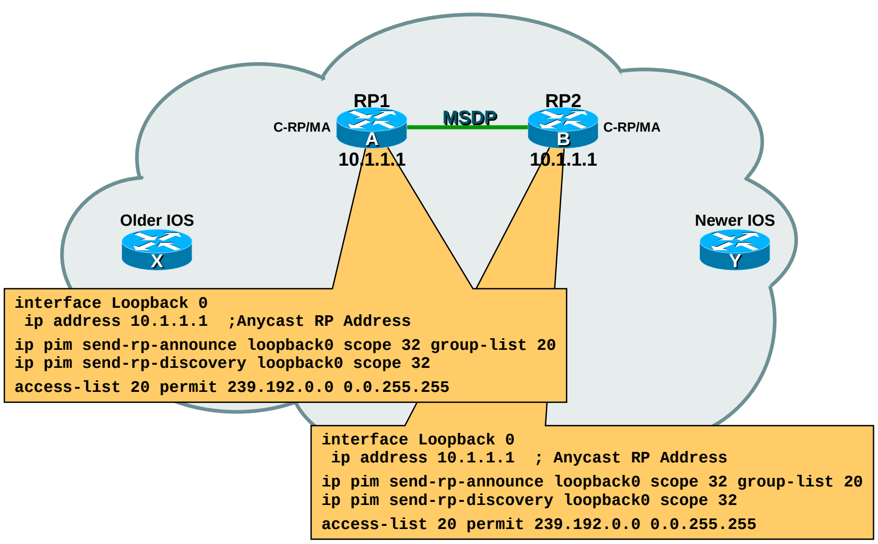

RST-2701
9799_05_2004_X

90

Using Admin. Scoped Zones

RST-2701
9799_05_2004_X

91

Administratively-Scoped Zones

Cisco.com

- Used to limit:
  - High-BW sources to local site
  - Control sensitive multicast traffic
- Simple scoped zone example:
  - 239.193.0.0/16 = Campus Scope
  - 239.194.0.0/16 = Region Scope
  - 239.195.0.0/16 = Organization-Local (Enterprise) Scope
  - 224.1.0.0 - 238.255.255.255 = Global scope (Internet) zone
    - High-BW sources use Site-Local scope
    - Low-Med. BW sources use Org.-Local scope
    - Internet-wide sources use Global scope

RST-2701
9799_05_2004_X

92

Administratively-Scoped Zones Example

Cisco.com

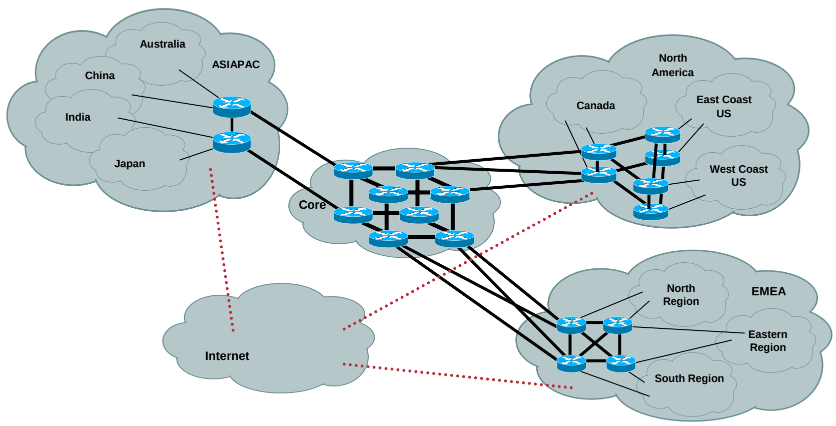

RST-2701
9799_05_2004_X

93

Administratively-Scoped Zones Example

Cisco.com

Level1: Campus Scope

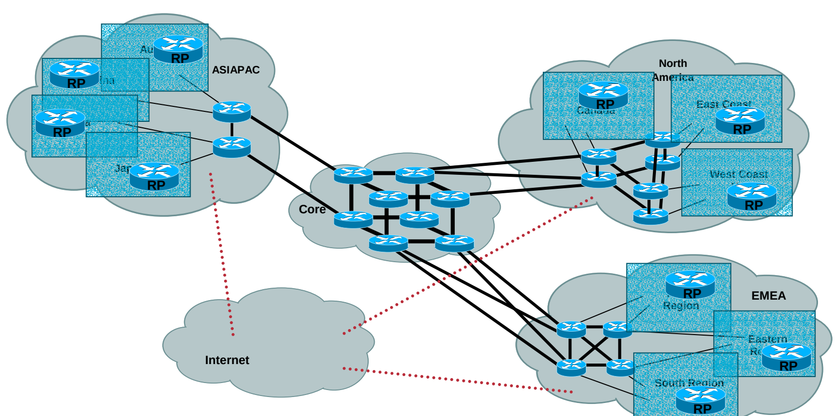

- Campus Scope: 239.193.x.x/16
- RP per Campus

RST-2701
9799_05_2004_X

94

Administratively-Scoped Zones Example

Cisco.com

Level2: Regional Scope

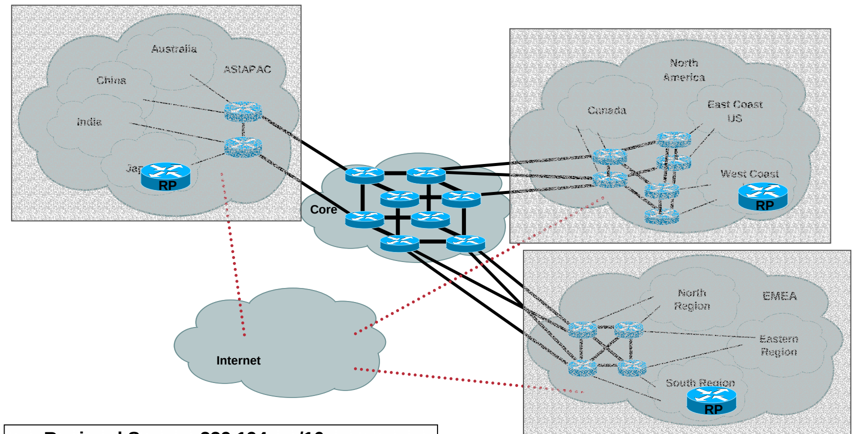

- Regional Scope : 239.194.x.x/16
- RP per Region

RST-2701
9799_05_2004_X

95

Administratively-Scoped Zones Example

Cisco.com

Level3: Enterprise Scope

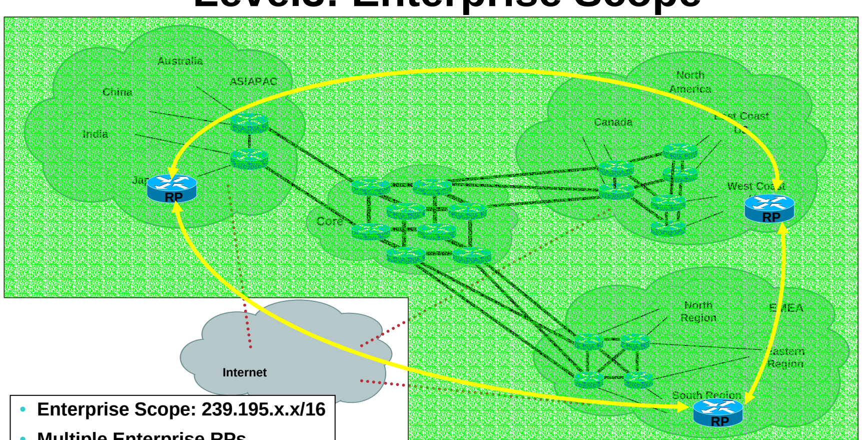

- Enterprise Scope: 239.195.x.x/16
- Multiple Enterprise RPs (via MSDP full mesh)

RST-2701
9799_05_2004_X

96

Administratively-Scoped Zones Example

Cisco.com

Level 4: Internet Global Scope

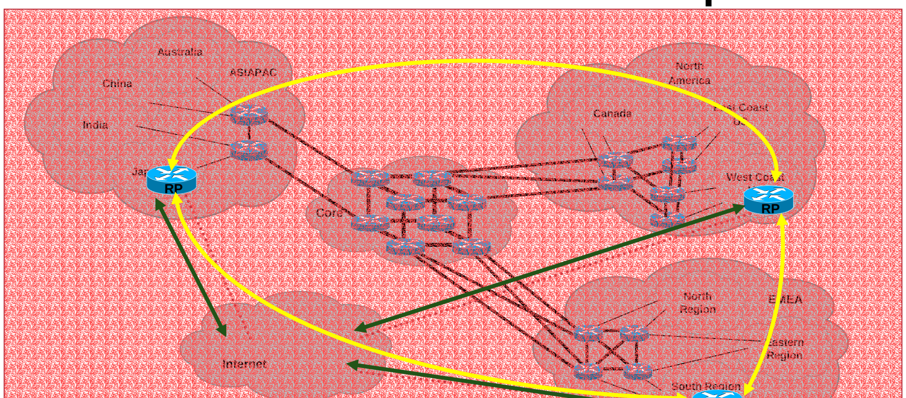

- Global Scope: 224.0.[2-255].x – 238.255.255.255
- Multiple Global RPs (via MSDP full mesh)
- MSDP connectivity to SP network

Administratively Scoped Address Range

Cisco.com

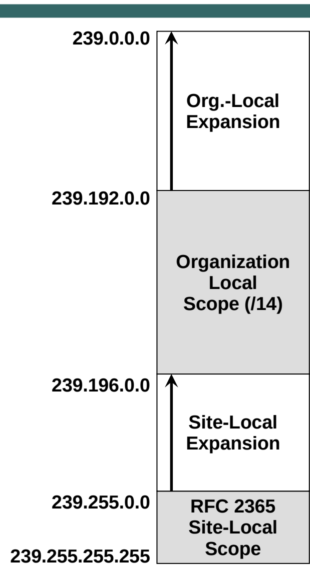

(Not to scale.)

- RFC 2365 Administratively Scoped Zones.
  - Organization-Local Scope (239.192/14)
    - Expands downward in address range.
  - Site-Local Scope (239.255/16)
    - Expands downward in address range.
    - Smallest possible scope.
    - Other scopes may be equal but not smaller.

Example Scope Address Assignments

Cisco.com

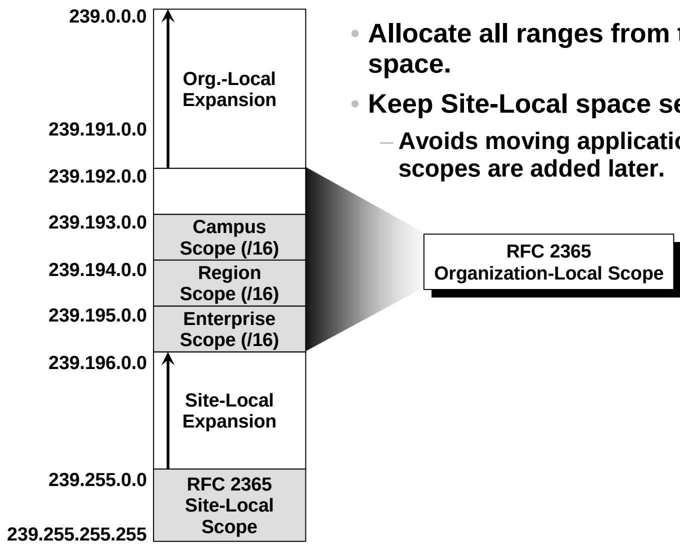

- Allocate all ranges from the Org-Local space.
- Keep Site-Local space separate.
  - Avoids moving applications when smaller scopes are added later.

RST-2701
9799_05_2004_X

99

Adding a Additional Scopes

Cisco.com

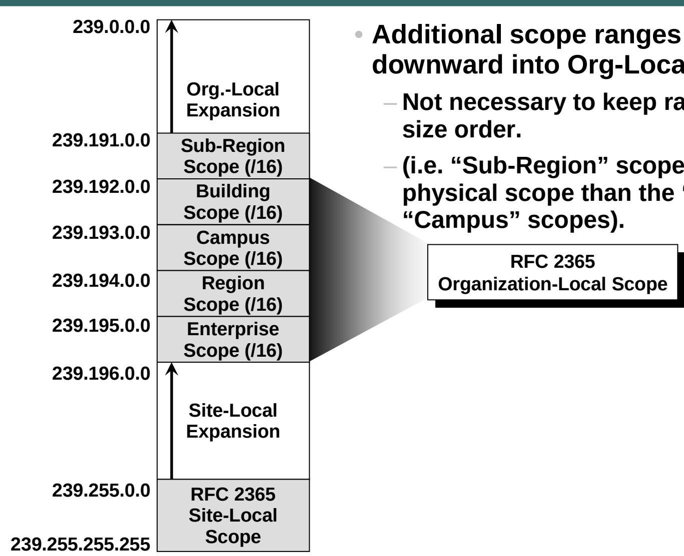

- Additional scope ranges are allocated downward into Org-Local Expansion.
  - Not necessary to keep ranges in scope size order.
  - (i.e. “Sub-Region” scope is a larger physical scope than the “Building” and “Campus” scopes).

RST-2701
9799_05_2004_X

100

Address Ranges to Avoid

Cisco.com

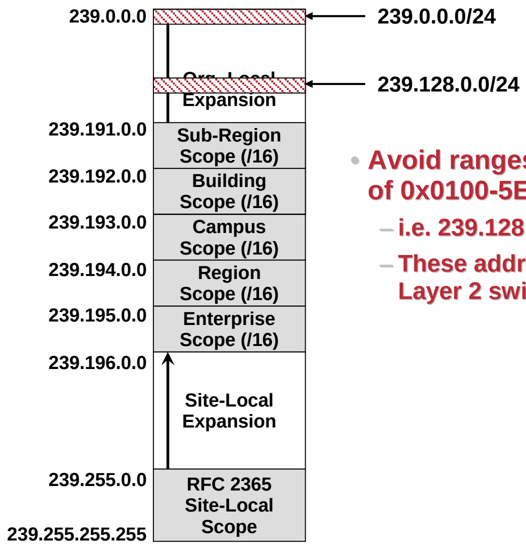

- Avoid ranges that map to a MAC address of 0x0100-5E00-00xx!
  - i.e. 239.128.0/24 & 239.0.0/24
  - These addresses are always flooded by Layer 2 switches!

RST-2701
9799_05_2004_X

101

Enterprise Scope Relative Range

Cisco.com

- Be sure to place Enterprise Scope at the high end of the Org-Local range.
  - Keeps Org-Local and Enterprise Scope Relative ranges identical.
  - Insures applications that use Org-Local Scope Relative addresses work correctly.

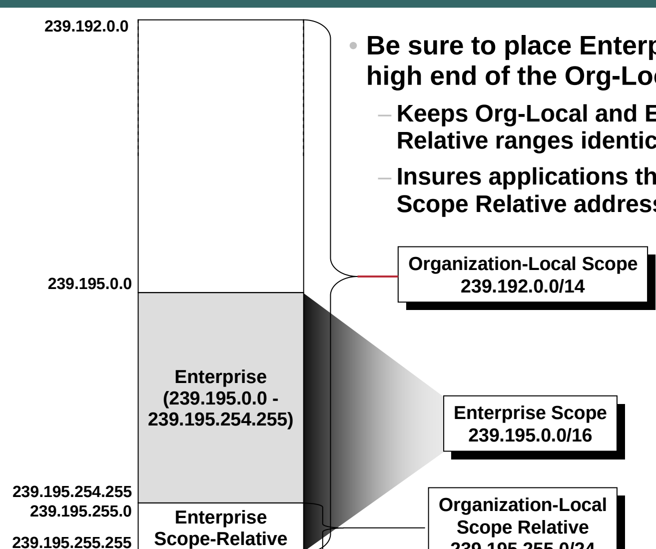

(Not to scale.)

RST-2701
9799_05_2004_X

102

Adding Bidir Ranges to each Scope

Cisco.com

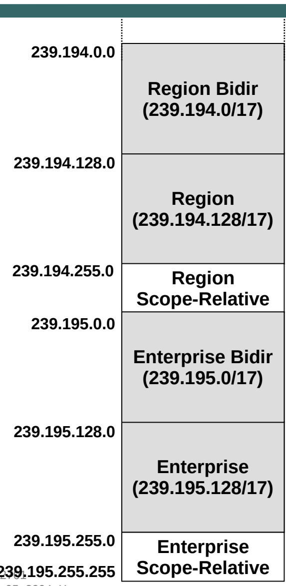

- Subdivide each scope's address range into Bidir and ASM ranges.
  - Keep ASM range at the upper end of the address range.
  - Keeps Scope-Relative multicast in ASM mode.

RST-2701
9799_05_2004_X

103

Adding Private SSM Space

Cisco.com

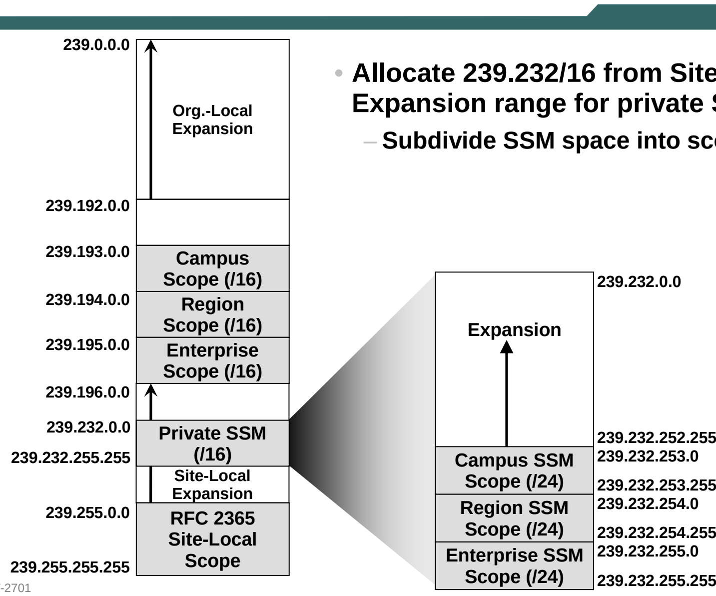

- Allocate 239.232/16 from Site-Local Expansion range for private SSM space.
  - Subdivide SSM space into scoped zones.

RST-2701
9799_05_2004_X

104

Deploying Administratively-Scoped Zones

Cisco.com

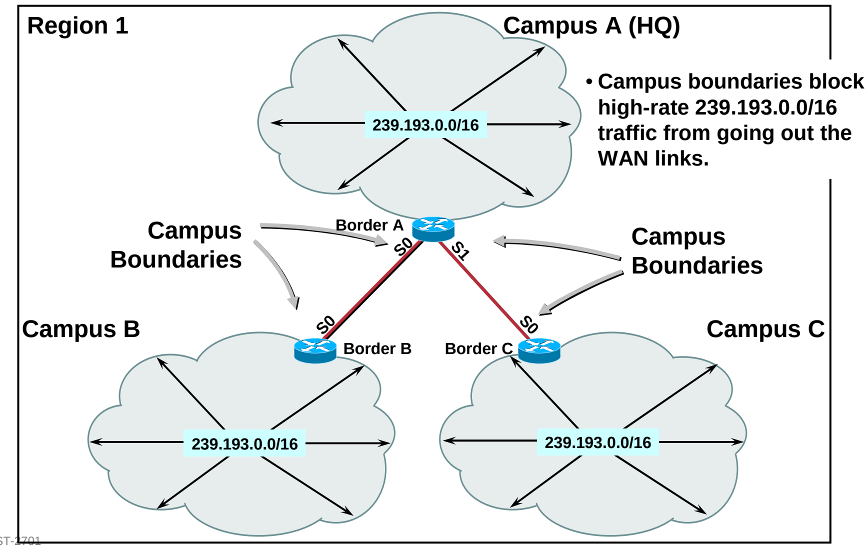

RST-2701
9799_05_2004_X

105

Deploying Administratively-Scoped Zones

Cisco.com

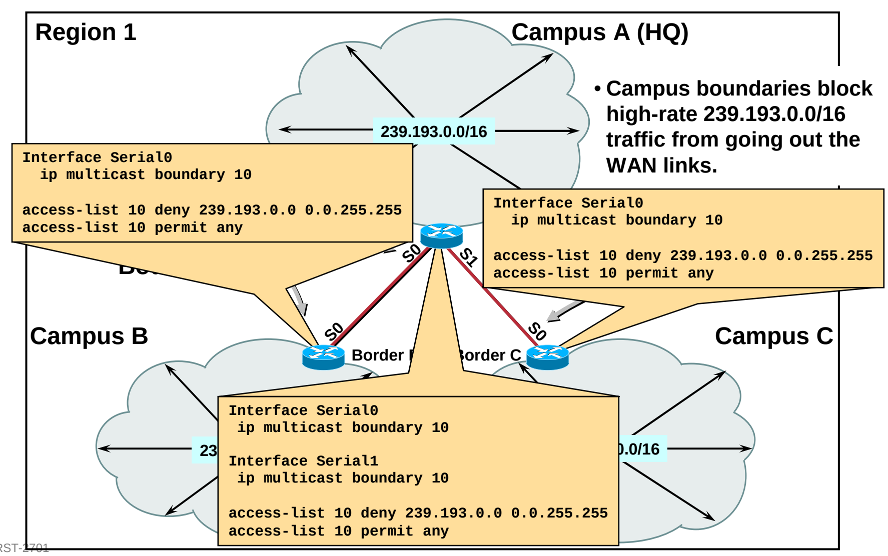

RST-2701
9799_05_2004_X

106

Deploying Administratively-Scoped Zones

Auto-RP Example

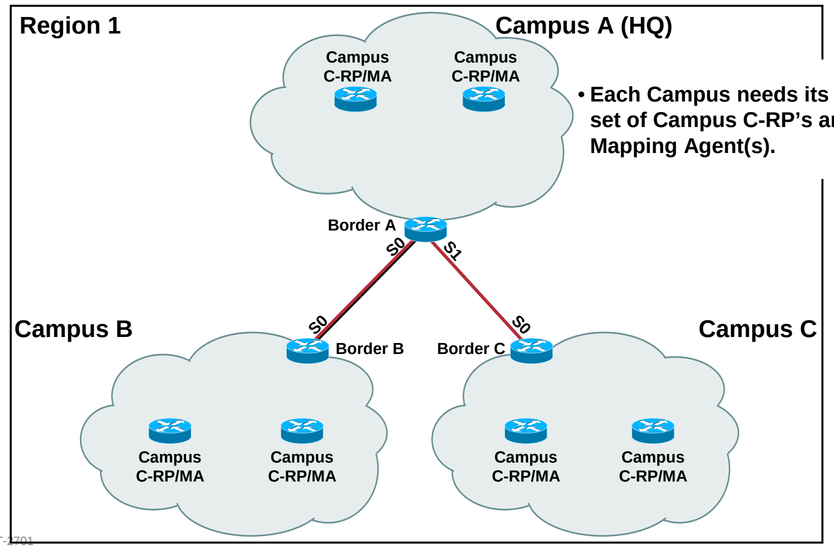

- Each Campus needs its own set of Campus C-RP's and Mapping Agent(s).

Deploying Administratively-Scoped Zones

Auto-RP Example

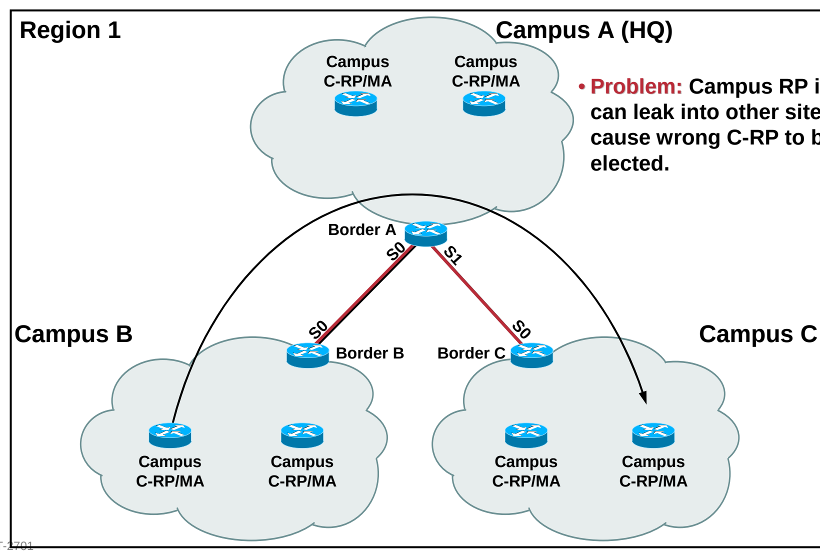

- **Problem:** Campus RP info can leak into other sites and cause wrong C-RP to be elected.

Deploying Administratively-Scoped Zones
Preventing Auto-RP Info Leakage

Cisco.com

- **Multicast Boundary Command**

  `ip multicast boundary <acl> [filter-autorp]`

  – **New 'filter-autorp' option**

    - **Filters contents of Auto-RP packets**
      - Filters both Announcement and Discovery messages
      - C-RP entries that fail <acl> are removed from packet
    - **Prevents C-RP information from leaking in/out of scoped zone.**
    - **Greatly simplifies Admin. Scoped Zone support in Auto-RP.**
    - **Available in 12.0(22)S, 12.2(12).**

RST-2701
9799_05_2004_X

109

Deploying Administratively-Scoped Zones
Preventing Auto-RP Info Leakage

Cisco.com

- **How 'filter-autorp' option works:**

  For each RP Entry in Auto-RP packet:

  If group-range in RP-Entry *'intersects'* any 'denied' group-range in the Multicast Boundary ACL, delete RP Entry from Auto-RP packet.

  If resulting Auto-RP packet is non-empty, forward across multicast boundary.

RST-2701
9799_05_2004_X

110

Deploying Administratively-Scoped Zones

Preventing Auto-RP Info Leakage

Cisco.com

- Using Multicast Boundary ‘filter-autorp’
  - Avoid Auto-RP Group-Range Overlaps
    - Overlapping ranges can “intersect” denied ranges at multicast boundaries.
      - Can cause unexpected Auto-RP info filtering at multicast boundaries.
      - Results in loss of Auto-RP info to other parts of network.
  - Rule of Thumb:
    - Make sure Auto-RP Group-Ranges match exactly any Multicast Boundary Ranges!
      (i.e. don’t use overlapping Auto-RP group ranges.)

RST-2701
9799_05_2004_X

111

Avoid Overlapping Group Ranges

Cisco.com

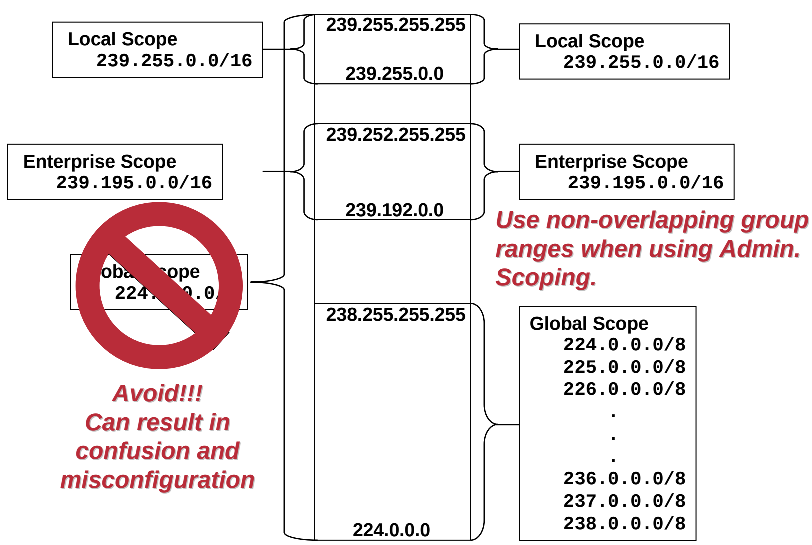

RST-2701
9799_05_2004_X

112

Avoid Overlapping Group Ranges

Cisco.com

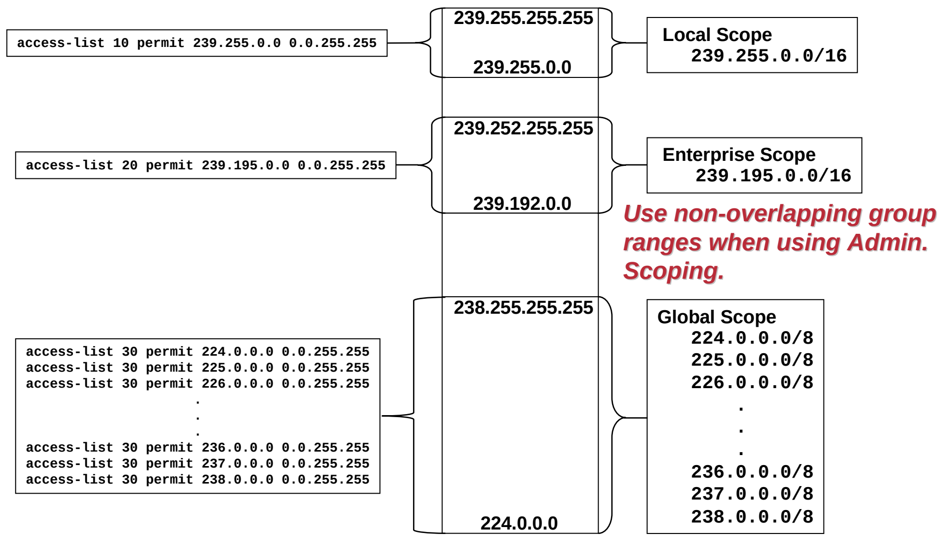

RST-2701
9799_05_2004_X

113

Avoid Overlapping Group Ranges

Cisco.com

- **Avoiding Overlapping Group Ranges**
  - **Can't use "deny" clause in C-RP ACL's**
    - **Implies "Dense-mode Override"**

```
ip pim send-rp-announce loopback0 scope 16 group-list 10
access-list 10 deny 239.0.0.0 0.255.255.255
access-list 10 permit 224.0.0.0 15.255.255.255
```

  - **Must only use "permit" clauses**

```
ip pim send-rp-announce loopback0 scope 16 group-list 10
access-list 10 permit 224.0.0.0 0.255.255.255
access-list 10 permit 225.0.0.0 0.255.255.255
                     .
                     .
                     .
access-list 10 permit 238.0.0.0 0.255.255.255
```

RST-2701
9799_05_2004_X

114

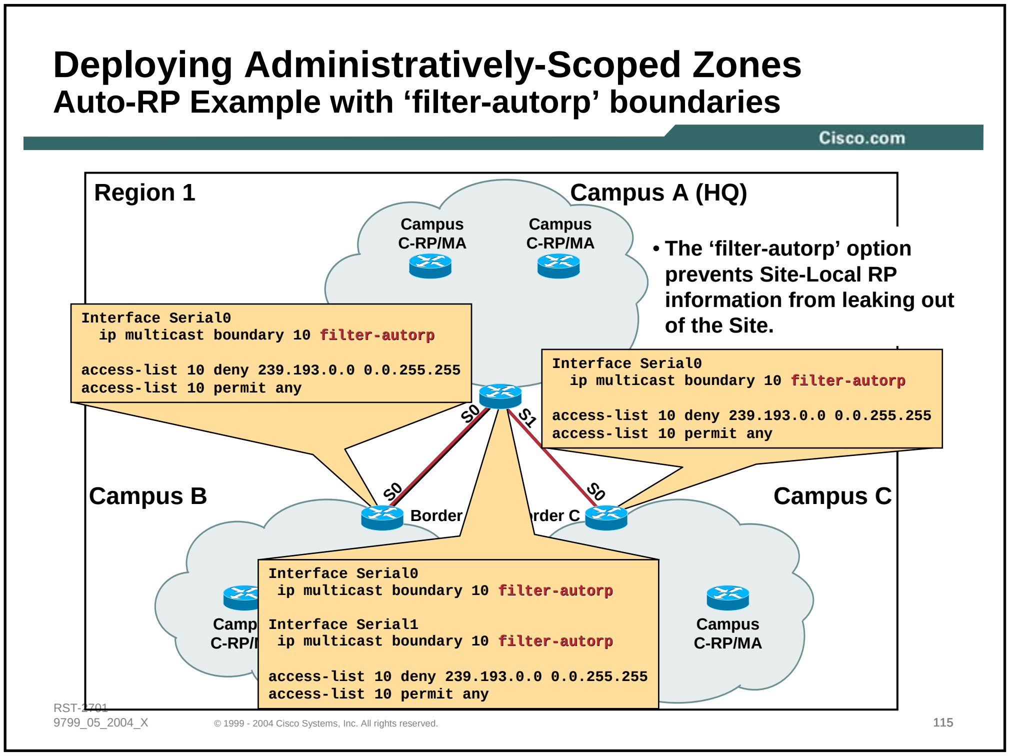
Deploying Administratively-Scoped Zones
Auto-RP Example with ‘filter-autorp’ boundaries
Cisco.com
Region 1
Campus A (HQ)
Campus C-RP/MA
Campus C-RP/MA
• The ‘filter-autorp’ option prevents Site-Local RP information from leaking out of the Site.
Interface Serial0
ip multicast boundary 10 filter-autorp
access-list 10 deny 239.193.0.0 0.0.255.255
access-list 10 permit any
Interface Serial0
ip multicast boundary 10 filter-autorp
access-list 10 deny 239.193.0.0 0.0.255.255
access-list 10 permit any
S0
S1
S0
S0
Campus B
Campus C
Border
rder C
Interface Serial0
ip multicast boundary 10 filter-autorp
Interface Serial1
ip multicast boundary 10 filter-autorp
access-list 10 deny 239.193.0.0 0.0.255.255
access-list 10 permit any
Campus C-RP/MA
RST-2701
9799_05_2004_X
© 1999 - 2004 Cisco Systems, Inc. All rights reserved.
115

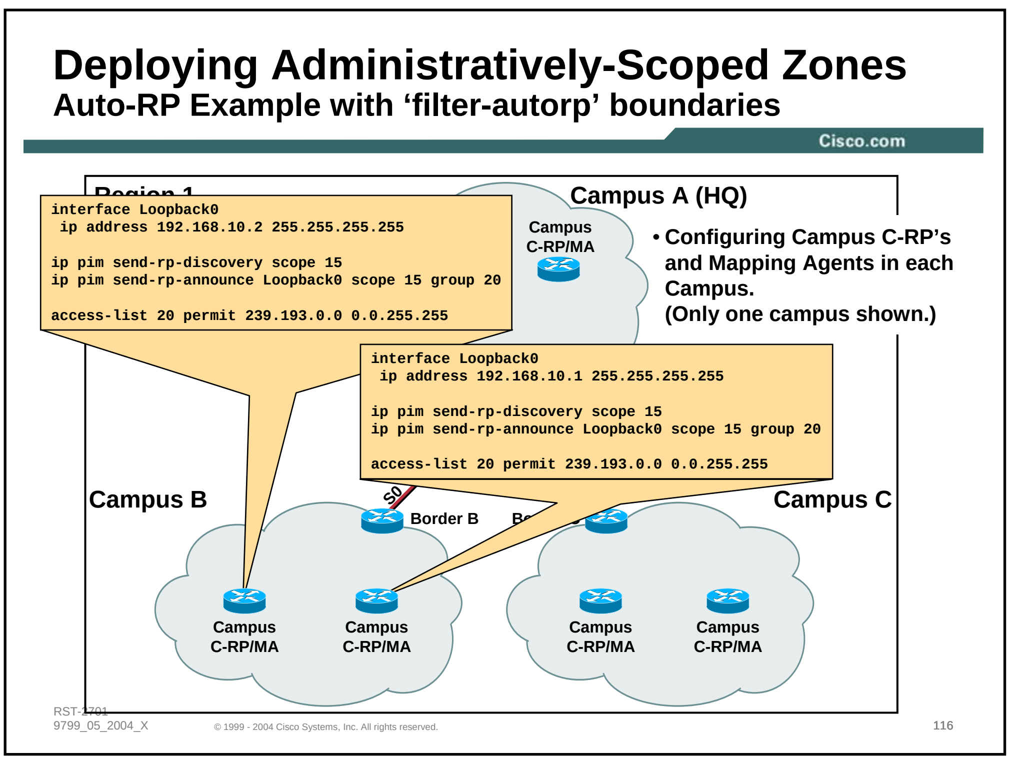
Deploying Administratively-Scoped Zones
Auto-RP Example with ‘filter-autorp’ boundaries
Cisco.com
Campus A (HQ)
Campus C-RP/MA
• Configuring Campus C-RP’s and Mapping Agents in each Campus.
(Only one campus shown.)
interface Loopback0
ip address 192.168.10.2 255.255.255.255
ip pim send-rp-discovery scope 15
ip pim send-rp-announce Loopback0 scope 15 group 20
access-list 20 permit 239.193.0.0 0.0.255.255
interface Loopback0
ip address 192.168.10.1 255.255.255.255
ip pim send-rp-discovery scope 15
ip pim send-rp-announce Loopback0 scope 15 group 20
access-list 20 permit 239.193.0.0 0.0.255.255
Campus B
Campus C
S0
Border B
Campus C-RP/MA
Campus C-RP/MA
Campus C-RP/MA
Campus C-RP/MA
RST-2701
9799_05_2004_X
© 1999 - 2004 Cisco Systems, Inc. All rights reserved.
116

Deploying Administratively-Scoped Zones

Auto-RP Example with ‘filter-autorp’ boundaries

Cisco.com

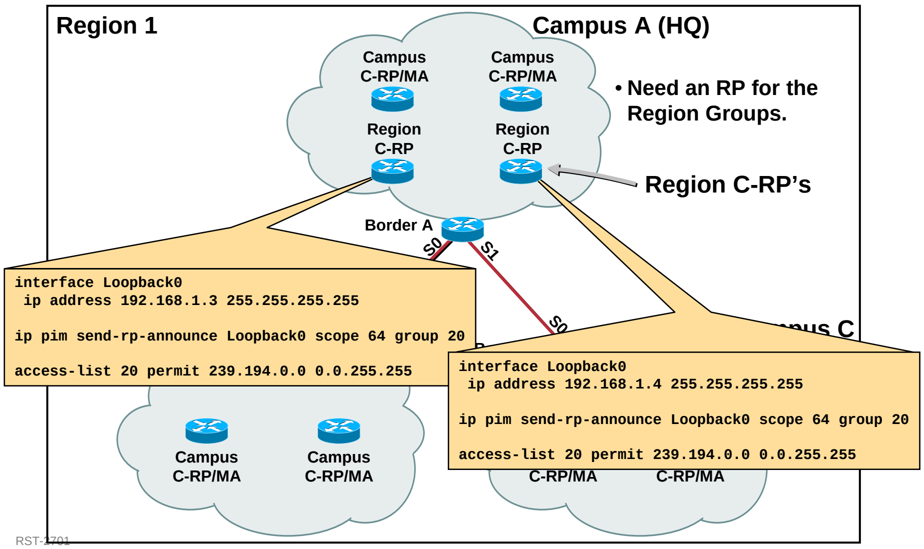

RST-2701
9799_05_2004_X

117

More Information

Cisco.com

- White Papers
- Web and Mailers
- Cisco Press

RTFB

RTFB = “Read the Fine Book”

CCO Multicast page:
http://www.cisco.com/go/ipmulticast

Questions:
cs-ipmulticast@cisco.com

Customer Support Mailing List:
tac@cisco.com

RST-2701
9799_05_2004_X

118

Wonderful Bedtime Stories

Cisco.com

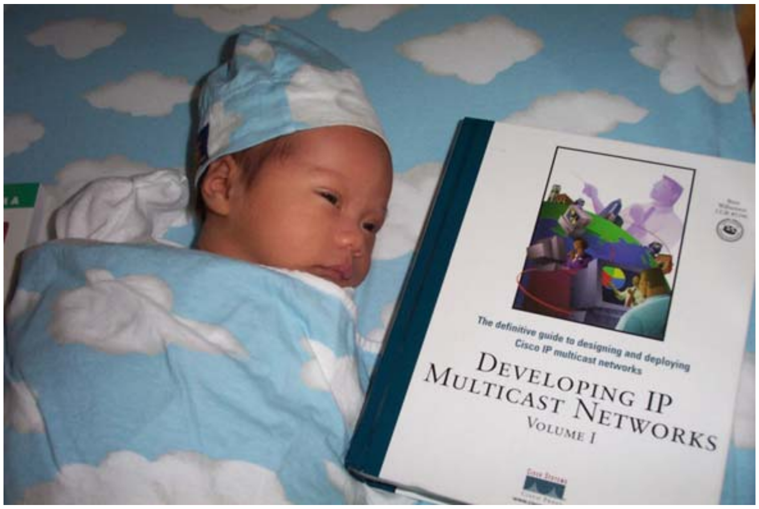

RST-2701
9799_05_2004_X

119

Complete Your Online Session Evaluation!

Cisco.com

WHAT: Complete an online session evaluation and your name will be entered into a daily drawing

WHY: Win fabulous prizes! Give us your feedback!

WHERE: Go to the Internet stations located throughout the Convention Center

HOW: Winners will be posted on the onsite Networkers Website; four winners per day

RST-2701
9799_05_2004_X

120

RST-2701
9799_05_2004_X

121