GUIDELINES FOR ENTERPRISE IP MULTICAST ADDRESS ALLOCATION

INTRODUCTION

With the growing popularity of IP multicast applications, many customers are considering deploying or have already deployed IP multicast in their networks. A prerequisite of successful IP multicast deployment is designing and documenting an appropriate multicast addressing scheme.

Some of the common problems that customers face during the IP multicast deployment are:

- Keeping administration and troubleshooting of IP multicast as simple as possible
- Controlling the distribution and the use of IP multicast group addresses within an organization, for instance, between Business Units
- Controlling the distribution and scope of multicast application data within an organization
- Locating Rendezvous Points with Protocol Independent Multicast (PIM) Sparse mode and which IP multicast groups each will serve
- Controlling IP multicast traffic on WAN links so that high rate groups cannot saturate low speed links
- For security, controlling who can send IP multicast and who can receive
- Readiness to deploy next generation IP multicast protocols like Bidirectional (Bidir) PIM and Source Specific Multicast (SSM)
- Linking to the Internet for multicast
- Allowing for future requirements and expansion so that readdressing is unnecessary at a later date

The key to solving or simplifying many of these issues is a solid addressing policy. Without a correctly planned and scoped IP multicast addressing scheme, customers will face more complex configurations, which significantly decrease the control of IP multicast over the network and increase administration and support overhead.

The objective of this document is to provide Cisco Systems® customers with several methodologies of allocating IP multicast group addresses. With this knowledge, the appropriate scheme or combinations of schemes can be chosen to fit specific customer requirements. It is assumed the reader is familiar with IP multicast terminology.

REFERENCES

- RFC 3171, *"IANA Guidelines for IPv4 Multicast Address Assignments,"* Best Current Practice, July 2001:
  http://search.ietf.org/internet-drafts/draft-ietf-mboned-iana-ipv4-mcast-guidelines-04.txt
- RFC 2365, *"Administratively Scoped IP Multicast,"* Best Current Practice, July 1998:
  http://www.ietf.org/rfc/rfc2365.txt
- The Multicast Addresses registry: http://www.iana.org/assignments/multicast-addresses

- Advanced Services RPs guidelines for Enterprise—Contact your Advanced Services Engineer—document is in the Cisco® Knowledge Base Management System (KBMS)
- RFC 2730, *"Multicast Address Dynamic Client Allocation Protocol (MADCAP)*, December 1999: http://www.ietf.org/rfc/rfc2730.txt
- RFC 2908, *"The Internet Multicast Address Allocation Architecture,"* September 2000: http://www.ietf.org/rfc/rfc2908.txt
- RFC 2770, *"GLOP Addressing in 233/8,"* February 2000: http://www.ietf.org/rfc/rfc2770.txt
- RFC 2776, *"Multicast-Scope Zone Announcement Protocol (MZAP),"* February 2000: http://www.ietf.org/rfc/rfc2776.txt
- RFC 2974, *"Session Announcement Protocol (SAP),"* October 2000: http://www.ietf.org/rfc/rfc2974.txt
- RFC 2327, *"SDP: Session Description Protocol,"* April 1998: http://www.ietf.org/rfc/rfc2327.txt
- Multicast Address Allocation (MALLOC) Working Group: http://www.aciri.org/malloc/
- *"Developing IP Multicast Networks,"* Beau Williamson, Copyright © 2000, Cisco Press, ISBN 1578700779: http://www2.ciscopress.com/search/index.asp?searchstring=Developing+IP+Multicast+Networks&searchgroup=Entire+Site&searchtype=Title

Multicast Address Fundamentals

By its nature, IP Multicast forwarding is different from IP Unicast and has additional addressing requirements to consider.

- Multicast sources and RPs (Rendezvous Points) are identified by their unique unicast address as a prerequisite.
- Multicast group addresses can be shared; for instance, many sources can send to the same address.
- Routers do not differentiate that many streams or channels may be multiplexed with different User Datagram Protocol port numbers, and the clients extract the data they want.
- Because any multicast source can send to any group address and any multicast client can receive any group without regard to geography, aggregation and summarization of multicast group addresses are meaningless.
- In the past, Time To Live field in the IP Multicast datagram was used for creating Auto-RP administrative boundaries using the **ttl-threshold** command. This has been superseded by the **ip multicast boundary** interface mode command, which filters IP multicast traffic and also Auto-RP messages. This is detailed in later sections of this paper.
- Administrative or private address space can and should be used within the enterprise unless multicast traffic will be sourced to the Internet, therefore requiring a unique group address.

Layer 3 IP Multicast Addresses

IP multicast addresses have been assigned to the old Class "D" address space by the Internet Assigned Number Authority (IANA). Addresses in this space are denoted with a binary "1110" prefix in the first four bits of the first octet, as shown in Figure 1. This results in IP multicast addresses spanning a range from 224.0.0.0 through 239.255.255.255.

Figure 1
IP Multicast Address Format

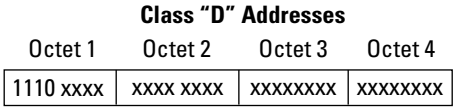

Layer 2 Multicast Addressing

The IEEE 802.2 specification makes provisions for the transmission of broadcast and multicast packets. As shown in Figure 2, Bit 0 of Octet 0 in an IEEE MAC address indicates whether the destination address is a broadcast/multicast address or a unicast address.

Figure 2
Ethernet MAC Address Format

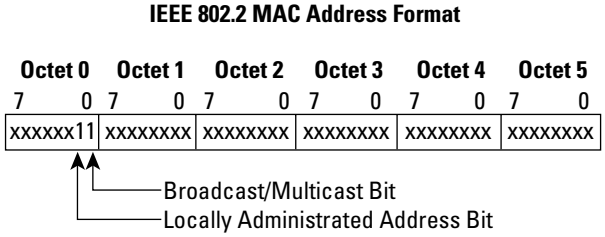

If this bit is set, the MAC frame is destined either for an arbitrary group of hosts or all hosts on the network (if the MAC destination address is the broadcast address 0xFFFF.FFFF.FFFF). IP multicasting at Layer 2 uses this capability to transmit IP multicast packets to a group of hosts on a LAN segment.

Layer 2 Multicast MAC Address Mapping

All IP multicast frames all MAC layer addresses beginning with the 24-bit prefix of 0x0100.5Exx.xxxx. With only half of these MAC addresses available for use by IP Multicast, 23 bits of MAC address space is available for mapping Layer 3 IP multicast addresses into Layer 2 MAC addresses. All Layer 3 IP multicast addresses have the first four of the 32 bits set to 0x1110, leaving 28 bits of meaningful IP multicast address information. These 28 bits must map into only 23 bits of the available MAC address. This mapping is shown graphically in Figure 3.

Figure 3
IP Multicast to Ethernet/FDDI MAC Address Mapping

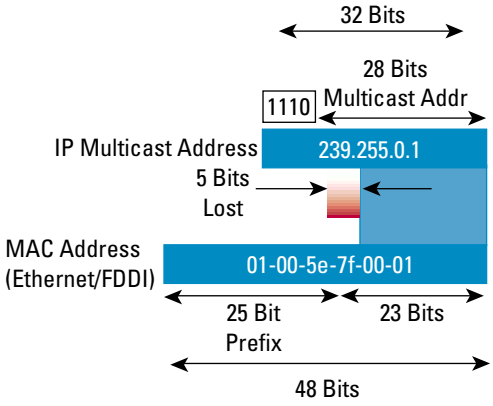

Performance Impact of Address Mapping

Because all 28 bits of the Layer 3 IP multicast address information cannot be mapped into the available 23 bits of MAC address space, five bits of address information are lost in the mapping process. This results in a 32:1 address ambiguity when a Layer 3 IP multicast address is mapped to a Layer 2 IEEE MAC address. This means that each IEEE IP multicast MAC address can represent 32 IP multicast addresses as shown in Figure 4.

This 32:1 address ambiguity has the potential to cause problems. For example, a host that wishes to receive multicast group 224.1.1.1 will program the hardware registers in the network interface card to interrupt the CPU when a frame with a destination multicast MAC address of 0x0100.5E00.0101 is received. Unfortunately, this same multicast MAC address is also used for 31 other IP multicast groups. If any of these 31 other groups are also active on the local LAN, the host CPU will receive interrupts when a frame is received for any of these other groups. The CPU must examine the IP portion of each of these received frames to determine if it is the desired group; for instance, 224.1.1.1. This can affect the host's available CPU power if the amount of "spurious" group traffic is high enough.

Figure 4
MAC Address Ambiguities

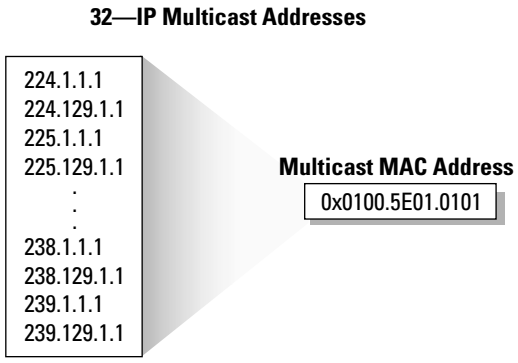

In addition to having a possible negative impact on host CPU power, this ambiguity can also cause problems when trying to constrain multicast flooding in Layer 2 LAN switches based solely on these multicast MAC addresses

IANA Assigned Multicast Address Blocks

The IETF has provided the IANA with guidance on how IP Multicast address space should be allocated in RFC 3171bis, *"IANA Guidelines for IPv4 Multicast Address Assignments."* Table 1 lists the current assignments blocks documented in RFC 3171bis.

Table 1 IANA Multicast Address Assignments

Range	Mask	Description
224.0.0.0–224.0.0.255	224.0.0/24	Local Network Control Block
224.0.1.0–224.0.1.255	224.0.1/24	Internetwork Control Block
224.0.2.0–224.0.255.255	–	Ad hoc Block
224.1.0.0–224.1.255.255	–	Unassigned
224.2.0.0–224.2.255.255	224.2/16	SDP/SAP Block
224.3.0.0–231.255.255.255	–	Unassigned
232.0.0.0–232.255.255.255	232/8	Source Specific Multicast Block
233.0.0.0–233.255.255.255	233/8	GLOP Block
234.0.0.0–238.255.255.255	–	Unassigned
239.0.0.0–239.255.255.255	239/8	Administratively Scoped Block

Of these address blocks, the IANA assigns addresses in the Local Network Control, Internetwork Control and Ad hoc blocks based on the guidelines also supplied by the IETF in RFC 3171bis. These guidelines call for Expert Review, Internet Engineering Steering Group (IESG) approval, or a Standards Action process before the IANA assigns addresses in these spaces.

Local Network Control Block

The range of 224.0.0.0 through 224.0.0.255 is considered the Local Network Control Block—more commonly known as the Link-Local Block—and is used by network protocols on a local subnet segment. Packets with an address in this range are local in scope and always should be transmitted with a Time To Live (TTL) of 1 so that they go no farther than the local subnet.

Table 2 is a list of the Link-Local multicast addresses taken directly from the IANA database at the time this paper was written. The table lists the reserved multicast addresses, the network protocol function to which it has been assigned, and the person who requested the address or the RFC associated with the protocol.

Table 2 Local Network Control Block

Address	Usage	Reference
224.0.0.1	All Hosts	[RFC 1112, JBP]
224.0.0.2	All Multicast Rrouters	[IANA]
224.0.0.3	Unassigned	[IANA]

Table 2 Local Network Control Block (Continued)

Address	Usage	Reference
224.0.0.4	DVMRP Routers	[RFC 1075, JBP]
224.0.0.5	OSPF Routers	[RFC 1583, JXM1]
224.0.0.6	OSPF Designated Rrouters	[RFC 1583, JXM1]
224.0.0.7	ST Routers	[RFC 1190, KS14]
224.0.0.8	ST Hosts	[RFC 1190, KS14]
224.0.0.9	RIP2 Routers	[RFC 1723, SM11]
224.0.0.10	IGRP Routers	[Farinacci]
224.0.0.11	Mobile-Agents	[Bill Simpson]
224.0.0.12	DHCP Server/Relay Agent	[RFC1884]
224.0.0.13	All PIM Routers	[Farinacci]
224.0.0.14	RSVP-ENCAPSULATION	[Braden]
224.0.0.15	all-cbt-routers	[Ballardie]
224.0.0.16	designated-sbm	[Baker]
224.0.0.17	all-sbms	[Baker]
224.0.0.18	VRRP	[Hinden]
224.0.0.19	IPAIlL1lss	[Przygienda]
224.0.0.20	IPAIlL2lss	[Przygienda]
224.0.0.21	IPAIlIntermediate Systems	[Przygienda]
224.0.0.22	IGMP	[Deering]
224.0.0.23	GLOBECAST-ID	[Scannell]
224.0.0.24	Unassigned	[IANA]
224.0.0.25	router-to-switch	[Wu]
224.0.0.26	Unassigned	[IANA]
224.0.0.27	Al MPP Hello	[Martinicky]
224.0.0.28	ETC Control	[Polishinski]
224.0.0.29	GE-FANUC	[Wacey]
224.0.0.30	indigo-vhdp	[Caughie]
224.0.0.31	Shinbroadband	[Kittivatcharapong]
224.0.0.32	Digistar	[Kerkan]
224.0.0.33	ff-system-management	[Glanzer]
224.0.0.34	Pt2-discover	[Kammerlander]

Table 2 Local Network Control Block (Continued)

Address	Usage	Reference
224.0.0.35	DXCLUSTER	[Koopman]
224.0.0.36	DTCP Announcement	[Cipiere]
224.0.0.37- 224.0.0.68	Zeroconfaddr	[Guttman]
224.0.0.69- 224.0.0.100	Unassigned	[IANA]
224.0.0.101	cisco-nhap	[Bakke]
224.0.0.102	HSRP	[Wilson]
224.0.0.103	MDAP	[Deleu]
224.0.0.104	Nokia MC CH	[Kalhour]
224.0.0.105	ff-lr-address	[Glanzer]
224.0.0.106- 224.0.0.250	Unassigned	[IANA]
224.0.0.251	MDNS	[Cheshire]
224.0.0.252- 224.0.0.255	Unassigned	[IANA]

The OSPF routing protocol is a good example where Local Network Control multicast addresses are employed by a protocol. If you use OSPF in your network, you may have seen packets addressed to the 224.0.0.5 and 224.0.0.6 multicast address on your networks. These addresses permit OSPF routers to communicate important OSPF data to All OSPF Routers or All OSPF Designated Routers, respectively.

The assignments in the Link-Local block already have taken nearly one-third of this very limited resource and care must be taken to insure that this critical IPv4 resource does not become exhausted. To avoid this, it may soon become necessary for the IANA to begin periodic reviews of these allocations and to reclaim some of these allocations that are no longer actively being used.

Layer 2 Flooding of Link-Local Multicast

IGMP Snooping normally is used by Layer 2 switches to constrain multicast traffic only to those ports that have hosts attached and that have signaled their desire to join the multicast group by sending IGMP Membership Reports. However, it is important to note that most Layer 2 switches flood all multicast traffic that falls within the MAC address range of 0x0100.5E00.00xx (which corresponds to Layer 3 addresses in the Link-Local block) to all ports

on the switch even if IGMP Snooping is enabled. This is true for the current suite of Cisco switches. The reason that this Link-Local multicast traffic is always flooded is that IGMP Membership Reports normally are never sent for multicast traffic in the Link-Local block. For example, routers do not send IGMP Membership Reports for the ALL-OSPF-ROUTERS group (224.0.0.5) when OSPF is enabled. Therefore, if Layer 2 switches were to constrain (that is, not flood) Link-Local packets in the 224.0.0.0/24 (0x0100.5E00.00xx) range to only those ports where IGMP Membership reports were received, Link-Local protocols such as OSPF would break.

The impact of this Link-Local flooding in combination with the 32:1 ambiguity that arises when Layer 3 multicast addresses are mapped to Layer 2 MAC addresses means that there are several multicast group ranges besides the 224.0.0.0/24 that will map to the 0x0100.5E00.00xx MAC address range and hence also will be flooded by most Layer 2 switches. *It is recommended that multicast addresses that map to the 0x0100.5E00.00xx MAC address range be avoided*. Table 3 lists all multicast address ranges that should not be used if Layer 2 flooding is to be avoided.

Table 3 Multicast Groups Flooded by Layer 2 Switches

Multicast Address Range	MAC Address Range
224.0.0.0/24	0x0100.5E00.00xx
224.128.0.0/24	0x0100.5E00.00xx
225.0.0.0/24	0x0100.5E00.00xx
225.128.0.0/24	0x0100.5E00.00xx
226.0.0.0/24	0x0100.5E00.00xx
226.128.0.0/24	0x0100.5E00.00xx
227.0.0.0/24	0x0100.5E00.00xx
227.128.0.0/24	0x0100.5E00.00xx
228.0.0.0/24	0x0100.5E00.00xx
228.128.0.0/24	0x0100.5E00.00xx
229.0.0.0/24	0x0100.5E00.00xx
229.128.0.0/24	0x0100.5E00.00xx
230.0.0.0/24	0x0100.5E00.00xx
230.128.0.0/24	0x0100.5E00.00xx
231.0.0.0/24	0x0100.5E00.00xx
231.128.0.0/24	0x0100.5E00.00xx
232.0.0.0/24	0x0100.5E00.00xx
232.128.0.0/24	0x0100.5E00.00xx
233.0.0.0/24	0x0100.5E00.00xx
233.128.0.0/24	0x0100.5E00.00xx

Table 3 Multicast Groups Flooded by Layer 2 Switches

Multicast Address Range	MAC Address Range
234.0.0.0/24	0x0100.5E00.00xx
234.128.0.0/24	0x0100.5E00.00xx
235.0.0.0/24	0x0100.5E00.00xx
235.128.0.0/24	0x0100.5E00.00xx
236.0.0.0/24	0x0100.5E00.00xx
236.128.0.0/24	0x0100.5E00.00xx
237.0.0.0/24	0x0100.5E00.00xx
237.128.0.0/24	0x0100.5E00.00xx
238.0.0.0/24	0x0100.5E00.00xx
238.128.0.0/24	0x0100.5E00.00xx
239.0.0.0/24	0x0100.5E00.00xx
239.128.0.0/24	0x0100.5E00.00xx

Inter-Network Control Block

The range of 224.0.1.0 through 224.0.1.255 is the Inter-Network Control Block. These addresses are similar to the Local Network Control Block except that they are used by network protocols when control messages need to be multicast beyond the local network segment.

Table 4 is a partial listing of the first 50 multicast addresses in the Inter-Network Control Block that have been assigned by the IANA. The table lists the address, the network protocol function to which it has been assigned, and the person who requested the address or the RFC associated with the protocol.

The Cisco Auto-RP protocol is a good example of Inter-Network Control multicast. Inter-Network Control multicast addresses 224.0.1.39 and 224.0.1.40 are assigned as the Cisco Announce and Cisco Discovery multicast groups, respectively. The Auto-RP mechanism uses these two multicast groups to communicate Rendezvous Point information through a PIM Sparse Mode (PIM-SM) domain.

At the time that this paper was written, address assignments in the Inter-Network Control Block have already used more than two-thirds of this limited resource. Many of these assignments may be for something other than what this address space was intended for or for protocols that have not (or possibly will not) see the light of day.

Table 4 Inter-Network Control Multicast Addresses

Address	Use	Reference
224.0.1.0	VMTP Managers Group	[RFC 1045,DRC3]
224.0.1.1	Network Time Protocol (NTP)	[RFC 1119,DLM1]
224.0.1.2	SGI-Dogfight	[AXC]
224.0.1.3	Rwhod	[SXD]

Table 4 Inter-Network Control Multicast Addresses (Continued)

Address	Use	Reference
224.0.1.4	VNP	[DRC3]
224.0.1.5	Artificial Horizons—Aviator	[BXF]
224.0.1.6	Name Service Server	[BXS2]
224.0.1.7	AUDIONEWS—Audio News Multicast	[MXF2]
224.0.1.8	SUN NIS+ Information Service	[CXM3]
224.0.1.9	Multicast Transport Protocol	[SXA]
224.0.1.10	IETF-1-LOW-AUDIO	[SC3]
224.0.1.11	IETF-1-AUDIO	[SC3]
224.0.1.12	IETF-1-VIDEO	[SC3]
224.0.1.13	IETF-2-LOW-AUDIO	[SC3]
224.0.1.14	IETF-2-AUDIO	[SC3]
224.0.1.15	IETF-2-VIDEO	[SC3]
224.0.1.16	MUSIC-SERVICE	[Guido van Rossum]
224.0.1.17	SEANET-TELEMETRY	[Andrew Maffei]
224.0.1.18	SEANET-IMAGE	[Andrew Maffei]
224.0.1.19	MLOADD	[Braden]
224.0.1.20	any private experiment	[IANA]
224.0.1.21	DVMRP on MOSPF	[John Moy]
224.0.1.22	SVRLOC	[Veizades]
224.0.1.23	XINGTV	[Gordon]
224.0.1.24	Microsoft-ds	–
224.0.1.25	nbc-pro	–
224.0.1.26	nbc-pfn	–
224.0.1.27	lmsc-calren-1	[Uang]
224.0.1.28	lmsc-calren-2	[Uang]
224.0.1.29	lmsc-calren-3	[Uang]
224.0.1.30	lmsc-calren-4	[Uang]
224.0.1.31	ampr-info	[Janssen]
224.0.1.32	Mtrace	[Casner]
224.0.1.33	RSVP-encap-1	[Braden]
224.0.1.34	RSVP-encap-2	[Braden]

Table 4 Inter-Network Control Multicast Addresses (Continued)

Address	Use	Reference
224.0.1.35	SVRLOC-DA	[Veizades]
224.0.1.36	rln-server	[Kean]
224.0.1.37	Proshare-mc	[Lewis]
224.0.1.38	Dantz	[Zulch]
224.0.1.39	cisco-rp-announce	[Farinacci]
224.0.1.40	cisco-rp-discovery	[Farinacci]
224.0.1.41	Gatekeeper	[Toga]
224.0.1.42	iberiagames	[Marocho]
224.0.1.43	nwn-discovery	[Zwemmer]
224.0.1.44	nwn-adaptor	[Zwemmer]
224.0.1.45	isma-1	[Dunne]
224.0.1.46	isma-2	[Dunne]
224.0.1.47	Telerate	[Peng]
224.0.1.48	Ciena	[Rodbell]
224.0.1.49	dcap-servers	[RFC 2114]
224.0.1.50	dcap-clients	[RFC 2114]

Ad Hoc Multicast Block

The multicast group range of 224.0.2.0 through 224.0.255.255 is the Ad Hoc Multicast Block. Historically, addresses in this range have been assigned to applications that do not clearly fall into the Local Network Control and Inter-Network Control categories. In general, the guidelines provided in RFC 3171bis for the assignment of addresses in this range state that the IANA should assign addresses in this range only under special circumstances. Even then, the assignment must undergo a strict Expert Review, IESG Approval, or Standards Action process before addresses are assigned.

Table 5 is a list of assignments in the Ad Hoc block at the time that this paper was written. Much of the address space in the Ad Hoc block was assigned by the IANA prior to receiving clear guidelines from the IETF, because most of the assignments in Table 5 *clearly* intend to reserve address space to source future commercial content. These sorts of commercial content address reservations are *highly* discouraged and are considered as address hoarding by the multicast community. This is especially true because most organizations can use GLOP addressing to source their commercial content to the Internet. Furthermore, some of these assignments are intended for multicast traffic that never will be transmitted to the Internet and do not need global multicast address space.

Finally, organizations that reserve large blocks of this address space to source commercial content should be aware that RFC 3171bis recommends an annual review of all assigned addresses and, when possible, reclaim improperly assigned addresses. It is hoped that the IANA soon will review these assignments and return most of this address space to “unassigned” status.

Table 5 Ad Hoc Multicast Addresses

Address Range	Usage	Reference
224.0.2.1	“rwho” Group (unofficial)	[IANA]
224.0.2.2	SUN RPC PMAPPROC_CALLIT	[BXE1]
224.0.2.064-224.0.2.095	SIAC MDD Service	[Tse]
224.0.2.096-224.0.2.127	CoolCast	[Ballister]
224.0.2.128-224.0.2.191	WOZ-Garage	[Marquardt]
224.0.2.192-224.0.2.255	SIAC MDD Market Service	[Lamberg]
224.0.3.0-224.0.3.255	RFE Generic Service	[DXS3]
224.0.4.0-224.0.4.255	RFE Individual Conf.	[DXS3]
224.0.5.0-224.0.5.127	CDPD Groups	[Bob Brenner]
224.0.5.128-224.0.5.191	SIAC Market Service	[Cho]
224.0.5.192-224.0.5.255	SIAC NYSE Order PDP protocol	[Chan]
224.0.6.0-224.0.6.127	Cornell ISIS Project	[Tim Clark]
224.0.6.128-224.0.6.255	Unassigned	[IANA]
224.0.7.0-224.0.7.255	Where-Are-You	[Simpson]
224.0.8.0-224.0.8.255	INTV	[Tynan]
224.0.9.0-224.0.9.255	Invisible Worlds	[Malamud]
224.0.10.0-224.0.10.255	DLSw Groups	[Lee]
224.0.11.0-224.0.11.255	NCC. NET Audio	[Rubin]
224.0.12.0-224.0.12.063	Microsoft and MSNBC	[Blank]
224.0.13.0-224.0.13.255	Worldcom Broadcast Services	[Barber]
224.0.14.0-224.0.14.255	NLANR	[Wessels]
224.0.15.0-224.0.15.255	Hewlett Packard	[van der Meulen]
224.0.16.0-224.0.16.255	XingNet	[Uusitalo]
224.0.17.0-224.0.17.031	Mercantile & Commodity Exchange	[Gilani]
224.0.17.032-224.0.17.063	NDQMD1	[Nelson]
224.0.17.064-224.0.17.127	ODN-DTV	[Hodges]
224.0.18.0-224.0.18.255	Dow Jones	[Peng]
224.0.19.0-224.0.19.063	Walt Disney Company	[Watson]
224.0.19.064-224.0.19.095	Cal Multicast	[Moran]
224.0.19.096-224.0.19.127	SIAC Market Service	[Roy]

Table 5 Ad Hoc Multicast Addresses (Continued)

Address Range	Usage	Reference
224.0.19.128-224.0.19.191	IIG Multicast	[Carr]
224.0.19.192-224.0.19.207	Metropol	[Crawford]
224.0.19.208-224.0.19.239	Xenoscience, Inc.	[Timm]
224.0.19.240-224.0.19.255	HYPERFEED	[Felix]
224.0.20.0-224.0.20.063	MS-IP/TV	[Wong]
224.0.20.064-224.0.20.127	Reliable Network Solutions	[Vogels]
224.0.20.128-224.0.20.143	TRACKTICKER Group	[Novick]
224.0.20.144-224.0.20.207	CNR Rebroadcast MCA	[Sautter]
224.0.21.0-224.0.21.127	Talarian MCAST	[Mendal]
224.0.22.0-224.0.22.255	WORLD MCAST	[Stewart]
224.0.23.0	ECHONET	[Saito]
224.0.23.1	Richo-device-ctrl	[Nishida]
224.0.23.2	Richo-device-ctrl	[Nishida]
224.0.23.3-224.0.23.10	Telefeed	[Beddoe]
224.0.23.11	SpectraTalk	[Karhade]
224.0.23.12	EIBNet/IP	[Goossens]
224.0.23.13	TVE-ANNOUNCE2	[Dolan]
224.0.23.14-224.0.255.255	Unassigned	[IANA]

SDP/SAP Multicast Block

The multicast group range of 224.2.0.0 through 224.2.255.255 (224.2/16) is the SDP/SAP Multicast Block, which is reserved for applications that send and receive multimedia session announcements using the SAP described in RFC 2974. An example of an application that uses SAP is the Session Directory tool (SDR), which transmits global scope SAP announcements on groups 224.2.127.254 and 224.2.127.255.

SSM Block

The multicast group range of 232.0.0.0 through 232.255.255.255 (232/8) is reserved for SSM. SSM is a new extension to PIM Sparse mode that eliminates the need for the Rendezvous Point and the Shared Tree and uses only the Shortest-Path Tree to the desired sources.

A key premise of SSM is that it is the responsibility of the host application program to determine the active source IP address and group multicast address of the desired multicast flow. Using IGMPv3, the host then signals the router with exactly which *specific source* (hence the name SSM) and group that it wishes to receive. Because PIM-SSM does not use a Rendezvous Point or a Shared Tree in the 232/8 range, a host will receive traffic only from sources that it has specifically requested. This eliminates interference or denial of service (DoS) attacks from unwanted sources sending to the same multicast group. Furthermore, the lack of Shared Trees in the SSM range means that two

different sources in the Internet can source traffic to the same group address in the 232/8 range and not worry about having a group address conflict. The reason that there is no conflict is that the hosts join only the Shortest-Path Tree of the desired source. Therefore, one host can receive Stock Quotes from (S1, 232.1.1.1) while at the same time another host can watch live video from (S2, 232.1.1.1) because separate Shortest-Path Trees are being used and no common Shared Tree exists that might accidentally deliver the unwanted source.

GLOP Multicast Block

This block of addresses has been assigned by the IANA as an experimental, statically assigned range of multicast addresses intended for use by Content Providers, ISPs, or anyone wishing to source content into the global Internet. This allocation methodology, called GLOP addressing, which is defined in RFC 2770, uses the multicast group range of 233.0.0.0 through 233.255.255.255 (233/8) and provides each Autonomous System with a block of 255 statically assigned multicast addresses. Content providers who wish to transmit multicast traffic to their customers in the Internet and that have an assigned Autonomous System Number (ASN) can use multicast addresses from their block of 255 static GLOP addresses to transmit content. If the content provider does not have its own assigned ASN, usually it can lease static GLOP addresses from their Internet Service Provider.

Administratively Scoped Block

In addition to the multicast address ranges previously described, the IANA has reserved the range of 239.0.0.0-239.255.255.255 as *Administratively Scoped* addresses for use in private multicast domains. These addresses are similar in nature to the reserved IP unicast ranges (such as 10.0.0.0/8) defined in RFC 1918 and will not be assigned by the IANA to any other group or protocol. This means that network administrators are free to use multicast addresses in this range inside of their domain without fear of conflicting with others elsewhere in the Internet. However, administratively scoped addresses should not be used when sourcing IP Multicast traffic to the Internet. This is because it has become a common practice to block multicast traffic in these ranges from entering or leaving an Autonomous Domain.

The use of administratively scoped addresses also helps to conserve the limited multicast address space because they can be reused in different regions of the network. Network administrators should configure their multicast routers to insure that multicast traffic in the Administratively Scoped address range does not cross into or out of their multicast domain.

It is becoming common practice for Enterprise network administrators to further subdivide this address range into smaller geographical Administrative Scopes within the Enterprise network to limit the *"scope"* of particular multicast applications. This is used to prevent high-rate multicast traffic from leaving a campus (where bandwidth is plentiful) and congesting the WAN links. (The topic of Administrative Scoping is addressed in more detail in the section on "Administratively Scoped IP Multicast")

ADDRESS ALLOCATION

Currently, the problem of allocating multicast addresses partially depends on whether the address is to be used strictly within an Enterprise or whether the address is to be used for global Internet multicasting. In the latter case, the Enterprise must allocate a global multicast address that does not conflict with someone else in the Internet. On the other hand, if the multicast is intended to remain inside of the Enterprise, an address may be allocated from the Administratively Scoped address range (239/8) without fear of conflict with other sources in the Internet. In fact, the normal procedure is for the Enterprise network to have "multicast boundaries" configured at the borders of the network so that traffic in the 239/8 address range can neither enter nor leave the Enterprise network.

The actual allocation of multicast addresses can be accomplished in the following ways:

- Static Address Allocation
- Scope Relative Address Allocation
- Dynamic Address Allocation

By further categorizing these methods by Global and Enterprise multicast, we can develop the matrix of protocols and methodologies (shown in Table 6) that are currently applicable today.

Table 6 Address Allocation Techniques

	Global	Enterprise
Static	IANA Assignment, GLOP	Internal Assignment
Scope Relative	IANA Assignment	IANA Assignment
Dynamic	SDR, Multicast Address Set Claim (MASC), SSM	MADCAP, SSM

Each of these methodologies and/or protocols are discussed in more detail in the following sections.

Static Address Allocation Methods

Statically allocated addresses are addresses that are assigned by a controlling authority for use by protocols or applications that require well-known multicast addresses.

Global IANA Assignment

Static multicast addresses that have been assigned by the IANA are considered to be permanent in nature and globally valid; therefore, they are valid everywhere and in all networks. This permits applications and hardware devices to have these addresses "hard-coded" into their software or microcode.

One example of a static multicast addresses that has been permanently assigned by the IANA is the Local Network Control multicast address 224.0.0.5, which is the "All OSPF Routers" multicast group address. Another example is multicast address 224.0.1.1, which is an Inter-Network Control multicast address that is permanently assigned by the IANA for use by the Network Time Protocol.

Enterprise Internal Assignment

Static address allocation methods are used by Enterprise network administrators to allocate specific addresses or address ranges from the Administratively Scoped address range, 239.0.0.0–239.255.255.255. In this case, the network administrator assumes the duties of the "Controlling Authority" for address assignment within the Enterprise.

GLOP Addressing

In the late 1990s when native multicast was beginning to be deployed in the Internet, several Content Providers planned to begin multicasting some of their audio and video content. Unfortunately, the state of dynamic address allocation at that time was such that no good solutions were available that permitted the Content Providers to uniquely allocate addresses. To work around this problem, an experimental form of static address allocation was proposed by the IETF. This allocation methodology, called GLOP addressing, which is defined in RFC 2770, uses

the multicast group range of 233.0.0.0 through 233.255.255.255 (233/8). This block was assigned by the IANA and is an experimental, statically assigned range of multicast addresses intended for use by Content Providers, ISPs, or anyone wishing to source content into the global Internet.

Note: The question is frequently asked, "What does the acronym GLOP stand for?" It turns out that this is not an acronym at all. The original authors of this RFC needed to refer to this mechanism by something other than "that address allocation method where you put your Autonomous System in the middle two octets." Lacking anything better to call it, one of the authors, David Meyer, simply began to refer to this as "GLOP" addressing and the name stuck.)

GLOP addresses (as shown in Figure 5) are constructed as follows: the high order octet is always 233 (decimal), followed by the next two octets which contain the 16-bit ASN of the Content Provider or ISP that is sourcing the multicast traffic.

Figure 5
GLOP Address Format

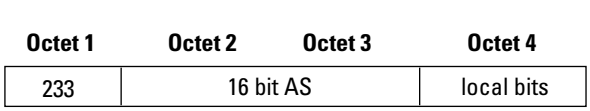

The advantage of this allocation mechanism is that for each registered Autonomous System that an entity owns, it automatically has a /24 worth of statically allocated multicast address space. No registration process is necessary because the allocation already is based on registered ASNs.

As an example of GLOP addressing, assume that Company XYZ wants to source various live video and audio multicast streams to the global Internet as part of their service offering. If Company XYZ has a registered ASN of 2109, they would be able to source this traffic using multicast addresses in the range of 233.8.61.0–233.8.61.255. (The decimal ASN 2109 converts to binary 100000111101 which, in turn, converts to 8.61 in dotted decimal format.)

It also might be the case that Company XYZ does not have a registered ASN. In that case, they could "lease" some GLOP address space from their ISP, who would allocate the leased addresses from their pool of statically assigned GLOP addresses based on their registered ASNs.

Extended GLOP Multicast Addresses

When GLOP addressing was initially introduced, several Content Providers immediately began to take advantage of this allocation method and started sourcing numerous streams of multicast video and audio into the global Internet. Unfortunately, some of these Content Providers quickly ran out of multicast addresses, as many had only a single registered ASN. As a result, they began to look for additional addresses that could be statically allocated to them for their multicast content streams.

RFC 2770, "GLOP Addressing in 233/8," explains that the GLOP address space consisting of private ASNs (64512 through 65545) in the middle two octets are reserved for future allocation. This effectively resulted in the address range of 233.252.0.0–233.255.255.255 (that is, four /16 blocks) being reserved. Given the need for more address space by some of the Content Providers, this space was a prime candidate for allocation of additional space to meet their needs. The problem was how to allocate it.

In June 2001, an "Informational" RFC was published (RFC 3138) that defined "Extended Assignments in 233/8" and referred to addresses in this range as Extended GLOP (EGLOP) address space. This RFC also states that a Regional Registry, such as the IANA, should control the assignment of address blocks from the EGLOP address space. Furthermore, applications for address space from this range must demonstrate that the request cannot be satisfied by other address means such as GLOP addressing, Administratively Scoped addressing, or SSM.

Scope Relative Address Allocation

Scope Relative multicast addresses (shown in Table 7) are static multicast addresses assigned by the IANA. These addresses are associated with Administratively Scoped multicast ranges and are used by protocols that require well-known addresses within the scope to perform their normal function. However, because the multicast addresses configured by a network administrator for an Administratively Scoped range can vary depending on network needs, Scope Relative addresses use offset address assignments. To facilitate this offset assignment approach, RFC 2365 (*"Administratively Scoped IP Multicast"*) reserves a block consisting of the highest 256 multicast addresses in every administratively scoped range for Scope Relative addressing.

Table 7 Scope Relative Multicast Addresses

Relative	Description	Reference
0	SAP (Session Announcement Protocol)	[Handley]
1	MADCAP Protocol	[RFC 2730]
2	SLPv2 (Service Location Protocol) Discovery	[Guttman]
3	MZAP	[Thaler]
4	Multicast Discovery of DNS Services	[Manning]
5	SSDP	[Goland]
6	DHCP v4	[Hall]
7	AAP (Authoritative access point)	[Hanna]
8	MBUS (Maintenance bus)	[Kutscher]
9-252	Reserved—To be assigned by the IANA	–
253	Reserved	–
254-255	Reserved—To be assigned by the IANA	–

Figure 6 lists the Scope Relative addresses that currently are assigned by the IANA. These addresses consist of a relative offset from the end of the reserved Scope Relative address block. A relative offset assignment of zero would correspond to the last address in the reserved Scope Relative address block, or x.x.x.255. Likewise, a relative offset of one would correspond to address x.x.x.254 in the reserved Scope Relative address block.

Figure 6
Scope Relative Addressing

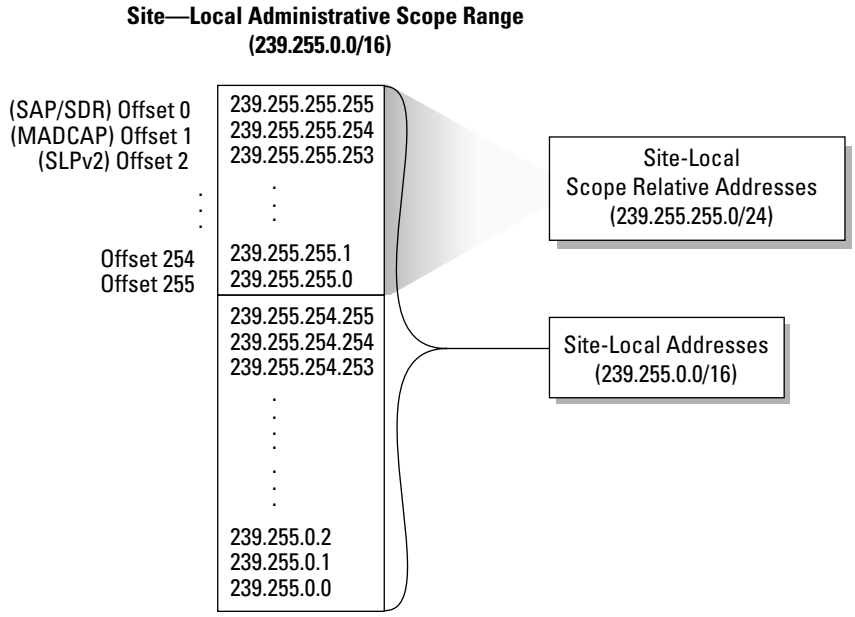

Figure 6 shows an example of Scope Relative addresses in use in the Site-Local zone consisting of the multicast address range of 239.255.0.0/16. Per RFC 2365, the subrange of 239.255.255.0/24 is reserved for Scope Relative multicast *within the Site-Local zone*. Because Scope Relative addresses are offset from the *end* of this range, Figure 6 shows that address 239.255.255.255 would be assigned for "SAP Session Announcements" and 239.255.255.254 would be assigned for use by the "MADCAP" protocol inside the Site-Local zone.

Dynamic Address Allocation

SDR—Session Directory

SDR is a multicast application that does both Dynamic Address Allocation and uses Scope Relative multicast. Although the SDR application is being phased out for other content announcement mechanisms (such as Web-based content servers), it makes a good case study for how dynamic multicast allocation initially was accomplished, as well as being a good example of how Scope Relative addressing is used. We will provide an overview of this application to complete our discussion of multicast addressing.

SDR listens for Session Announcement Protocol (SAP) packets on a well-known, global IP multicast group (for instance, 224.2.127.254). These SAP packets periodically are multicast by other SDR hosts that have *created* a multimedia session. Each SAP packet contain a session description, the time the session is active, its IP multicast group addresses, media format, contact person, and other information about the advertised multimedia session.

Figure 7 shows the SDR SAP window after it has listened to the SDR multicast group (224.2.127.254) for a period of time. Each line corresponds to a multimedia session announcement that SDR received in a SAP packet. As these packets are received, SDR updates this window by adding the multimedia session name to the list and storing this information in an SDR cache file.

Figure 7
SDR Session Announcement Window

The frequency of these SDR announcements depends on the total number of sessions being announced in the Internet. As the list of announced sessions grows, each SDR sender slows its transmission rate so that the total bandwidth consumed by SDR announcements is kept to a very low rate. The end result is that the announcement period of any one session easily can increase to several minutes.

As SDR continues to run and collect all session announcements from the Internet, it builds a database of cache files that describe every multicast session being announced in the Internet. The information in these cache files includes the multicast addresses in use that, in turn, allows SDR to determine what addresses are *not* in use. As a result, when users want to create a new session and allocate one or more multicast addresses, they click "New" at the top of the Session Announcement window shown in Figure 7. This launches the SDR "Create New Session" window shown in Figure 8.

Notice that SDR is proposing multicast address 224.2.234.131 for use with this audio session. Theoretically, this should be an unused global address that should not conflict with anyone else in the Internet.

The SDR multicast application is an example of an application that uses Scope Relative multicast addresses. SDR uses SAP to announce the existence of multimedia multicast conferences to other workstations in the network running the SDR application. Normally, SDR would announce these sessions by multicast to the well-known multicast group 224.2.127.254, which would travel throughout the Internet because it falls in the global Internet scope. However, when a multimedia conference is created inside a designated scoped zone such as the Site-Local zone, SDR would announce this session using the SAP scope relative address. In the case of the Site-Local zone, this would

correspond to the scope relative multicast address of 239.255.255.255. Assuming that multicast boundaries were in place around the Site-Local zone to block traffic in the 239.255.0.0/16 multicast address range from leaving the Site-Local zone, these Site-Local SDR announcements to 239.255.255.255 would not leave the zone.

Figure 8
SDR Create Menu

Continuing with this SDR example of Scope Relative multicast addressing, assume that an Enterprise has deployed both Site-Local and Organization-Local (for instance, Enterprisewide) scoped zones in the 239.255.0.0/16 and 239.192.0.0/14 address ranges, respectively. In this case, an SDR workstation would use a Scope Relative multicast address of 239.255.255.255 to announce Site-Local sessions and a Scope Relative multicast address of 239.192.255.255 to announce Organization-Local (Enterprisewide) sessions[1].

SDR served its purpose in the early days of the DVMRP-based Multicast Backbone, or MBone for short, where the primary multicast content was limited to a few multimedia conferences.

1. Manual configuration of the SDR application is required to make it aware of the scoped zones in which it resides for it to make use of scope relative SAP announcements. Otherwise, it will default to using the SAPv1 global address of 224.2.127.254 to announce all sessions.

SDR is being phased out by most Enterprise network administrators as the preferred method of multimedia session announcement in deference to other methods, such as well known Content Managers (IP/TV) or Web links. Therefore, until the take up of a zone protocol like MZAP (RFC 2776), these assignments have historically only been used in the Site-Local scope (with some minor exceptions), although good practice is to reserve the relative assignments within each scope for future use.

SDR has some serious drawbacks. First, unlike the hierarchical address space of the DNS, SDR uses a flat address space that cannot scale to the projected numbers of content sources in the Internet. This is particularly true if you consider that as the number of announced sessions grows, the period of the announcements also grows. Finally, not everyone who wishes to source multicast content to the Internet wants or needs to use SDR.

MASC

As far as "global dynamic multicast address allocation" methods go, the only other *proposed* solution is MASC. MASC is a hierarchical address allocation protocol that is defined in RFC 2909.

Figure 9 shows how MASC will work at a very high level. At key locations in the Internet (possibly at certain Internet Exchange points), top-level MASC nodes will reside in the top level, or root MASC domain. The MASC nodes in the next lower domain in the hierarchy are connected to the root-level nodes in a parent-child relationship. These second layer domain MASC nodes are connected to the MASC nodes in the third layer of hierarchy, and so on. All of these connections use TCP and are configured in a similar fashion as Border Gateway Protocol peer connections.

The root domain MASC nodes are responsible for the allocation of the global multicast address range into smaller contiguous blocks of addresses to the MASC nodes in the next lower MASC domain. Currently, it is envisioned that this first level domain of MASC nodes primarily would consist of first tier ISPs.

Figure 9
MASC Hierarchy

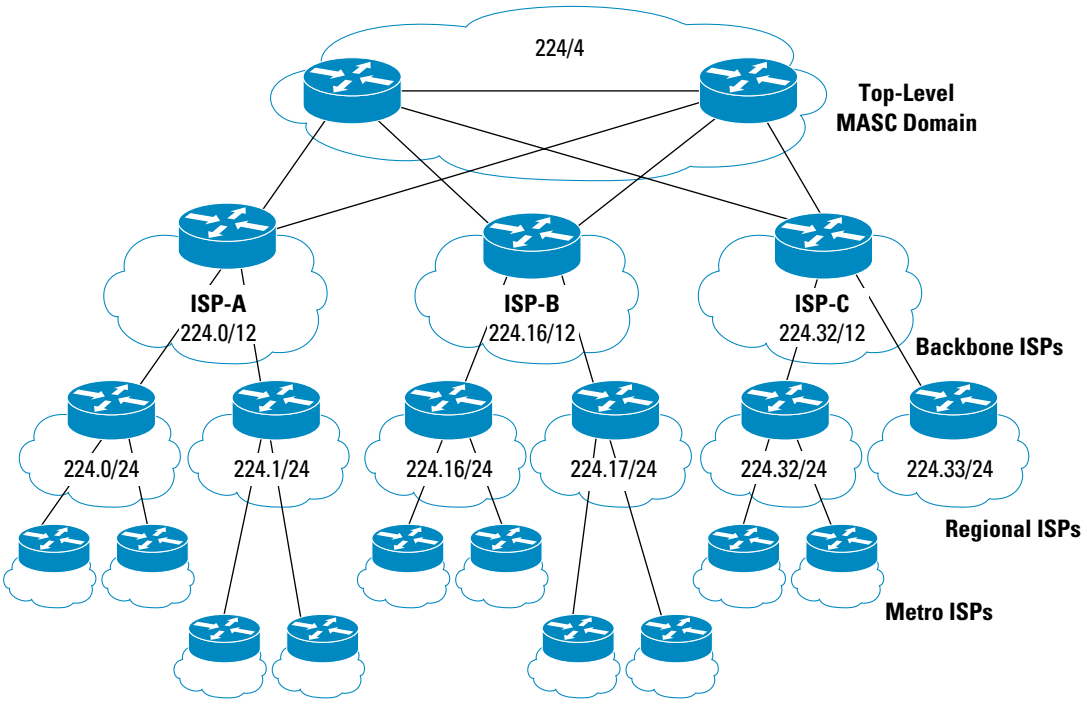

Each of the MASC nodes in the first tier ISPs would send a *Set-Claim* message containing a requested range of multicast addresses to the root domain MASC nodes. If this range is not available, the root domain MASC nodes would acknowledge the Set-Claim message and mark the range as allocated in its database. If a portion of the range already has been successfully claimed by another MASC node in another domain, the root domain node will propose an alternative range of available addresses to the requestor. This back and forth negotiation continues until the requestor had successfully claimed an acceptable range of addresses (or finally timed out).

After a MASC node has claimed a range of addresses, it can use these addresses in its internal network as well as allocate smaller contiguous blocks of addresses to lower tier domain of MASC nodes as shown in Figure 9. (The actual protocol that an end station will use inside of a domain to request a multicast address from a MASC server or one of its agents will be MADCAP. This protocol is discussed in the next section.)

Because a successful MASC Set-Claim on a range of addresses typically is valid only for a finite period of time, a MASC domain periodically must renew its claim. If the parent MASC domain experiences address exhaustion, it may reduce the size of a lower tier domain range in an attempt to free up some address space. Alternatively, a MASC domain that is beginning to experience address exhaustion also can issue additional Set-Claims to its parent domain in an attempt to allocate more space.

MASC is a nontrivial protocol, which must be carefully designed and implemented to avoid causing massive fragmentation of the limited resource of multicast addresses. If serious fragmentation does occur, one quickly can envision one of the most incredibly complex, distributed, *garbage collection* problems that the computer industry has seen to date.

Unfortunately, the deployment of MASC in the Internet simply has not happened. This may be partly due to the fact that the protocol is so complex. Additionally, the introduction of GLOP, EGLOP addressing, as well as the new Source-Specific Multicast extension to PIM Sparse mode has resulted in other solutions to the global address allocation problem, at least for the short term.

MADCAP

Multicast Address Dynamic Client Allocation Protocol (MADCAP) allows a client workstation to "lease" a multicast address from a MADCAP server in a manner similar to how it "leases" an IP address from a DHCP server.

When a MADCAP client workstation wants to "lease" a multicast address, first it must locate the nearest MADCAP Servers. This is accomplished by multicasting a MADCAP DISCOVER message to the MADCAP Scope Relative multicast address (-1) in the Site-Local scope, (for instance, 239.255.255.254). (**Note: This is why it is important to adhere to the conventions for the Site-Local scope defined in RFC 2365.**) Any MADCAP Servers that hear this Scope Relative multicast message and wish to participate in the allocation process identify themselves by sending back an OFFER message to the MADCAP client. After the client has discovered the appropriate MADCAP server, it can send REQUEST messages to the server to request the "lease" of a multicast address from the server.

MADCAP supports the concept of Administratively Scoped Zones. Workstations can request a list of known scope ranges by sending a GETINFO message to one or more MADCAP servers, which respond by sending back a list of its configured multicast scope ranges. This allows the MADCAP client to select the scope range that is most appropriate to the application and subsequently request an address from this range in a REQUEST message.

To use MADCAP for multicast address allocation in an Enterprise network, the following two requirements must be met:

- The multicast application program must be written to make use of MADCAP APIs.
- The network administrator must deploy and configure one or more MADCAP servers in the network.

The first requirement normally is beyond the control of the network administrator and, to date, few applications exist with this capability. However, if there are at least one or more applications being deployed (or expected to be deployed) in the network that will use MADCAP, it is good practice for the network administrator to deploy MADCAP servers in the network.

Beginning with Microsoft Server 2000 operating system, MADCAP Server functions have been added under the DHCP Service. Figure 10 shows an example of a "Campus" scope Multicast Scope range being configured in a Microsoft 2003 Server.

Figure 10
MADCAP Scope Configuration in Windows Server 2003

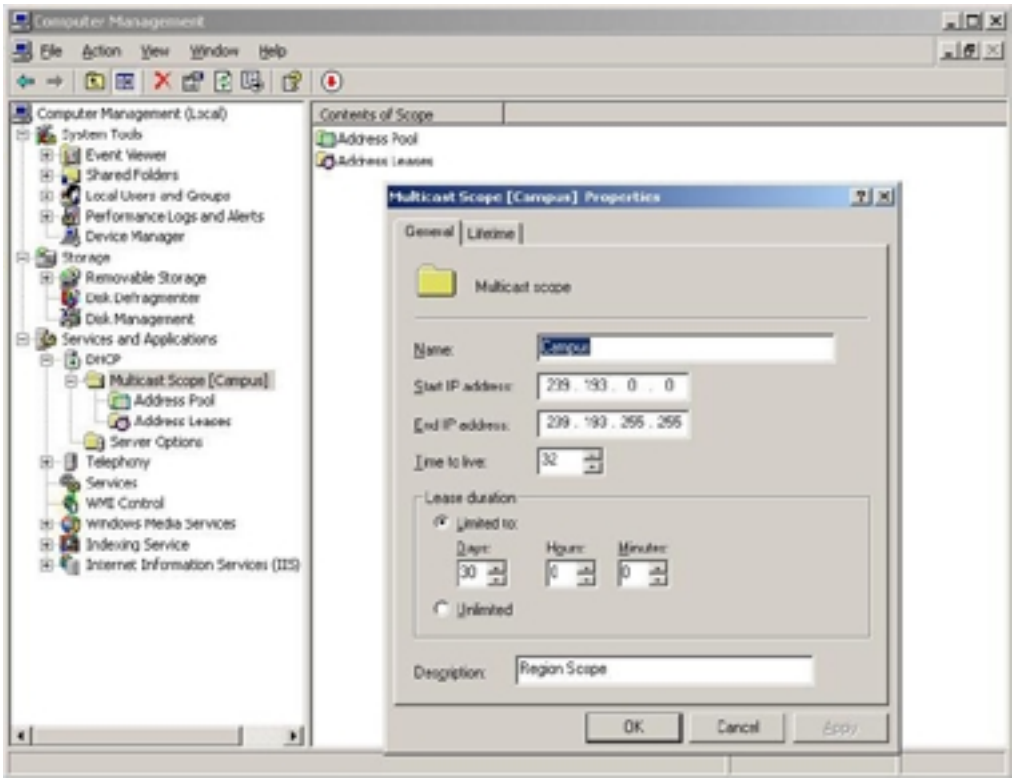

Notice that Microsoft has chosen to include the MADCAP configuration functions inside the DHCP Service even though MADCAP is a separate protocol. This is probably because the concepts of configuring a MADCAP Multicast Scope are like the configuration of a DHCP scope. However, other operating system implementations of MADCAP may take a different approach.

Addresses Allocation Considerations

The allocation of IP multicast group addresses is a complex process. Allocation of addresses needs to take into account multiple factors such as:

- **Size of the organization now and in the future**—Sufficient address allocation or expansion should accommodate future growth and acquisitions or mergers.
- **Organizational structure and relations between Business Units**—There may be specific demarcation points of administrative control within an organization that need to be considered when allocating addresses. Each Business Unit may require its own address range.
- **Scale of the IP multicast deployment currently and in future**—Many organizations underestimate the growth of multicast throughout the business and do not create a comprehensive addressing scheme from the beginning. This often can lead to readdressing at a later stage (in the same way as unicast addressing).
- **Internal policies on the control and deployment of network applications**—Depending on the network topology, certain constrains and protection mechanisms may need to be put in place to protect network resources. A well-summarized address range can help to simplify this process.
- **Scope of the applications**—Will they be mostly local to a site, city, or country? How does this relate to the network topology? Will quality of service need to be put in place for specific application?
- **Security policy**—Many multicast application lack real security. Organizations may find the need to impose restrictions by other means. A carefully designed addressing plan and Rendezvous Point scoping will aid the implementation of these restrictions if required. By using the Administratively Scoped addresses defined in RFC 2365, these addresses cannot traverse the Internet. This prevents outside sources from accessing organization data using multicast addresses.
- **Application flexibility**—It is unfortunate that some multicast applications that clearly are not intended to be run Enterprisewide (or on the global Internet), often make use of hard-coded multicast addresses and have no provision to reconfigure the application to use another multicast address that is more applicable in scope. This can create serious issues when attempting to develop a rational address scoping plan, particularly if the application is hard coded to use an address that does not fall within the desired scope.
- **Avoiding problematic address ranges**—Multicast addresses that map to MAC addresses in the 0x0100.5E00.00xx range normally are flooded by Layer 2 switches. These addresses should be avoided.
- **Readiness for future use of new multicast delivery methods such as Bidir PIM and SSM.**

Overall, many different requirements need to be considered. There is not a single best way to allocate multicast addresses for all organizations. Administrators need to take their own unique requirements into account and design the best addressing policy for their needs.

The following sections outline some examples that should be considered. The final addressing decision probably will be made up of some or all of these considerations ordered in the most appropriate way for your organization requirements.

Allocation Based on Application Scope

The most common method of allocating address is by the geographical scope of the application. Figure 11 shows a simple three-layer geographical scope based on Sites, Regions, and the entire Organization (for instance, the entire enterprise network). Multicast boundaries are configured around the edges of each site such that "Site" multicast traffic in the range of 239.193.0.0/16 will not leave the site. Similarly, a multicast boundary around each region prevents "Region" multicast traffic in the range of 239.194.0.0/16 from leaving the region. Finally, a multicast boundary is configured around the entire enterprise network so that "Organization" (for instance, Enterprise) multicast traffic in the range of 239.195.0.0/16 is not permitted to leave the Enterprise network.

Figure 11
Application Scope Example

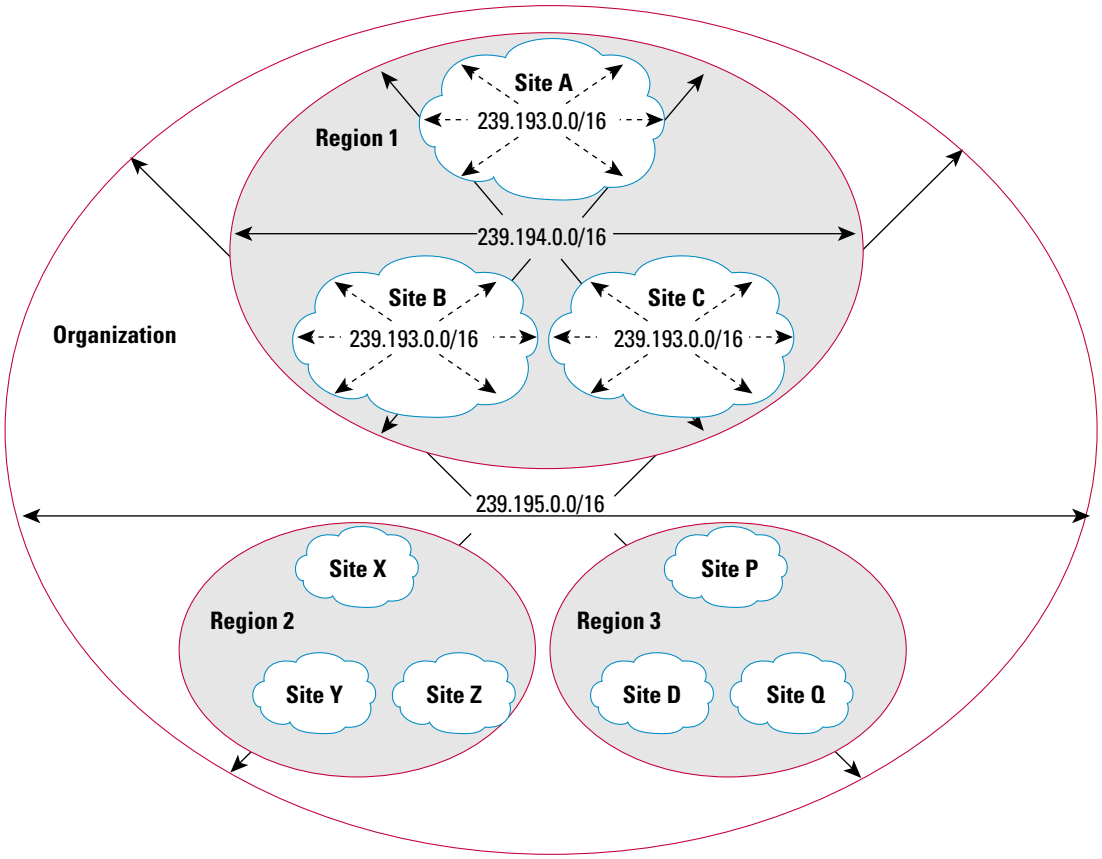

Allocation Based on Multicast Application Rates

Addresses may be subdivided into application data rates. This approach normally is combined with a geographical scope to prevent high rate traffic streams from congesting WAN links. It is generally a good idea to make the boundaries of these bandwidth based scopes fall on geographical scope boundaries to implement this and minimize the number of scopes in the network. This subdividing normally would be required only on regional and enterprise local scopes because that is where the potential for congestion resides. Table 8 shows a four rate bandwidth scoping example. The actual number of groups and the specific rates depend on the organization network topology.

To successfully deploy Bandwidth Scoping it is necessary that each multicast application is fully understood in terms of bandwidth before deployment and that the proper multicast addresses appropriate to the desired bandwidth is used by the application.

Table 8 Bandwidth Scoping

Scope Boundary	Bandwidth	Example or Rate Range
Building	Very High Rate	>2 Mbps
Site	High Rate	768 Kbps—2 Mbps
Region	Medium Rate	64 Kbps—768 Kbps
Enterprise	Low Rate	< 64 Kbps

Allocation Based on Type of Multicast Protocol

Addressing may want to encompass a method of identifying the type of multicast delivery method that is being used. For example, this could be PIM Dense mode, PIM Sparse mode, SSM, or Bidir PIM. This may be an important consideration to allow preconfiguration of the network. Cisco recommends that, at a minimum, address ranges are reserved for SSM groups, which are expected to gain in popularity as IGMP Version 3 becomes more common in host systems.

SUMMARY

Given the scarcity of available multicast addresses, it is understandable that some commercial organizations have tried to pressure the IANA into "reserving" blocks of multicast addresses to guarantee that they will be able to source multicast content at will. Some of the address assignments listed previously in this chapter show that several commercial organizations have succeeded in accomplishing this despite the IANA's attempts to resist this sort of misallocation of a valuable Internet resource.

In some cases, application developers have taken the "hard-wired" multicast address approach to the development of client-server application software to avoid the work of the address configuration step of the client workstations. (This does not include legitimate network protocols that require "well-known" multicast addresses to function properly.) The use of hard-wired multicast addresses in application software should be weighed against the loss of flexibility and usability of the software itself. (Consider what would have happened if the authors of the multimedia, multicast conferencing tools such as the Robust Audio Conference Tool (RAT), and the Video Conference Tool (VIC) would have hard wired these applications to a single fixed multicast address.)

In summary, most of these static assignments are unnecessary because of the allocation methods such as GLOP addressing and Administratively Scoped addressing. In addition, when these allocation methods are combined with new multicast models such as Source-Specific Multicast that virtually guarantee conflict free addressing, the need for organizations to "block-out" ranges of IPv4 multicast addresses so that they can source content nearly has been eliminated. In addition, current traffic statistics show that there is little activity in the Internet on these multicast addresses assigned to specific organizations. Some might argue that the reason that these assigned addresses are not seen in the Internet is that they are being used to distribute content inside of private domains. If this is true, there is no reason that Administratively Scoped addressing could not be used inside the domain to accomplish the same thing without wasting valuable Internet address space. These reservations could and should be released.

ADMINISTRATIVELY SCOPED IP MULTICAST

Before we proceed with the discussion of Administratively Scoped IP Multicast, it is important to understand that *all* IP Multicast is scoped. Even multicast that falls into the range of 224.1.0.0–238.255.255.255 that is destined for the global Internet has a scope; that is, the entire Internet. Multicasts that fall into the IANA's Administratively Scoped Block of 239/8 as described in previous "Administratively Scoped Block" section are considered private addresses that should be used by Enterprise networks and never should be forwarded outside the Enterprise. This is the next level of multicast scoping below the global Internet scope and is typically the initial implementation of scoping that an Internet multicast connected Enterprise would encounter because they should block the 239/8 address range from entering and leaving their domain.

The following sections address the concept of multicast address scoping in more detail and provide some of the fundamental concepts employed to further "scope" multicast traffic beyond the simple one to two scope model mentioned previously. Nearly all Enterprise networks eventually will require the use of these additional levels of scoping to scale multicast within their Enterprise.

RFC 2365—Administratively Scoped Addresses

RFC 2365 provides limited guidelines on how the multicast address space can be divided and used privately by enterprises. The terminology "Administratively Scoped IPv4 multicast space" relates to the group address range of 239.0.0.0 to 239.255.255.255. The key properties of Administratively Scoped IP Multicast are that:

- Packets addressed to Administratively Scoped multicast addresses do not cross configured administrative boundaries. The limits of these scope boundaries often are called "Zones" or "Scoped Zones."
- Administratively Scoped multicast addresses are locally assigned and are not required to be unique across administrative boundaries.

Table 9 and the bullet items that follow it summarize the recommendations of RFC 2365.

Table 9 Administratively Scoped Addresses 239.0.0.0/8

Range	Description	Reference
239.000.000.000-239.191.255.255	Organization-Local Scope Expansion Space	[Meyer, RFC 2365]
239.192.000.000-239.195.255.255	Organization-Local Scope	[Meyer, RFC 2365]
239.195.000.000-239.254.255.255	Site-Local Scope Expansion Space	[Meyer, RFC 2365]
239.255.000.000-239.255.255.255	Site-Local Scope	[Meyer, RFC 2365]

- Organization-Local Scope addresses are recommended for private use within an organization for intersite applications that will be run regionally or globally.
- The address range numerically below the Organization-Local Scope is intended as the expansion space for the Organization-Local Scope. Organizations can allocate or subdivide this range as needed either to extend the Organization-Local Scope or to create other geographically smaller subscopes within the Enterprise.

- Site-Local Scope addresses represent the smallest possible scope in the network. More applications are being developed that default to using this scope (unless configured otherwise) to insure that the scope of their application is limited to the smallest scope size. This is why it is important to adhere to RFC 2365 guidelines for the Site-Local Scope.

Note: It is unfortunate that many applications do not behave in this manner and instead often default to using addresses in a *global* scope instead. This results in their application traffic being multicast *far* beyond where it is desired by the network administrator.

- The address range numerically below the Site-Local Scope is intended as expansion space for the Site-Local Scope. Organizations can allocate these ranges for private use if they exceed the 239.255.0.0/16 Organization-Local range.
- The Site-Local Scope must not span any other scope boundary and must be completely contained within or equal to any larger scope.

In general, the recommendations in RFC 2365 provide only basic information regarding how the administratively scoped address space can be allocated by an enterprise. To date, few applications assume that the address space is carved up in this manner. Therefore, network administrators can allocate this address space as they see fit to build the scope hierarchy that best fits their needs. The exception to this is the Site-Local Scope, which merits special attention and is discussed in more detail later in the Site-Local Scope section.

Finally, it is important to note that the common practice is to reuse scope address ranges. For example, assume that the address range of 239.193.0.0/16 has been designated as the Campus Scope address range. This range would be used at all Campus sites in the network. This practice has the advantage of saving address space and simplifying network administration. (This does mean that if a local application needed to be extended to reach other domains the only option would be to change the application multicast address. This may or may not be a concern.)

Organization-Local Scope

The Organization-Local scope, 239.192.0.0/14, is defined by RFC 2365 as *"the space from which an organization should allocate sub-ranges when defining scopes for private use."* In general, this space (239.192.0.0–239.195.255.255) and its expansion space numerically below the Organization-Local range (239.0.0.0–239.191.255.255) are available for use by the network administrator for private multicast within the private enterprise. The above quotation (taken from RFC 2365) means that this range is to be used to define other scopes smaller than the Enterprise Scope such as a Region Scope, Campus Scope, Building Scope, and so on.

RFC 2365 provides an Organization-Local expansion space (239.0.0.0-239.191.255.255) that can be used to expand the Organization-Local space. However, the RFC states that this range *"should be left unassigned until the 239.192.0.0/14 space is no longer sufficient. This is to allow for the possibility that future revisions of this document may define additional scopes on a scale larger than organizations."* Although the likelihood of this occurring is small, this address space should be unused in the enterprise network. In addition, most Layer 2 switches flood all multicast that maps to the MAC address range of 0x0100.5E00.00xx. The 239.0.0.0/24 and 239.128.0.0/24 subranges of the Organization-Local Expansion space corresponds to this MAC address range. *Therefore, avoid the 239.0.0.0/24 and 239.128.0.0/24 subranges.*

Allocating the Organization-Local Space

If we assume that each scope will be allocated a Class B sized (/16) block of addresses, the size of the Organization-Local Scope range allows a network administrator to define up to four different scope sizes. This is generally sufficient for all but the largest networks. If this not is insufficient, alternative allocations strategies can be used, such as using smaller block sizes of /17 down to /25 for each scope. As an example, reducing the scope size down to a /18 range would allow 16 scope ranges to be defined each with a complement of 4096 multicast addresses. Another alternative is to use variable sized ranges. Typically, the larger the geographic scope, the smaller the number of actual multicast addresses that are needed for the scope.

When using smaller allocation sizes for scopes, remember that Scope Relative Addressing reserves the upper 256 address of any scope range. The smallest block size that should be used is a /25 allocation because this provides a total of 512 addresses—256 assignable multicast addresses plus 256 reserved Scope Relative addresses.

Site-Local Scope

Per RFC 2365, the Site-Local Scope, 239.255.0.0/16, is the smallest possible scope. This means that all other scopes must be equal to or greater in size than the Site-Local Scope. Although network administrators can architect the scoping hierarchy within the Enterprise in any manner that seems appropriate, it is advisable to adhere to this one key principle from RFC 2365. This may mean that for a particular hierarchy the Site-Local Scope range of 239.255.0.0/16 actually may be applied to a smaller area than an enterprise "site" such as a building. However, this SHOULD be the smallest scoped zone in the hierarchy.

This is important because more applications are being written that make the assumption that the 239.255.0.0/16 range is the smallest scope and attempt to locate specific services within this scope before searching outside of the scope. Applications written by Microsoft to use MADCAP make this assumption and under certain circumstances will attempt to locate the nearest MADCAP server using a Relative multicast address within the Site-Local Scope. Network administrators should maintain the Site-Local Scope as the smallest scope within their scope hierarchy.

Allocating the Site-Local Space

To maintain the Site-Local Scope as the smallest scope and yet allow for new, possibly smaller scopes to be defined, it is good to keep the Site-Local Scope separate and not assign this address range (239.255.0.0/16) to any other scopes. For example, assume that the smallest scope initially desired in a particular network is the "Campus Scope." Because this is the smallest scope, the Site-Local Scope cannot be larger than this scope. Although the Campus Scope and the Site-Local *could* be assigned to the same scope range (for instance, the Campus Scope assigned the 239.255.0.0/16 range), it is much better to define separate Campus and Site-Local Scopes using different group ranges, although identical in geographical scope. This provides flexibility if there is a future need to define a smaller scope—say a Building Scope—that is smaller than the Campus Scope. For example, if the Campus Scope is assigned the Site-Local address range and a smaller scope is needed, this new reduced size scope would have to be assigned to the Site-Local Scope range, 239.255.0.0/16, and a new group range would have to be assigned to the original Campus Scope. This would require all existing Campus Scope applications to be moved to this new group range, which could cause problems.

Administrative Scoped Zones Example

The following example shows three different levels of Enterprise scopes and three levels of boundaries. This normally is sufficient for most organizations. The following terms are used:

- **Campus Scope**—Identifies applications local for buildings or building complexes (from single building up to Campus network). Generally the physical topology that this scope covers depends on the network and business topology. For example, if several buildings are interconnected using high-speed links and business divisions span these buildings, they may be considered as a single Campus Scope.
- **Regional Scope**—Identifies applications for intraregion IP multicast (for example, Europe, Middle East, and Africa (EMEA), North America, South America, Asia Pacific, etc.).
- **Enterprise Scope**—Identifies applications for interregion IP Multicast communication between any sites within the organization.

To deploy this scoping strategy, the Campus, Regional, and Enterprise Scopes are subscopes of the Organization-Local scope (239.192.0.0/14) defined by RFC 2365, while a separate Site-Local Scope is deployed using group range 239.255.0.0/16, also per RFC 2365. The address allocations for these ranges are shown in Figure 12.

Figure 12
Scoped Zone Address Ranges

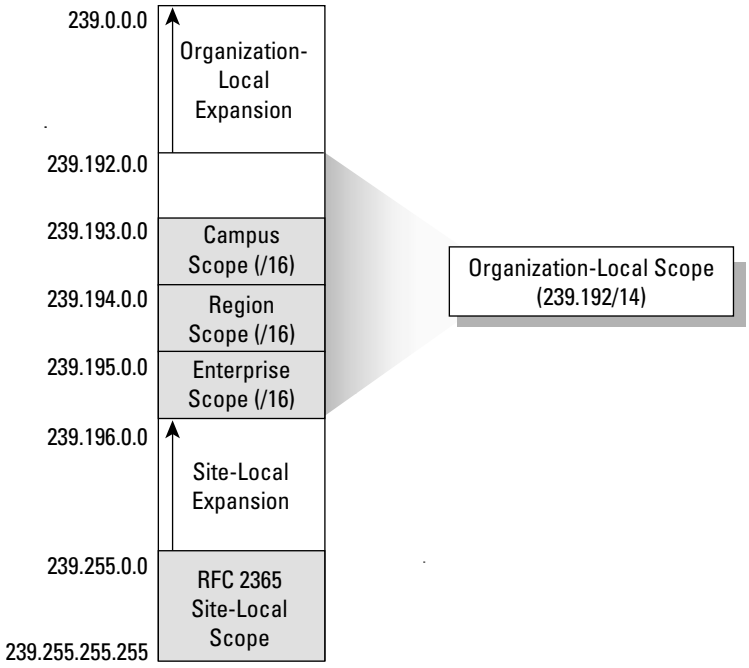

Notice that the Enterprise Scope is allocated at the high end of the Organization-Local Scope address space at 239.195.0.0/16, and the Region and Campus Scopes also are /16 ranges allocated from the Organization-Local space immediately below the Enterprise range. Although this example has allocated the Enterprise, Region, and Campus Scopes in order of decreasing scope size, it is not a requirement. In fact, the Campus and Region ranges could be reversed without problems. In addition, the size of these address ranges are only a suggestion and individual network administrators may use different group range sizes and expansion ranges. However, to keep the Scope Relative

address block for the Organization-Local range (239.195.255/8) the same as the Scope Relative block for the Enterprise range, *the Enterprise range should be allocated at the top of the 239.192/14 block*. Remember, the Enterprise scope extends to the edge of the network; therefore, its boundaries are identical to the Organization-Local Scope defined in RFC 2365. It is important to maintain this relationship so that any application that wishes to make use of Scope Relative addressing in the Organization-Local Scope will assume that this address space falls in the range of 239.195.255/8, which is the highest that 256 addresses in the Organizational-Local group range. By assigning the Enterprise Scope to the 239.195/16 range, we maintain this important relationship as shown in Figure 13.

Figure 13
Correct Placement of the Enterprise Scope Relative Range

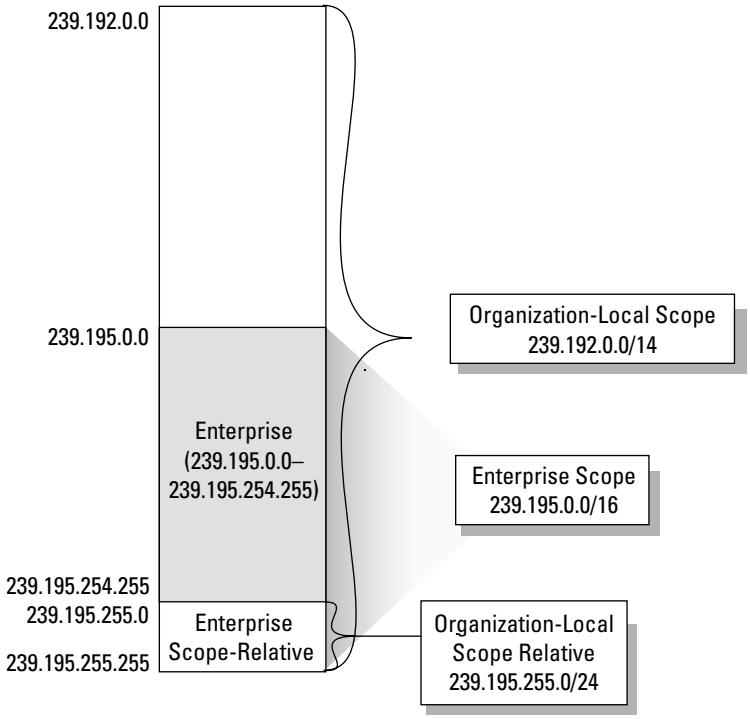

Network administrators often have chosen to use the Site-Local range for the smallest scope in use. (In this example, the Campus Scope would have been configured to use the Site-Local address block 239.255/24 instead of having its own separate address range.) However, it is Cisco's recommendation that this not be done. Instead, it is recommended that the Site-Local Scope range of 239.255.0.0/16 be kept separate and independent from the address range used for the smallest scope. (Remember, there are applications that use the well-known Site-Local address range and assume that this is the physically smallest scope in the network. The MADCAP protocol is just one example.) In this case, the address range of 239.193/16 is used for the Campus Scope. Note, however, that in this example the Campus Scope is currently the smallest defined scope and therefore its boundaries would be identical to the boundaries of the of the Site-Local Scope. The separation of the smallest scope and the Site-Local Scope is important because it allows additional smaller scopes to be defined without readdressing the Campus Scope and yet still comply with RFC 2365.

For example, assume that at some point in the future we want to define a new "Building" Scope that is smaller than the Campus Scope. If we were using the Site-Local address range for the Campus Scope, we would have to reassign this scope to a different range (as well as moving all Campus Scoped applications to this new range) to assign the

Site-Local address range to the new smaller Building Scope to comply with RFC 2365. Using the Cisco recommended approach of assigning the smallest scope in its own address range independent of the Site-Local Scope range, we can avoid this issue. One way to accomplish this is to allocate the address range of 239.192/16 for the new Building Scope, as shown in Figure 14, and move the boundaries of the Site-Local Scope to match the boundaries of the smallest scope, which is now the Building Scope.

The drawback to assigning the new Building Scope to this new 239.192/16 block from the Organization-Local range is that it requires coordination at the Enterprise administrator level to insure that all network administrators (Region, Campus, etc.) are in sync with the address use plan. If only one or two Campus sites in the Enterprise require Building Scopes, it may be better to allow Campus administrators at these locations to allocate the Building Scope range from the Campus Scope range. However, care must be taken to redefine the Campus range access control lists (ACLs) so that they do not overlap the newly allocated Building range, because this will cause problems with the Candidate-RP filtering at multicast boundaries if the Auto-RP **filter-autorp** option is in use on the **ip multicast boundary** interface command.

Figure 14
Adding a Smaller Scope

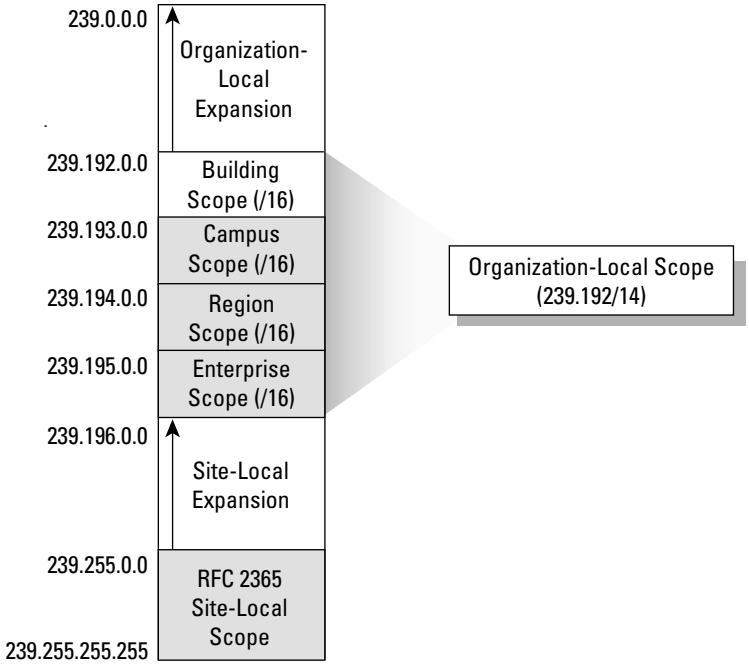

The key points to follow are:

- The Site-Local Scope should be assigned to address range 239.255.0.0/16 and should not be used by any other operational scopes.
- The Site-Local Scope should have the same *boundaries* as the smallest operational scope.
- The address range for operational scopes (except for SSM) should be allocated from the Organization-Local group range 239.192.0.0/14 or its expansion space.

- The Enterprise Scope should be allocated at the top of the Organization-Local address space.
- All scopes should be assigned a unique (that is, nonoverlapping) group range.

When possible, it is recommended that applications use logical names, corresponding to dynamically assigned group addresses, where possible. This allows applications groups to be changed from a central point in terms of not only address but by scope. Options that could be investigated to achieve this include DNS, MADCAP (RFC 2730), RFC 2908, and Microsoft MDHCP (Multicast Dynamic Host Configuration Protocol) extensions.

Figure 15 shows an example network along with the associated Enterprise, Region, and Campus Scopes. The Campus Scope is used for IP Multicast applications that remain within the Campus. The same IP Multicast range is used within each Campus and should be blocked on the Campus boundaries into the region. The Region Scope is used for IP Multicast applications that remain within the Region. The same IP Multicast range is used within each Region (normally local campus to local campus within the region) and should be blocked on the Region boundaries into the Enterprise. Finally, the Enterprise Scope is intended for enterprisewide (interregional) IP Multicast applications (for instance, global to the entire company). This range will be blocked from leaving the enterprise at any border routers that connect to any outside networks including the Internet.

Figure 15
Administrative Scoped Zone Example

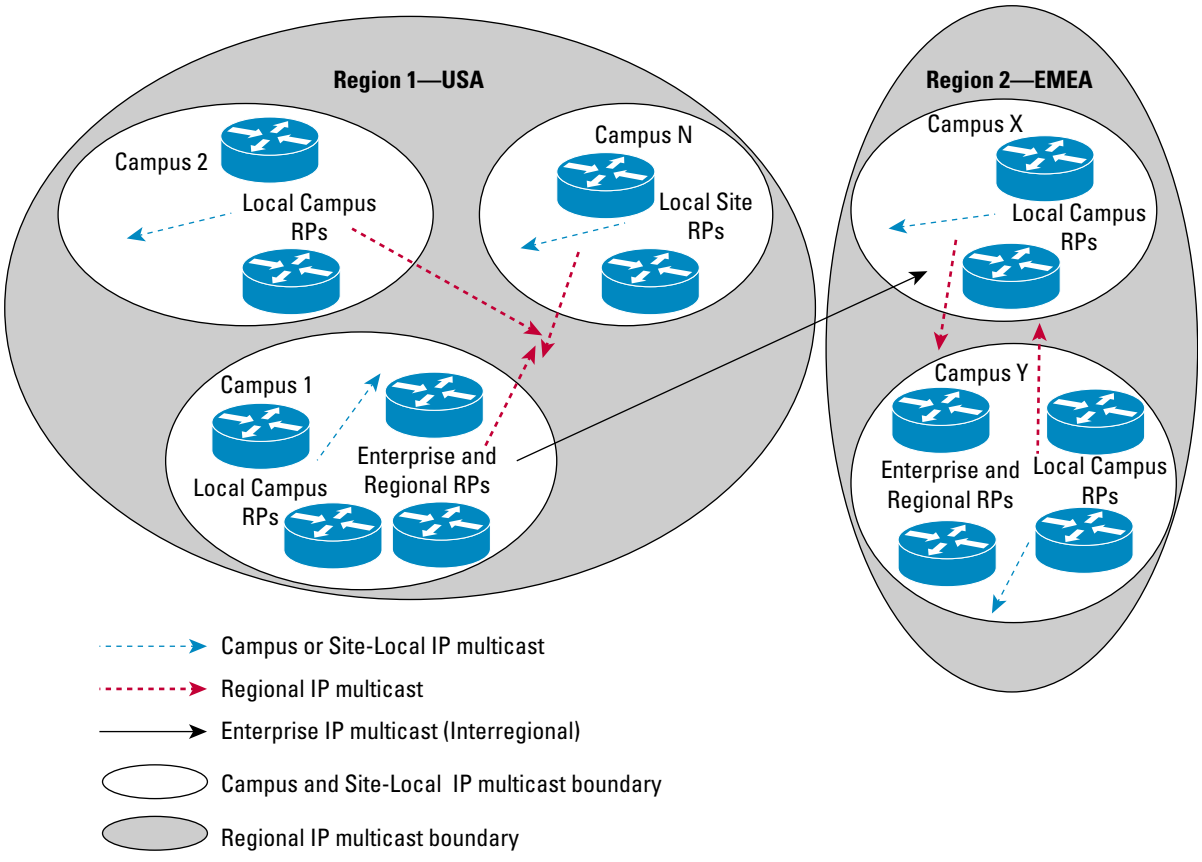

Extending the Administrative Scope Model

The model described in the previous sections assumes that only classic PIM-SM is in use. However, most networks will benefit from the use of either SSM or Bidir PIM or possibly both. This allocates address space along with the interaction with Administrative Scopes. These topics are covered in the following sections that describe how the previous model can be extended to cover these protocol extensions.

Bidir PIM

Bidir PIM (also called Bidir) is an extension to classic PIM-SM that uses only the Shared Tree to transmit multicast traffic. The Shared Tree operates in a bidirectional fashion, allowing multicast sources to transmit their traffic to the Rendezvous Point by forwarding the multicast traffic *up* the branch of the Shared Tree. This is important for applications in the Many-to-Many or Many-to-Few categories, where the number of sources in the group can grow quite large. The advantage to Bidir PIM is that only a single (*, G) multicast forwarding entry is required to forward the group's traffic regardless of the number of active sources in the group.

In Cisco IOS, Bidir groups are designated by the addition of the bidir keyword to either Auto-RP or BSR Candidate-RP configuration commands or to static RP definitions as shown in Example 1:

Example 1 Bidir Group Configuration Commands

Auto-RP

```
ip pim send-rp-announce loopback0 scope 32 group-list bidir-groups bidir
```

BSR

```
ip pim rp-candidate loopback0 group-list bidir-groups bidir
```

Static

```
ip pim rp-address 192.1.1.1 bidir
```

An easy way to add Bidir operation to our existing model is to designate some portion of each scope address range to Bidir. Figure 16 shows a /17 allocation for Bidir operation from the beginning of the Enterprise Scope range. (Note that the Enterprise Scope Relative range remains at the top of the Organization-Local block at 239.195.255/8.) This technique would be repeated for the Region and Campus Scopes in our example.

Figure 16
Bidir Range Allocation Within the Enterprise Range

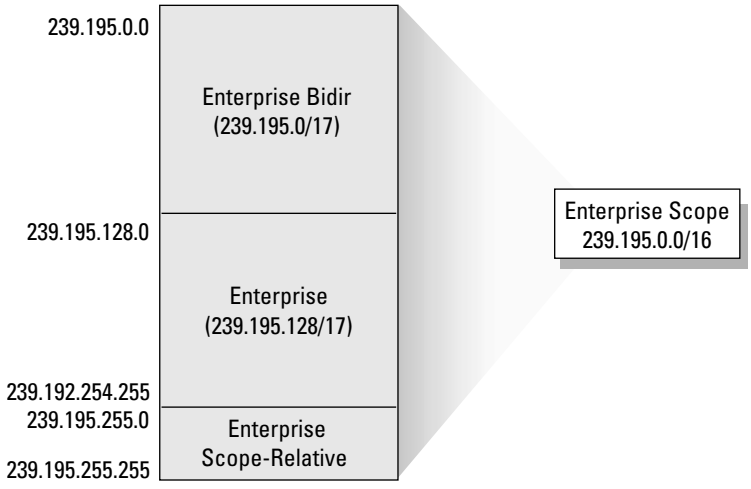

Initially, Bidir PIM may not be deployed in the network. However, it is a good idea to plan for its future deployment and it is recommended that addresses in each scope range be allocated from the top down. Using the Enterprise Scope as an example, addresses are allocated beginning at address 239.195.254.255, which is the highest address in the Enterprise group range without getting into the reserved Scope Relative range in the Enterprise Scope. If this strategy is followed, the Enterprise Bidir range later can be allocated as shown in Figure 16 with a minimum chance of any previously assigned enterprise groups falling into the newly defined Enterprise Bidir range.

SSM

At the request of the IETF, the IANA has reserved the group range of 232/8 for SSM. Because this range falls outside of the 239/8 private address space, one might assume erroneously that SSM was intended for use only in the global Internet scope. However, this is not true, because SSM has numerous advantages that simplify and improve multicast security and scalability. Application developers and network administrators should consider deploying SSM within the enterprise network. After the decision to deploy SSM within the enterprise network has been made, the combination of SSM *and* Administrative Scoping poses some interesting scenarios. First of all, the network administrator must decide which address ranges are to be enabled for SSM mode within the network. If we desire either to send or to receive SSM traffic to and from the global Internet, the default range of 232/8 must be enabled for SSM. However, assume the network administrator also wants to support SSM at some scope smaller than the global Internet scope, perhaps at the Enterprise (Organization-Local) Scope or even smaller scopes within the enterprise network. Because of the way SSM works, there is no reason that the same 232/8 SSM address range could not be used by SSM applications within the enterprise network for private enterprise-only SSM multicast. However, if this approach is taken, it is difficult to configure the boundary routers at the edge of the enterprise to prevent hosts outside of the network from joining the Shortest Path Tree for internal SSM sessions in the 232/8 range while at the same time allowing other public SSM sessions to cross the boundary. It is better to dedicate a portion of the private 239/8 address space for an Organization-Local SSM Scope for private SSM traffic inside of the enterprise network. This allows the 239/8 multicast boundary that is normally configured at the edge of the enterprise network to prevent private SSM traffic from leaving the network. Additionally, assume that the network administrator also wants to establish other SSM Scopes that are smaller than the Organization-Local SSM Scope just described. These smaller SSM Scopes would restrict SSM traffic from crossing an SSM Scope boundary.

The recommended method is to allocate the 239.232.0.0/16 address space from the Site-Local Expansion block for private SSM multicast. Continuing with our previous example, this results in the Private SSM block being allocated as shown in the address map in Figure 17.

Figure 17
Private SSM Address Range

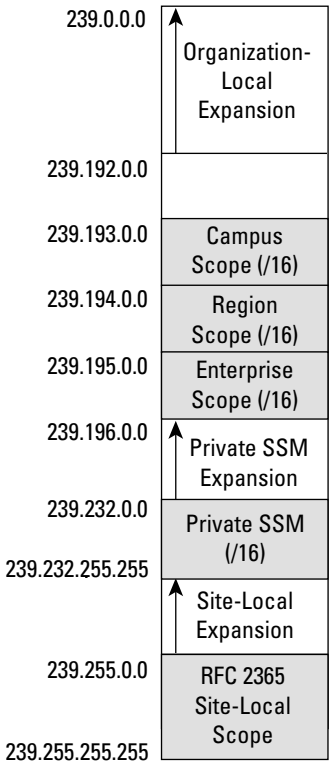

The Private SSM block can be subdivided into smaller blocks for SSM Scopes that correspond to our Region and Campus Scopes as shown in Figure 18. Notice that the Enterprise SSM Scope is allocated from the top of the Private SSM block at 239.255.0.0/24, and the Region and Campus Scopes are also /24 ranges allocated from the Private SSM block immediately below the Enterprise SSM range. The approach used in this example has the advantage that the configuration for the SSM address range in each router in the network remains constant (ranges 232/8 and 239.232/16) even if new SSM Scopes are defined later.

The size of these SSM Scope group ranges is only a suggestion, and individual network administrators may use different group range sizes as desired. The choice of a /24 allocation for each of the SSM Scopes was made to make the ranges fall on an octet boundary to keep the address masks simple. Smaller mask sizes can be used (for instance, /20 masks) if additional addresses are desired in each of the SSM Scopes. However, one of the key advantages of SSM is that multiple sources can use the same group address without causing their traffic to be merged. (That is because SSM does not use a Shared Tree and only Shortest-Path Trees are used.) This means that the number of addresses in an SSM range is less critical than for other PIM multicast modes such as classic PIM-SM (also called Any Source Multicast) or Bidir PIM.

Figure 18
Subdividing the Private SSM Block

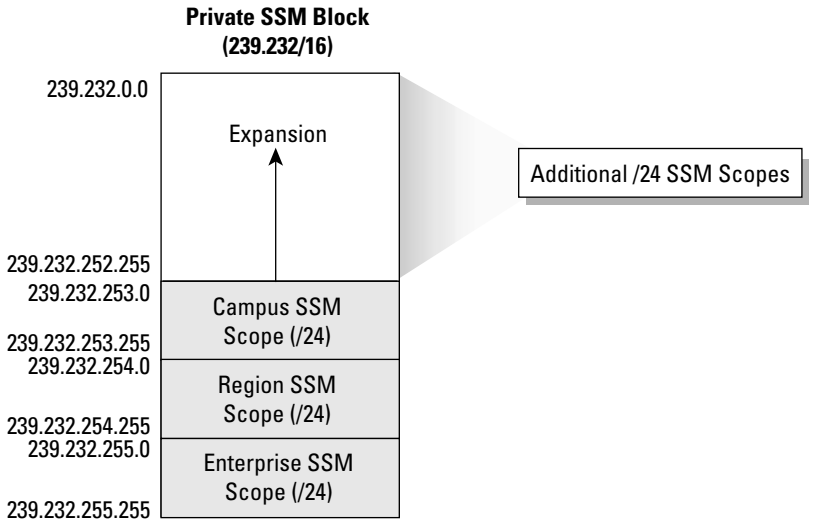

Finally, this example has allocated the Enterprise, Region, and Campus SSM scope ranges in order of their decreasing scope size. This is not mandatory and these SSM Scope ranges may be allocated in any order.

Final Address Allocation Map

In this section we have outlined a recommended strategy for allocating the various group ranges for Administrative Scoped Zones. Our Administrative Scoped Zone example given in this section defined three geographical scopes, Enterprise, Region, and Campus, and then added allocations for along with the group ranges for SSM and Bidir operation. Table 10 is a complete listing of the Administratively Scoped group address range and its allocations as defined in this example.

Table 10 Example Private Group Address Allocation Map

Range	Use	Subrange	Use
239.0.0.0–239.191.255.255	Unused	—	—
239.193.0.–239.193.255.255	Campus Scope	239.193.0.0–239.193.253.255	Campus Scope
		239.193.254.0–239.193.254.255	Campus Bidir Scope
		239.193.255.0–239.193.255.255	Campus Scope Relative
239.194.0.0-239.194.255.255	Region Scope	239.194.0.0–239.194.253.255	Region Scope
		239.194.254.0–239.194.254.255	Region Bidir Scope
		239.194.255.0–239.194.255.255	Region Scope Relative
239.195.0.0-239.195.255.255	Enterprise Scope	239.195.0.0–239.195.253.255	Enterprise Scope
		239.195.254.0–239.195.254.255	Enterprise Bidir Scope
		239.195.255.0–239.195.255.255	Enterprise Scope Relative
239.196.0.0–239.231.255.255	Unused	—	—

Table 10 Example Private Group Address Allocation Map (Continued)

Range	Use	Subrange	Use
239.232.0.0–239.232.252.255	SSM Expansion	—	—
239.232.253.0-239.232.255.255	SSM	239.232.253.0–239.232.223.255	Campus SSM
		239.232.254.0–239.232.239.255	Region SSM
		239.232.255.0–239.232.255.255	Enterprise SSM
239.233.0.0–239.254.255.255	Unused	—	—
239.255.255.0-239.255.255.255	Site-Local Scope	239.255.0.0–239.255.254.255	Site-Local Scope
		239.255.255.0–239.255.255.255	Site-Local Scope Relative

DEPLOYING ADMINISTRATIVE SCOPING

Enterprise networks have evolved from simple campus networks to global networks. Many large enterprise networks are beginning to look like Service Provider networks and the allocation of multicast addresses is becoming more challenging. Any deployment of IP multicast in a large-scale enterprise network needs to consider:

- How to control IP multicast traffic propagation within and between domains
- How to control dynamic distribution of Rendezvous Point information within and between domains

Because both considerations also may have an impact on addressing, they are discussed in the following sections.

Controlling IP Multicast Traffic

Generally IP Multicast traffic control within and between domains is achieved by using the **ip multicast boundary** command. This command filters IP Multicast data and controls traffic inbound and outbound that matches the optional ACL assigned to the interface command. Instead, it acts on PIM messages to stop the multicast forwarding table being populated in such a way that it violates the configured ACL. This has no effect on the multicast forwarding performance of the device.

TTL scoping of application traffic and the use of the **ip multicast ttl-threshold** interface command to create boarders was an old method of scoping multicast traffic and is not recommended for the following reasons:

- The use of the command is platform dependent and generally has an adverse effect on the IP Multicast forwarding performance on the platforms that perform hardware switching
- Difficult to plan and maintain as sources are generally not located at equal distances (in hops) from the boarders; redundant paths also need to be considered during failure scenarios as this will change the hop count from source to boundary.
- Relies on the configuration of the TTL on the application level that is outside the control of the Network Operations team

Controlling Rendezvous Point Information Distribution

Multiple options exist for configuring Rendezvous Point's and how Rendezvous Point information can be distributed within and between administrative IP multicast domains. These options include but are not limited to:

- Static RPs
- Anycast RPs
- BSR
- Auto-RP

This document is not intended to discuss each of the different options in detail. However, the following sections provide a basic overview of the various options and make some basic recommendations as to how administrative scopes can be controlled when Rendezvous Point information is dynamically distributed.

Static RPs

This method requires the network administrator to manually configure the IP address of the Rendezvous Point associated with each particular group range on each router in the network. The disadvantage of this approach is that the network administrator must make any group range to address changes on every router in the network. The second disadvantage is that the Static RP method alone does not provide a way to failover to a backup Rendezvous Point if the configured Rendezvous Point fails. Some administrators prefer having the address of the Rendezvous Point "locked-down" and not be subject to possible DoS attacks by devices attempting to send bogus Auto-RP or BSR messages. One advantage of the Static RP method is that it prevents the network from experiencing "Dense Mode Fallback" when older versions of Cisco IOS are in use. This is because the Rendezvous Point information is statically defined in each router and cannot timeout as is the case in Auto-RP or BSR due to Rendezvous Point failure or problems with the Auto-RP or BSR mechanisms due to network outage or congestion.

The Static RP approach can be used with Administratively Scoped Zones by configuring each router within a particular scoped zone with the IP address of the router that is serving as the Rendezvous Point within that zone.

BSR

BSR is a protocol for dynamic Rendezvous Point distribution and Rendezvous Point redundancy that originally was defined in the PIMv2 specification. This method uses a hop-by-hop flooding mechanism to distribute Bootstrap messages throughout the network. Using this hop-by-hop flooding of BSR messages, Candidate BSRs (C-BSRs) participate in a BSR election based on their configured BSR priority. The highest priority C-BSR is elected and begins functioning as the active BSR for the network. (This election mechanism is similar to the Root Bridge election mechanism used in the Spanning-Tree Protocol, which also uses hop-by-hop flooding of bridge protocol data unit messages to perform the election.)

After the BSR election has occurred, Candidate-RPs send their Candidate-RP announcement messages to the BSR using unicast. The BSR, in turn, collects all Candidate-RP announcements from all Candidate-RPs in the network into a collection called the RP-Set. The RP-Set is transmitted in its periodic BSR messages. Because these messages are flooded hop-by-hop through the network, all PIM routers receive the RP-Set and use a well-known hashing algorithm to select the active Rendezvous Point.

Many network administrators would like to use BSR as the dynamic Rendezvous Point mechanism in their network. However, the current implementation of BSR does not provide the ability to support Administratively Scoped Zones[1]. The first key reason that the current implementation of BSR cannot be used for dynamic scoping is that Candidate-RP messages are *unicast* to the BSR without any deference to scope boundaries. The second is that BSR messages are flooded through out the entire PIM domain. Either one of these two functions can result in the leakage of Rendezvous Point information from one zone into another. Future specifications and implementations of BSR may provide extensions to support Administratively Scoped Zones.

Auto-RP

Auto-RP is a Cisco proprietary protocol for dynamic Rendezvous Point information distribution and Rendezvous Point redundancy. This method makes use of the two the IANA assigned multicast groups, 224.0.1.39 (Cisco Announce) and 224.0.1.40 (Cisco Discovery) to perform dynamic Rendezvous Point election and information distribution within the PIM domain. In Auto-RP, one or more Cisco routers are configured by the network administrator as Candidate-RPs for a particular group range using the Cisco IOS **ip pim send-rp-announce** global command. (Generally two or more Candidate-RPs are used to provide Rendezvous Point redundancy if the active Rendezvous Point fails.) In addition, one or more Mapping Agents are configured by the network administrator using the Cisco IOS **ip pim send-rp-discovery** global command. (Generally two or more Mapping Agents are used to provide Rendezvous Point redundancy if the active Rendezvous Point fails.) The Candidate-RPs *multicast* their Candidate-RP announcements to the 224.0.1.39 Cisco Announce group. The Mapping Agents serve as the Rendezvous Point election agency by listening to the Cisco Announce group to receive Candidate-RP Announcements, selecting the highest IP address Candidate-RP for a group range and then *multicast* the election results to the 224.0.1.40 Cisco Discovery group. All Cisco routers in the network automatically join the 224.0.1.40 group so that they can learn the IP address of the elected Rendezvous Point for each group range.

One of the problems with Auto-RP when combined with Administrative Scoping is that Auto-RP information from one scope can leak into another scope. Figure 19 is an example of Auto-RP information for Campus A leaking into the Headquarters Campus and Campus B. This can occur as a result of either Campus A Candidate-RP Announcements or Mapping Agent Discovery messages leaking across the Campus A boundary. This can cause routers in Campuses other than Campus A erroneously to elect a Candidate-RP in Campus A as their local Campus Rendezvous Point.

1. This was an oversight when the original version of the BSR specification was published in the PIMv2 draft. The IETF PIM working group is actively working on a new version of the BSR specification that will address this deficiency.

Figure 19
Auto-RP Information Leakage

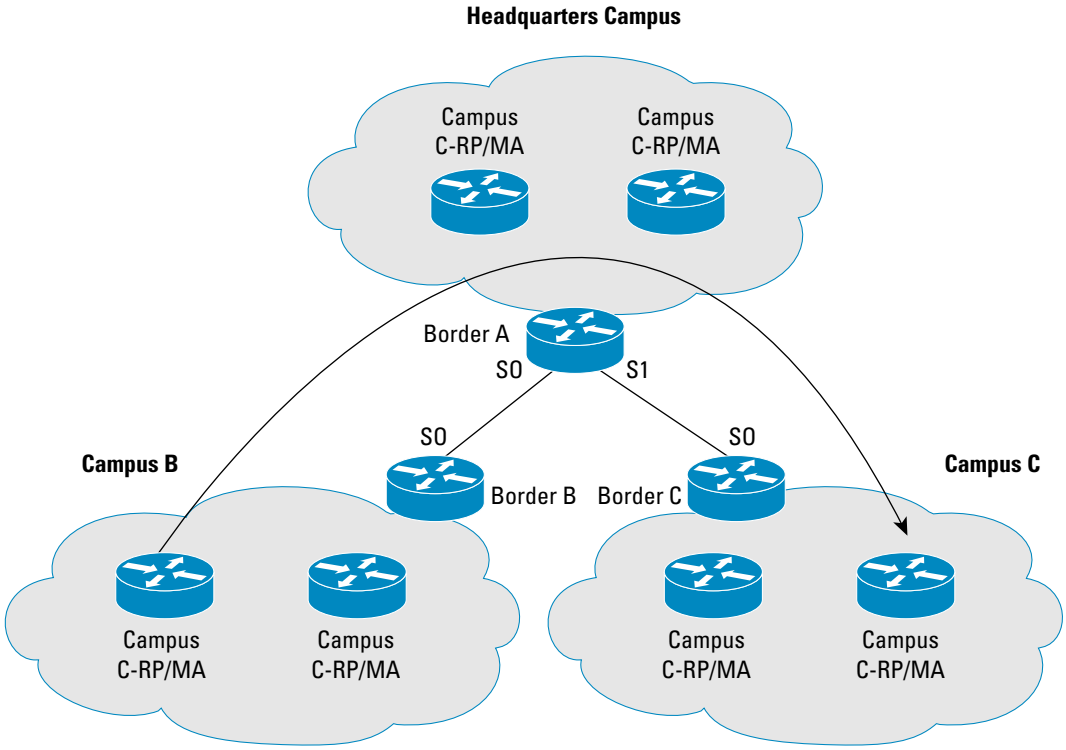

The original solution to this problem was to control the propagation of Auto-RP information by the use of specific TTL values on Auto-RP Announcement and Discovery messages for the scope. These messages would, in turn, be stopped from leaving the scope using **ip multicast ttl-threshold** commands at boundaries edges. This approach creates unnecessary administrative overhead and is difficult to plan and support. Since Cisco IOS Software Versions 12.2.(12), 12.2.(12)S and 12.1.(13)E, there is now a new **filter-autorp** option on the **ip multicast boundary** interface command, which is detailed below:

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

When the **filter-autorp** option is added, the multicast boundary will examine RP-Discovery and RP-Announcement messages and filter (remove) RP group-range announcements from these messages if they are denied by the boundary ACL. An RP group-range announcement in an Auto-RP message is permitted and passed by the boundary only if ALL addresses in the RP group-range are permitted by the multicast boundary ACL; otherwise, the whole RP group-range announcement is removed from the message before the message is passed on.

As an example, if an Auto-RP message contains a RP group-range of 239.100.0.0/16 and tries to pass a multicast boundary with a deny clause in the ACL, such as **deny 239.100.5.0/24**, then the RP group-range entry for 239.100.0.0/16 is removed from the Auto-RP packet regardless of the Auto-RP packet type (RP-Announce or RP-Discovery). The inverse of this situation is also true. For instance, if an Auto-RP message contains a RP group-range of 239.100.5.0/24 and tries to pass a multicast boundary with a deny clause in the ACL, such as **deny 239.100.0.0/16**, then the RP group-range entry for 239.100.5.0/24 is removed from the Auto-RP packet.

Because of this "all-or-nothing" Auto-RP group-range filtering mechanism, care should be taken to insure that the group-range ACLs in Auto-RP Candidate RP configurations do not unintentionally "intersect" or overlap group-ranges denied in the multicast boundary ACL. The best way to avoid this problem is to insure that the group-ranges used in multicast boundaries and Candidate RP configurations match exactly.

Anycast-RPs

Anycast-RP is an extension of the Static RP technique that also allows multiple Rendezvous Points for a group range to be deployed. This allows the network to continue to operate if a Rendezvous Point fails. The basic concept of Anycast-RP is shown in Figure 20. The idea is to configure two or more routers in the network to be the Rendezvous Point. Each of these Anycast-RP routers will be configured with the same Rendezvous Point address (in this case 10.1.1.1) on one of their Loopback interfaces. Each router also will advertise this address (the Rendezvous Point address) as a /32 host route. This will result in the other routers in the network using the closest Anycast-RP as their Rendezvous Point based on the unicast routing metrics. Normally, this would split the network into multiple PIM-SM domains that would not talk to each other. However, the Multicast Source Discovery Protocol (MSDP) is used to communicate active source information from one Anycast-RP to the other in Source Active (SA) messages. This allows active sources in one half of the network to be learned and joined by the Rendezvous Point in the other half of the network.

Figure 20
Anycast-RP

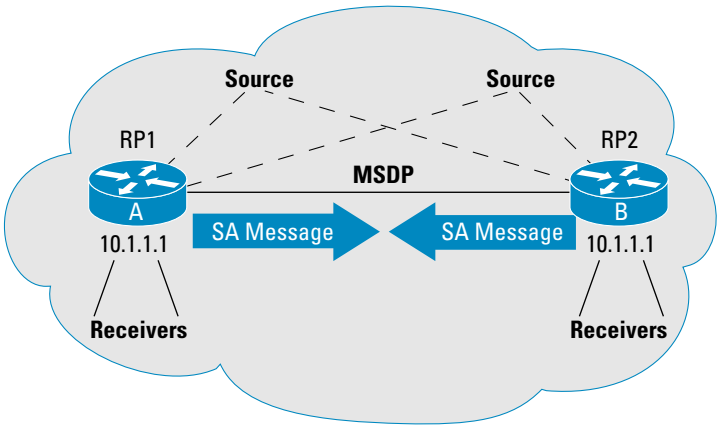

If an Anycast-RP were to fail, its host route would cease being advertised to the network and the unicast routing will reconverge on the remaining Anycast-RP. This will cause the routers in the network to rejoin and reregister receivers and sources to the remaining Anycast-RP to maintain multicast traffic flow. This process occurs in approximately the time that it takes unicast routing to converge which means that Anycast-RP has one of the fastest Rendezvous Point failover times of all of the Rendezvous Point configuration methods. This is why many network administrators prefer this method over Auto-RP and BSR.

In addition to providing rapid Rendezvous Point failover, the Anycast-RP method reduces the possibility of DoS attacks that can occur when rogue devices send bogus Auto-RP or BSR messages. Finally, because the Anycast-RP is based on a statically defined Rendezvous Point address, the network cannot "fallback" into Dense mode as it can when older versions of Cisco IOS are used with Auto-RP or BSR.

Administrative Scoping may be used with Anycast-RP, although the configuration requirements are more complex than Auto-RP and require the use of unique Anycast-RP addresses within physical scope. This is necessary to prevent routers in one scope from trying to use the Anycast-RP of an adjacent scope.

Figure 21 shows an example of Anycast-RPs being used for the Campus scope at three different Campus sites. It is important that routers in each Campus use a unique Anycast-RP address so that they do not accidentally attempt to connect to Anycast-RPs in one of the other two Campus sites.

Figure 21
Campus Scopes Using Anycast-RP

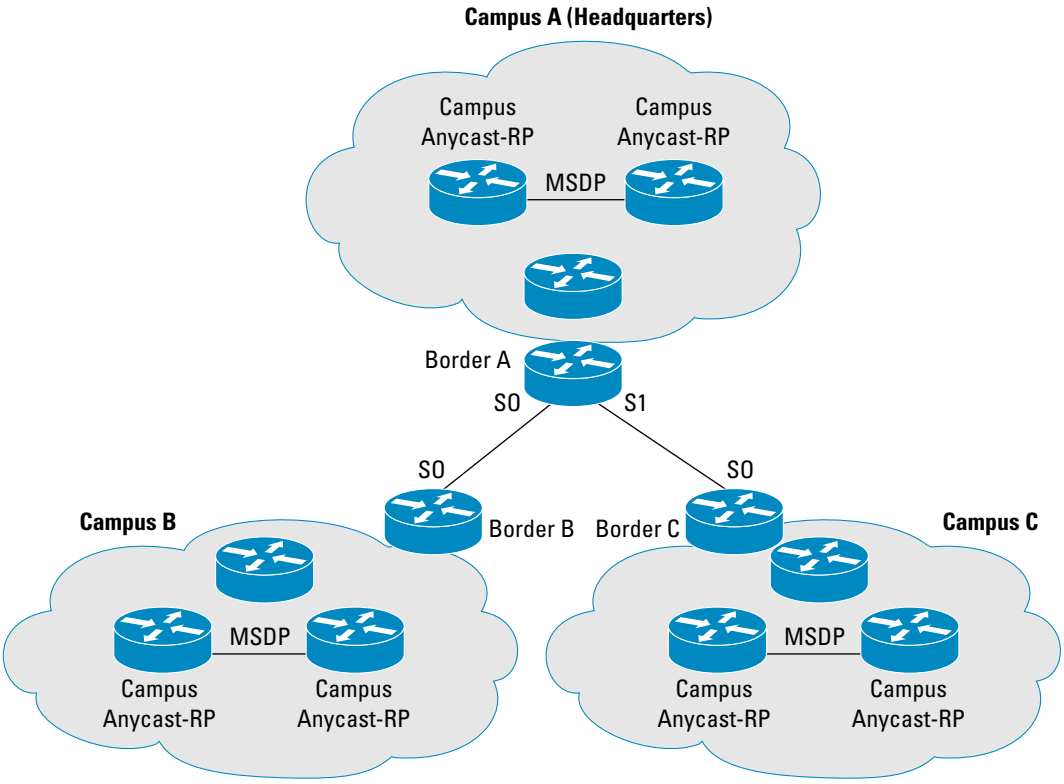

Example 2

Campus A Anycast-RP Configuration

```
!
! Campus A Left Anycast-RP
!
Interface loopback 0
 ip address 10.0.10.1 255.255.255.255
Interface loopback 1
 ip address 10.0.10.2 255.255.255.255
!
ip msdp peer 10.0.10.3 connect-source loopback 1
ip msdp originator-id loopback 1
!
! Campus A Right Anycast-RP
!
```

```
Interface loopback 0
 ip address 10.0.10.1 255.255.255.255
Interface loopback 1
 ip address 10.0.10.3 255.255.255.255
!
ip msdp peer 10.0.10.2 connect-source loopback 1
ip msdp originator-id loopback 1
!
! All Campus A Routers
!
ip pim rp-address 10.0.10.1
```

Example 3

Campus B Anycast-RP Configuration

```
!
! Campus B Left Anycast-RP
!
Interface loopback 0
 ip address 10.0.20.1 255.255.255.255
Interface loopback 1
 ip address 10.0.20.2 255.255.255.255
!
ip msdp peer 10.0.20.3 connect-source loopback 1
ip msdp originator-id loopback 1
!
! Campus B Right Anycast-RP
!
Interface loopback 0
 ip address 10.0.20.1 255.255.255.255
Interface loopback 1
 ip address 10.0.20.3 255.255.255.255
!
ip msdp peer 10.0.20.2 connect-source loopback 1
ip msdp originator-id loopback 1
!
! All Campus B Routers
!
ip pim rp-address 10.0.20.1
```

Example 4

Campus C Anycast-RP Configuration

```
!
! Campus C Left Anycast-RP
!
Interface loopback 0
 ip address 10.0.30.1 255.255.255.255
Interface loopback 1
 ip address 10.0.30.2 255.255.255.255
!
ip msdp peer 10.0.30.3 connect-source loopback 1
ip msdp originator-id loopback 1
!
! Campus C Right Anycast-RP
!
Interface loopback 0
 ip address 10.0.30.1 255.255.255.255
```

```
Interface loopback 1
 ip address 10.0.30.3 255.255.255.255
!
ip msdp peer 10.0.30.2 connect-source loopback 1
ip msdp originator-id loopback 1
!
! All Campus C Routers
!
ip pim rp-address 10.0.30.1
```

The details of Anycast-RP configuration for these three Campus scopes are shown in Example 2 through Example 4. It is important to note that when we use Anycast-RP with administrative scoping we cannot use a "cookie-cutter" configuration for Anycast-RP. Instead, we must insure that every router within the scope has the correct Anycast-RP address configured and that this address will vary depending on the router's location.

Combining Auto-RP and Anycast-RP

The use of Auto-RP and Anycast-RP techniques are not mutually exclusive. In fact, these techniques can be used together to reap the benefits of both. This can be true especially if we want to deploy Administrative Scoping in an Anycast-RP environment to reap the benefits of rapid Rendezvous Point failover. As discussed in the previous section, this results in some rather tedious configuration tasks, because a unique Anycast-RP address must be assigned and configured inside of every scope. However, we can use Auto-RP to distribute the unique Anycast-RP addresses inside of each scope.

This is accomplished by configuring Auto-RP Mapping Agents inside each of the smallest scopes. In the example in the previous section, we would configure Mapping Agents inside of each Campus scope using the **ip pim send-rp-discovery** command. (Most likely we also would configure more than one Mapping Agent to provide redundancy.) Next, we would configure the Anycast-RP routers as Auto-RP Candidate RPs using the **ip pim send-rp-announce** command. Finally, we would configure the **ip multicast boundary <acl> filter-autorp** command at the edges of the scopes to insure that Auto-RP information does not leak outside of the scope. This will result in the Anycast-RP address being advertised to all routers inside the scope and will eliminate the need to uniquely configure every router in the scope. Using this technique allows one to easily define new Rendezvous Points and scopes without having to reconfigure every router in the network.

Example Deployment

Figure 22 shows a global enterprise network where geographical scoping would be useful to control and monitor traffic. In this example, a three-tiered scoped zone architecture is used, which is similar to the methodology outlined in the "Deploying Administrative Scoping" section as follows:

- **Site-Local Scope** is a reserved scope that allows us to comply with RFC 2365 for applications (such as MADCAP and others) that assume that this is the smallest possible scope.

  Note: The Site-Local zone should never be subdivided into smaller scopes because this violates RFC 2365. Although it is possible to assign the Site-Local Scope range to the smallest known geographical scope, it is better to keep this as an independent range.

- **Campus Scope** identifies applications local to a campus (for instance, buildings or a building complex). Generally, the physical topology that this scope covers depends on the network and business topology. For example, if several buildings are interconnected by high-speed links, they may be considered as a single Campus Scope.

- **Regional Scope** identifies applications for intraregion IP Multicast (for example, EMEA, North America, South America, Asia Pacific, etc.).
- **Enterprise Scope** identifies applications for IP Multicast communication to all locations within the organization.
- **Global Scope** identifies applications for IP Multicast communication to and from the global Internet that include all addresses in the range of 224.0.1.0 through 238.255.255.255.

Figure 22
Example Global Network

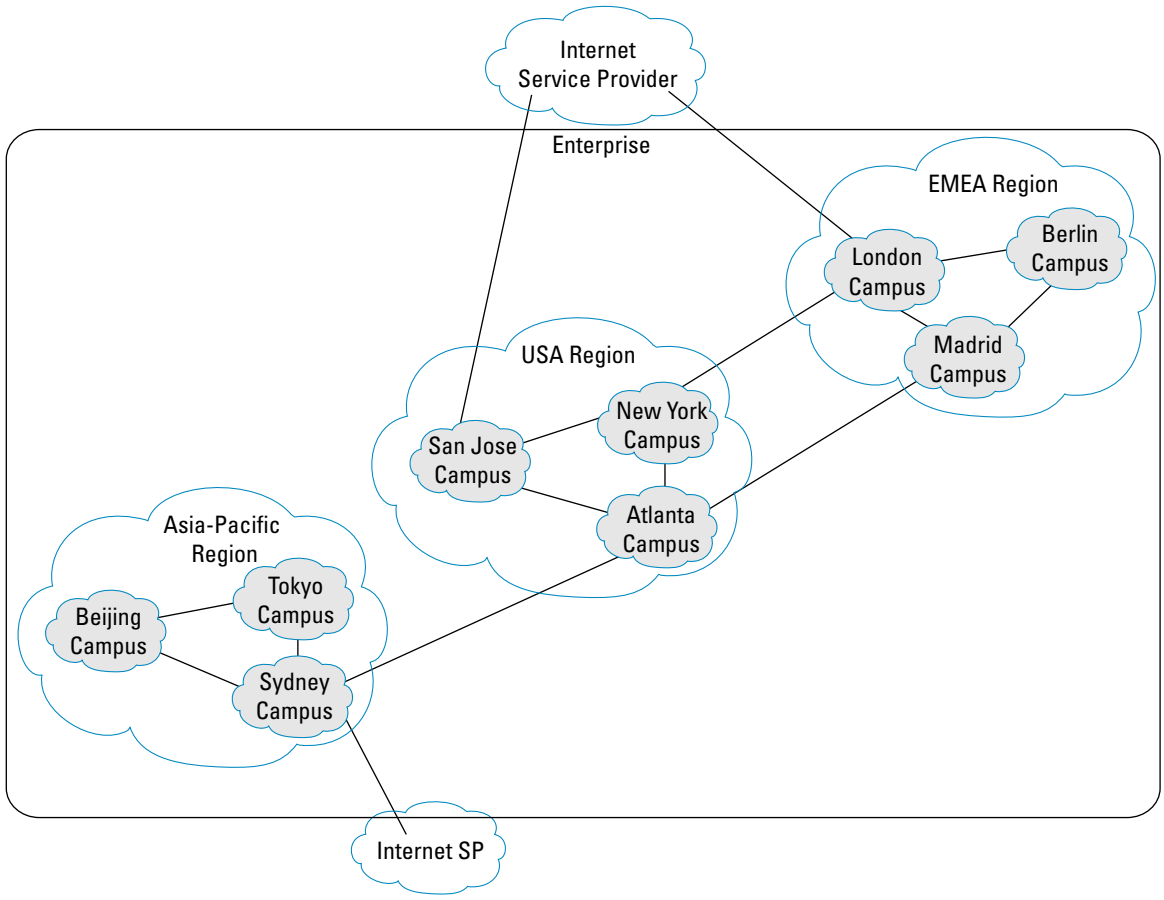

Configuring Scope Boundaries

Table 11 lists scoped zones within the enterprise along with their assigned address ranges used in our example network. (**Note:** The Global Scope addresses are not shown in this table but include addresses in the range of 224.0.1.0–238.255.255.255.)

Table 11 Example Scoped Zone Address Blocks

Byte 1	Byte2	Byte3								Byte4								Notes
239	193	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	Campus Scope
239	194	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	Region Scope
239	195	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	Enterprise Scope
239	232	253								*	*	*	*	*	*	*		SSM Campus Scope
239	232	254								*	*	*	*	*	*	*		SSM Region Scope
239	232	255								*	*	*	*	*	*	*		SSM Enterprise Scope
239	255	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	Site-Local Scope

The expansions shown in Table 11 take into consideration the ease of configuration of scope boundaries by providing contiguous masks for filtering. This is reflected by the simplicity of the corresponding boundary access list definitions shown in Table 12.

As shown in Table 12, using a clean octet boundary to define the scopes enhances the readability of these ACLs. The Enterprise ACL can be simplified to block all multicast in the 239/8 range from entering and leaving the enterprise network. This is considered good practice because the Enterprise Scope is the largest possible scope short of the global Internet Scope.

Table 12 Recommended Multicast Boundary Access Lists

Scoped Zone	Multicast Boundary Access List

Enterprise	`ip access-list standard Enterprise`
	`deny 239.195.0.0 0.0.255.255`
	`deny 239.232.255.0 0.0.0.255`
	`permit 224.0.0.0 15.255.255.255`
	`or`
	`ip access-list standard Enterprise`
	`deny 239.0.0.0 0.255.255.255`
	`permit 224.0.0.0 15.255.255.255`

Region	`ip access-list standard Region`
	`deny 239.194.0.0 0.0.255.255`
	`deny 239.232.254.0 0.0.0.255`
	`permit 224.0.0.0 15.255.255.255`

Campus	`ip access-list standard Campus`
	`deny 239.193.0.0 0.0.255.255`
	`deny 239.232.253.0 0.0.0.255`
	`permit 224.0.0.0 15.255.255.255`

Site-Local	`ip access-list standard Site-Local`
	`deny 239.255.0.0 0.0.255.255`
	`permit 224.0.0.0 15.255.255.255`

On initial inspection, one might be concerned that each ACL denies only the address range associated with the scope and permits all others. This seemingly "permissive" approach to ACL construction may appear backwards compared to typical ACL definition techniques. Such techniques would use the opposite approach, where only the desired traffic is explicitly permitted while all other traffic is implicitly denied, as shown in Table 13. *Although this alternative approach can be used, it is not recommended because these ACLs are less intuitive and more difficult to create and manage.*

For example, the ACLs shown in Table 13 are not clear as to what address range is being blocked at each boundary. Adding new scopes to the network requires revisiting ***ALL*** router configurations in the network to reconstruct new ACLs that accommodate the new scopes with implicit **permit** statements.

Finally, one should remember that the basic nested nature of the scopes from smaller scope size to larger scope size accomplishes the same result while allowing a simplified ACL construct. For example, Campus traffic cannot leave a Region Scope because traffic from hosts inside of Campus cannot cross the Campus boundary in the first place. Because the Campus boundaries are completely contained inside of the Region boundaries, then by definition, Campus traffic cannot escape the Region.

Table 13 Alternative Multicast Boundary Access List

Scoped Zone	Multicast Boundary Access-list

Enterprise	`ip access-list standard Enterprise`
	`permit 224.0.0.0 7.255.255.255`
	`permit 232.0.0.0 3.255.255.255`
	`permit 236.0.0.0 1.255.255.255`
	`permit 238.0.0.0 0.255.255.255`
	`deny 224.0.0.0 15.255.255.255`

Region	`ip access-list standard Region`
	`permit 239.195.0.0 0.0.255.255`
	`permit 239.232.255.0 0.0.0.255`
	`permit 224.0.0.0 7.255.255.255`
	`permit 232.0.0.0 3.255.255.255`
	`permit 236.0.0.0 1.255.255.255`
	`permit 238.0.0.0 0.255.255.255`
	`deny 224.0.0.0 15.255.255.255`

Campus	`ip access-list standard Campus`
	`permit 239.195.0.0 0.0.255.255`
	`permit 239.194.0.0 0.0.255.255`
	`permit 239.232.255.0 0.0.0.255`
	`permit 239.232.254.0 0.0.0.255`
	`permit 224.0.0.0 7.255.255.255`
	`permit 232.0.0.0 3.255.255.255`
	`permit 236.0.0.0 1.255.255.255`
	`permit 238.0.0.0 0.255.255.255`
	`deny 224.0.0.0 15.255.255.255`

Table 13 Alternative Multicast Boundary Access List (Continued)

Scoped Zone	Multicast Boundary Access-list

Site-Local	`ip access-list standard Site-Local`
	`permit 239.195.0.0 0.0.255.255`
	`permit 239.194.0.0 0.0.255.255`
	`permit 239.193.0.0 0.0.255.255`
	`permit 239.232.255.0 0.0.0.255`
	`permit 239.232.254.0 0.0.0.255`
	`permit 230.232.253.0 0.0.0.255`
	`permit 224.0.0.0 7.255.255.255`
	`permit 232.0.0.0 3.255.255.255`
	`permit 236.0.0.0 1.255.255.255`
	`permit 238.0.0.0 0.255.255.255`
	`deny 239.255.0.0 0.0.255.255`
	`deny 224.0.0.0 15.255.255.255`

Configuring Auto-RP Candidate-RP ACLs

Table 14 shows the Auto-RP Candidate-RP group-range ACLs that correspond to our example network scopes for both the Bidir and non-Bidir (ASM) group-ranges. The definitions of the Candidate-RP ACLs in Table 14 and the Boundary ACLs in Table 12 have been chosen carefully so that Auto-RP information will be propagated correctly within a scope but will not accidentally leak out of the scope.

Table 14 Auto-RP Candidate-RP Access Lists

Candidate-RP	Auto-RP Group Range Access List
Enterprise	`ip access-list standard Enterprise-RP` `permit 239.195.128.0 0.0.127.255`
Enterprise-Bidir	`ip access-list standard Enterprise-bidir-RP` `permit 239.195.0.0 0.0.127.255`
Region	`ip access-list standard Region-RP` `permit 239.194.128.0 0.0.127.255`
Region-Bidir	`ip access-list standard Region-bidir-RP` `permit 239.194.0.0 0.0.127.255`
Campus	`ip access-list standard Campus-RP` `permit 239.193.128.0 0.0.127.255`
Campus-Bidir	`ip access-list standard Campus-bidir-RP` `permit 239.193.0.0 0.0.127.255`
Site-Local	`ip access-list standard Site-Local-RP` `permit 239.255.0.0 0.0.255.255`

REAL-WORLD ISSUES

The examples presented in the previous sections are based on ideal situations. In practice, frequently this is not the case and network administrators must deal with special case scenarios over which they have little or no control. This section describes some of the more common issues and how they impact multicast address allocation and scoping.

The one common solution to most of these issues is related to the ability to change the multicast addresses used by the applications. This should be a primary consideration when accepting applications onto the production network. Network and application administrators should understand the importance of this requirement to avoid the scenarios discussed in the following sections.

Bandwidth Scoping

In addition to scoping on an application or geographical basis, sometimes it is desirable to provide scoping based on a bandwidth or data rate basis. For example, consider the case where a real-time video broadcast has multiple feeds of the same content, although at two different data rates as shown in Figure 23. In this case, remote Campus A is connected to the Headquarters site using a DS-3 (45 Mbit) WAN link while remote Campus B is connected using only a DS-1 (1.5 Mbit) link. Given a Campus, Region, or Enterprise Scoping plan, the normal situation would be to assign the 2 MB video stream to the Campus Scope range and the 128 KB video stream to the Region range. However, because remote Campus A has sufficient bandwidth to receive the 2 MB stream, it is desirable to allow this stream to cross the Headquarters Campus boundary and flow to Campus A but block it from flowing to Campus B as shown in Figure 23.

The quick and convenient way to accomplish this is to modify the Campus **ip multicast boundary** ACLs at both ends of the DS-3 link to **permit** this one particular Campus multicast stream to cross the normal Campus boundaries. This approach creates several problems. First, our carefully crafted Campus boundary plan is becoming less clear. We have effectively "poked a hole" in the Campus boundaries between the Headquarters Campus and Campus A. If we continue to make exceptions like this for multiple applications, then our carefully crafted scoping plan will dissolve into chaos.

Second, this approach is useful only where we have a One-to-Many multicast application (as in this example) and a Hub-and-Spoke network topology (again as shown in this example). Consider what would happen if the application falls into the Many-to-Many category (such as if RTCP multicast feedback were in use by clients receiving the video stream). By just "poking a hole" in the Campus boundaries for this multicast group, we have bifurcated the PIM domain for this multicast group and we now have two active Rendezvous Points for the group; one at the Headquarters Campus and one in Campus A. Clients in Campus A would have their RTCP multicasts "Registered" to the Campus Rendezvous Point in Campus A while clients in the Headquarters Campus would be registered to the Rendezvous Point in the Headquarters Campus. A worse situation is if Auto-RP is in use it is possible for the Candidate-RP information for the Campus Rendezvous Points to begin to leak between the Headquarters Campus and Campus A, possibly causing problems in normal Campus Rendezvous Point election. *In short, this method can be made to work only if it is used on a limited basis and care is used to insure that the application characteristics fit this model and that no changes are ever made.*

Another approach would be to move the 2 MB stream out of the Campus scope range to the Region Scope range so that the stream can flow to all locations within the Region. However, since Campus B does not have sufficient bandwidth to accommodate the 2 MB stream, it must be blocked from flowing over the T1 link. Again, the quick and convenient fix would be to modify the Region **ip multicast boundary** ACL at the Headquarters end of the T1 link to **deny** this one particular Region multicast stream from crossing the T1 link even though Campus B still is

inside of the normal Region boundary. While this "hack" certainly appears on the surface to accomplish our goal, it has many of the same problems as the previous approach in that our carefully crafted scoping plan quickly can be reduced to chaos if this approach is used to any extent.

Figure 23
Scoping by Bandwidth and Data Rate

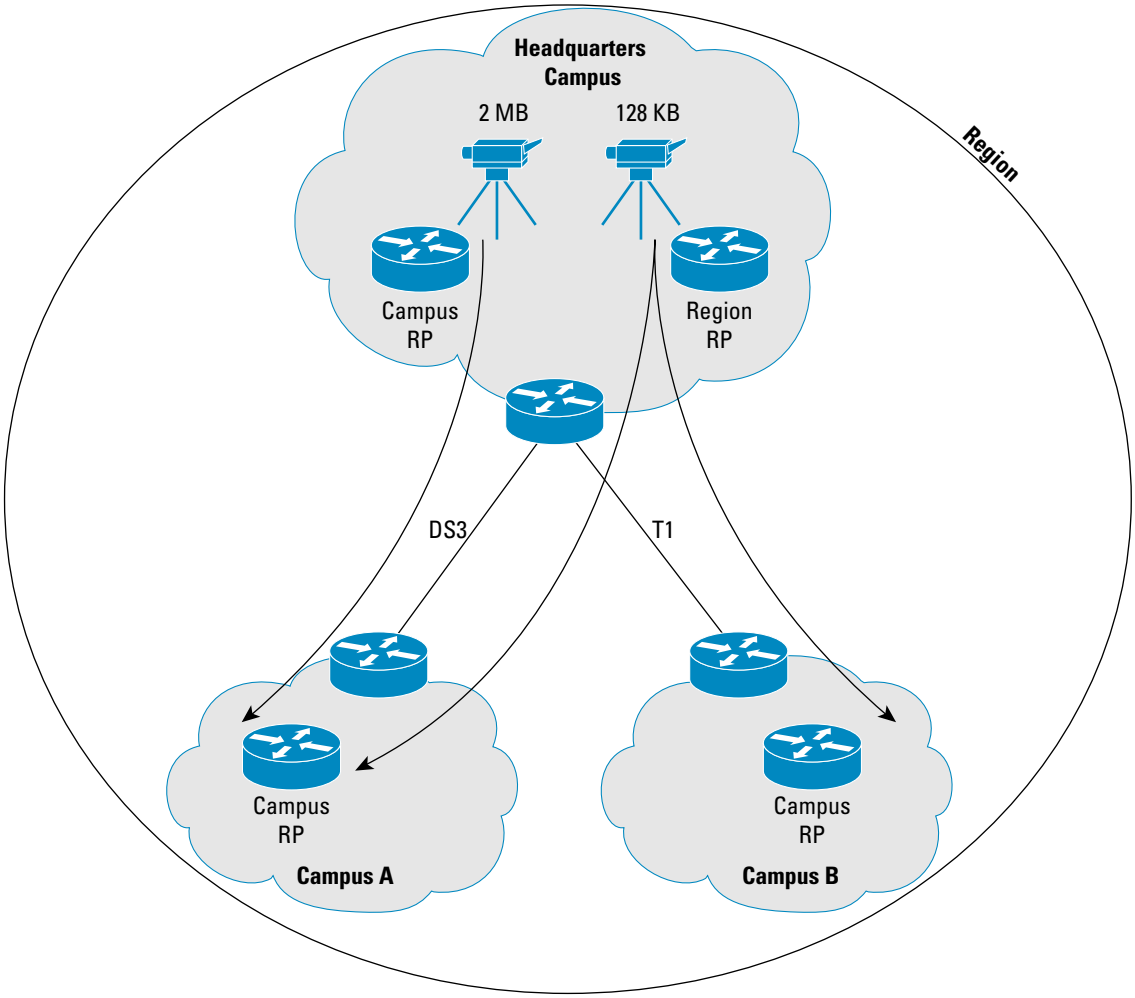

For example, this approach effectively has disabled the use of this address for multicast within Campus B. This is because the Rendezvous Point for the group is in the Headquarters Campus and any PIM Joins traveling from Campus B to the Rendezvous Point in the Headquarters Campus will be blocked by the **deny** that we have added at the Headquarters end of the T1 link. (Remember, the **ip multicast boundary** command not only blocks data, it also blocks PIM control traffic for any denied group.) What is worse, if Auto-RP is in use, the normal **filter-autorp** keyword that we would normally use on the **ip multicast boundary** command to prevent Candidate-RP information from leaking out of the scope would result in loss of Region Rendezvous Point information reaching Campus B. *This would cause the Region group range in Campus B to fallback into Dense mode.* Cisco Although this limitation can be overcome by using static Rendezvous Point definitions, it adds to the increase in chaos and overall management burden as we apply this approach to more and more groups.

Instead of taking this approach (which quickly can lead to total network management chaos), it is better to define a completely new scope that is bigger than the Campus Scope and yet smaller than the Region Scope. This could be accomplished easily by defining a new "Subregion" Scope and configuring the appropriate Subregion **ip multicast boundary** ACLs and adding the Subregion group ACL to the Rendezvous Point definitions as shown in Figure 24.

The group range for the new Subregion scope would be allocated from the Organization-Local range below the Campus Scope range as shown in Figure 25. This range can be allocated from any unused space in the Organization-Local range or its expansion range. It is not necessary to readdress the ranges to have the Subregion range fall between the Region and Campus address ranges.

Figure 24
Defining a Subregion Scope

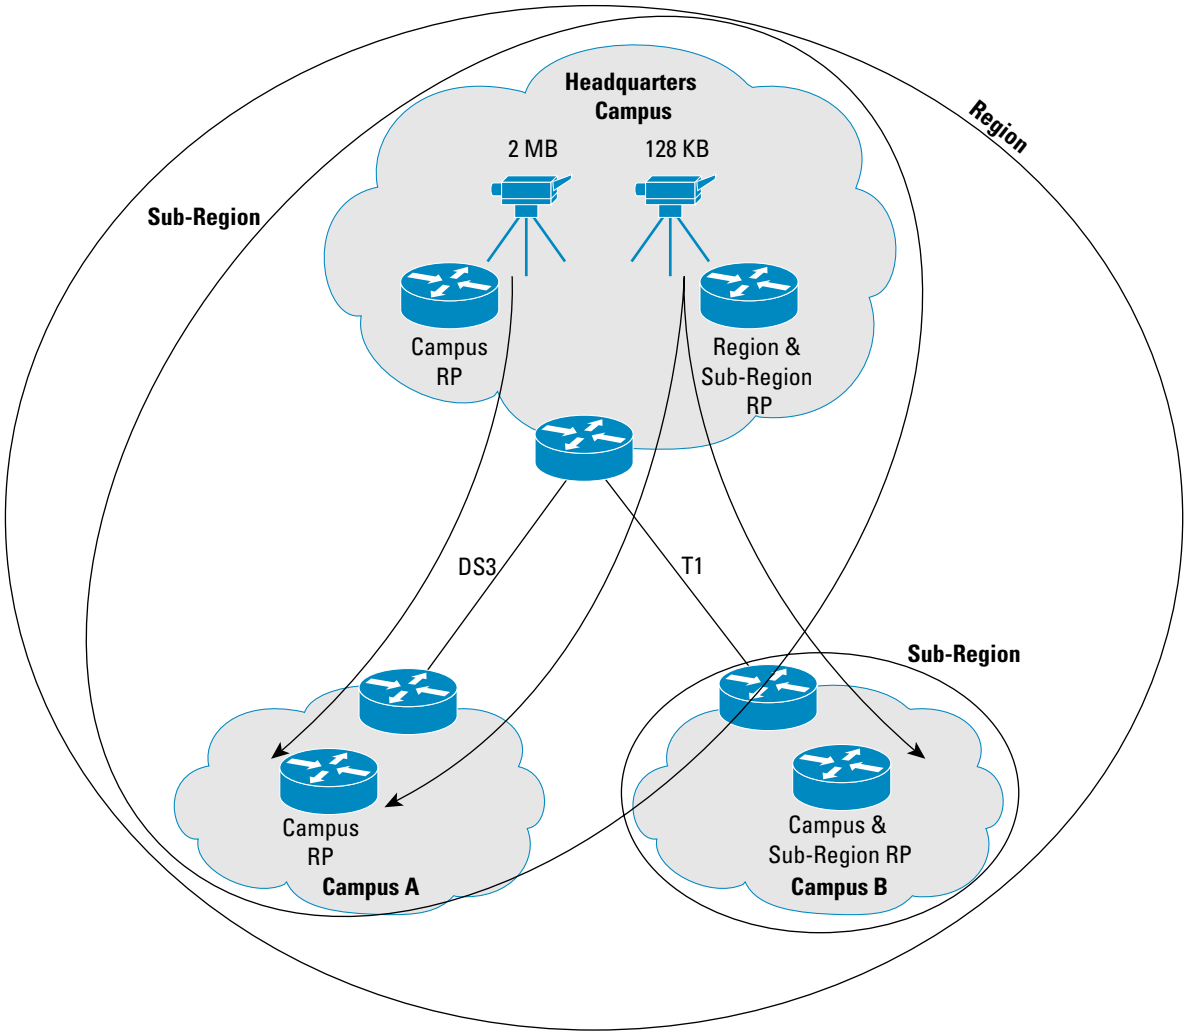

The problems described in the previous two methods are not evident in this approach. However, it is necessary to consult the global entity responsible for administrating multicast address space within the Enterprise prior to deploying this approach because it may affect other Region administrators. It is preferable for the administrator responsible for this Region's address allocation to use some portion of the Region address space to accomplish the same thing without the need for global approval.

Figure 25
Allocating Addresses for the Subregion Scope

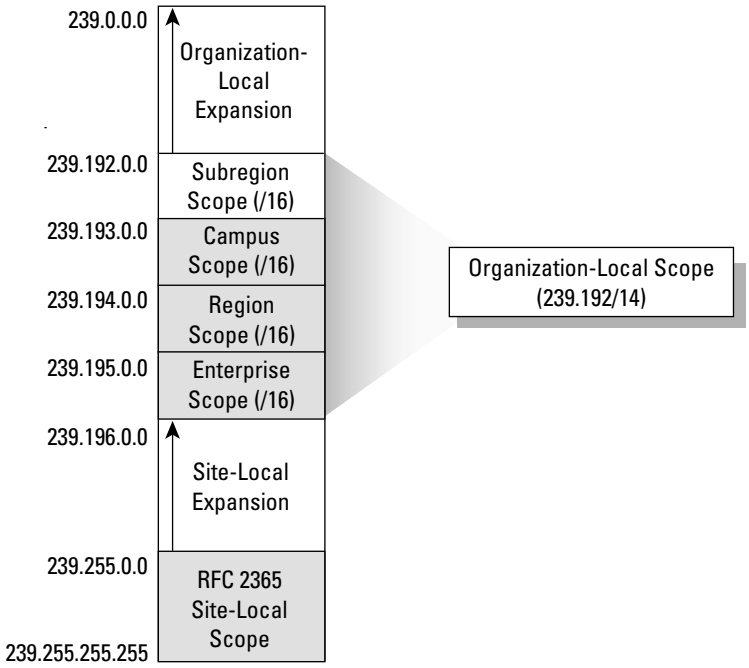

Figure 26 shows an example of how Region administrators can subdivide the Region space in use inside of their Region into other Subregion scopes based on data rate. The advantage is that Region administrators can allocate this space as they desire because the normal Region boundaries will isolate multicast use within each Region, thereby giving each Region administrator a large degree of autonomy in Region multicast address allocation.

In the example shown in Figure 26, the Region administrator has chosen to subdivide the Region address space into four data rate scopes to accommodate Low, Medium, High, and Unlimited data rate applications in both the Bidir and classic PIM-SM (ASM) operations.

Figure 26
Extending the Region Scope by Data Rate

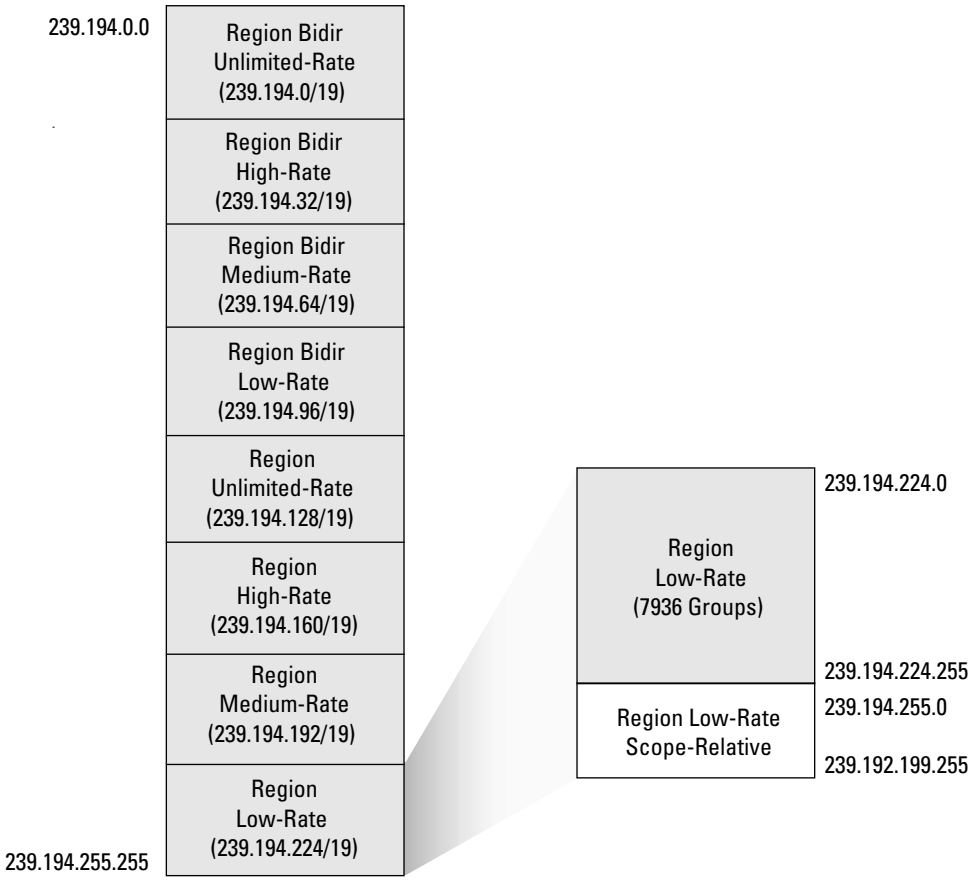

This four-rate by two-mode approach provides a high degree of flexibility for scoping multicast traffic within the Region, although at the price of more Rendezvous Point and multicast boundary configuration within the Region. However, once deployed these configurations generally are static and do not need to be modified frequently.

Hard-Coded Multicast Address Applications

Most of the common issues affecting address allocation are caused by IP multicast applications that use hard-coded IP Multicast addresses. The severity of the problem increases with globalization of those applications. In many cases if such problems arises, especially if they are applications managed by third parties, the only option for administrators is to make corresponding network design changes (multicast boundary ACLs, Rendezvous Points configurations in case of PIM SM, etc.) to accommodate the application. This makes administration and troubleshooting of the network complex. Another variation of this problem is unwillingness of the Application support engineers to change IP Multicast group addresses of existing applications to the addresses from the new allocated addressing scheme. This condition happens when a customer had limited deployment of IP Multicast and moves towards large scale IP Multicast and redesigns IP Multicast addressing scheme for this purpose.

Overlapping Addresses

The two most common cases of address overlapping are:

- Connection to third-party-managed applications
- Merging of organizations

In the first case, a third-party application provider may be using IP Multicast group address allocated from Organizational Local Scope (239.0.0.0/8) for its IP Multicast data stream. This can cause a problem with overlapping addresses with a customer's private addressing scheme. The ideal solution will be to request the third-party provider to change the multicast address in use. If this is not possible, an alternative network configuration will need to be found.

The second case, two merging organizations, is becoming more common with the wide scale adoption of IP Multicast. The long-term solution will be to readdress applications that require global connectivity. This shows the importance of an addressing scheme that is designed with expansion in mind.

Application Scope Exceptions

Well-defined IP Multicast scopes help to simplify network configurations, administration, and troubleshooting of IP Multicast. Any exceptions to these scopes begin to erode these advantages and introduce additional complexity to the network. A common issue relates to existing applications that require their original scope to be extended. Often this is only for a single user or department that relocates outside the originally intended scope of the application. The correct way to manage this would be to change the multicast address range used by the application to cover the now extended scope. If this cannot be done, it will require the current scope to be split and network configuration changes to be made to Rendezvous Point definitions and boundary ACLs throughout the network. This introduces additional complexity to the network and breaks the hierarchical structure that has been defined.

Although defining unique address scopes may give some "flexibility" to extend boundary scopes at short notice, this should be strongly discouraged in favor of application readdressing where possible.

Congruent Scope Boundaries

Congruent Scope Boundaries can occur when the interface on a router corresponds to the boundary of two or more scopes. For example, in Figure 27 interface **Serial0/1** on the border router in the Headquarters Campus is a boundary interface for both the Campus and the Region scopes.

Figure 27
Congruent Campus and Region Boundary

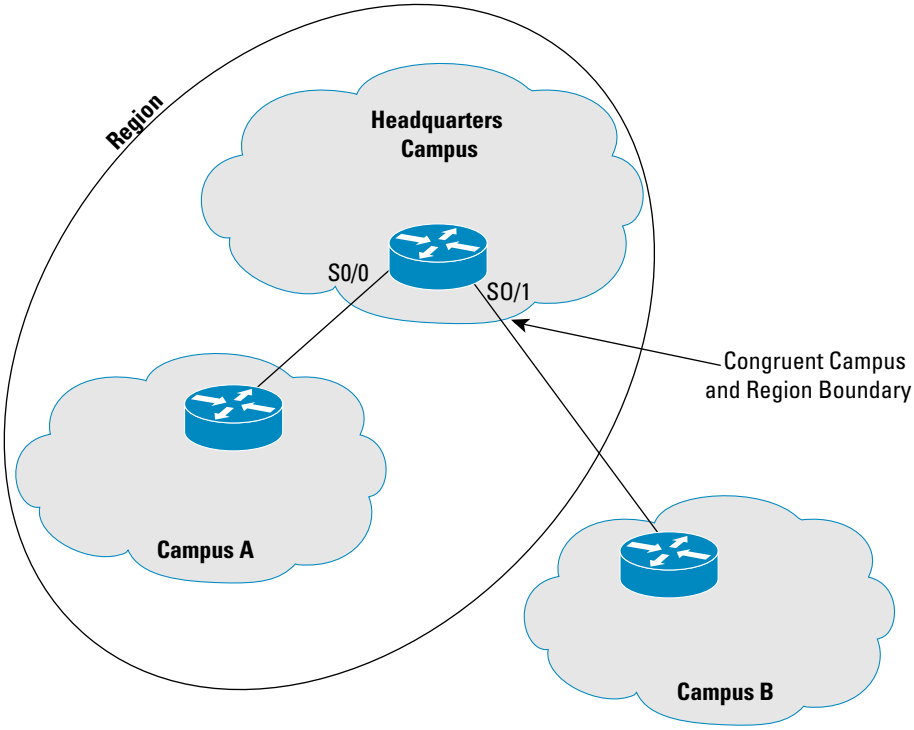

As a result, the multicast boundary ACLs of both the Campus and Region Scopes must be combined as shown in Example 5.

Example 5

Congruent Campus-Region Boundary ACL

```
Interface Serial0/1
. . .
ip multicast boundary Campus-Region filter-autorp
ip access-list standard Campus-Region
deny 239.194.0.0 0.0.255.255
deny 239.193.0.0 0.0.255.255
deny 239.232.254.0 0.0.0.255
deny 239.232.253.0 0.0.0.255
permit 224.0.0.0 15.255.255.255
```

Multiscope Rendezvous Points

Most of the examples shown in the previous sections depict separate Rendezvous Points (or Candidate-RPs in the case of Auto-RP) for each scope range. In certain cases, these Rendezvous Points can be combined when desired. For example, in the Headquarters Campus shown in Figure 27, the same Rendezvous Point can be used as the Headquarters Campus Rendezvous Point as well as the Region Rendezvous Point. Assuming that Auto-RP is in use, this could be configured as shown in Example 6.

Example 6

Combined Campus-Region Candidate-RP

```
ip pim send-rp-announce loopback0 scope 64 group-list Campus-Region-RP
ip access-list standard Campus-Region-RP
permit 239.194.0.0 0.0.255.255
permit 239.193.0.0 0.0.255.255
```

SUMMARY

This document does not attempt to provide a rigid addressing policy for each organization. It is intended to provide an overview of many factors that should be considered when designing a global addressing policy. Each individual organization needs to consider these with respect to their network, processes, and future plans. However, this document does provide a basic model from which network administrators may develop their own address and scoping plans to meet their individual network needs.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • V enezuela • Vietnam • Zimbabwe

204044_ETMG_AE_08.04