Secure IP Multicast Deployments

Contents

Introduction

This document describes general guidance on best practices to secure an IP multicast network infrastructure.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- IP multicast

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

This document covers some basic concepts, terminology, and discusses the listed topics:

- Mechanisms to secure a specific platform and the network in general.
- Any Source Multicast (ASM) and Source Specific Multicast (SSM) models.
- Multicast Virtual Private Network (MVPN) security.
- Group Encrypted Transport (GET) Virtual Private Network (VPN) architecture which provides confidentiality and integrity for multicast data plane or control plane traffic.

Terminology

In IP multicast there are two classic service models:

1. Any Source Multicast (ASM)

2. Source Specific Multicast (SSM)

In ASM, the receiver joins a group G via an Internet Group Membership Protocol (IGMP) or Multicast Listener Discovery (MLD) membership report to indicate the group. This report requests traffic sent by any source to the group G, and hence the name “any source.” By contrast, in SSM, the receiver joins a specific channel defined by a source S, which sends to a group G. Each of these service models is described in detail below.

Any Source Multicast

The ASM model is characterised by two classes of protocol: “dense mode flood-and-prune” and “sparse mode explicit join”:

i) Dense Mode Flood-and-Prune Protocols (DVMRP / MOSPF / PIM-DM)

In dense mode protocols, all routers in the network are aware of all trees, their sources and receivers. Protocols such as Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM) dense mode flood “active source” information across the whole network and build trees via the creation of “Prune State” in parts of the topology where traffic for a specific tree is unwanted. They are also called flood-and-prune protocols. In Multicast Open Shortest Path First (MOSPF), information about receivers is flooded throughout the network to support the build out of trees.

Dense mode protocols are undesirable because every tree built in some part of the network can always cause resource utilization (with convergence impact) on all routers in the network (or within the administrative scope, if configured). These protocols are not discussed further in the remainder of this article.

ii) Sparse Mode Explicit Join Protocols (PIM-SM/PIM-BiDir)

With sparse mode explicit join protocols, devices do not create a group-specific state in the network unless a receiver has sent an explicit IGMP/MLD membership report (or “join”) for a group. This variant of ASM is known to scale well and is the multicast paradigm of focus.

This is the basis for PIM-Sparse Mode, which most multicast deployments have used to this point. This is also the basis for Bidirectional PIM (PIM-BiDir), which is increasingly deployed for MANY (sources) to MANY (receivers) applications.

These protocols are called sparse mode because they efficiently support IP multicast delivery trees with a “sparse” receiver population and create a control plane state only on routers in the path between sources and receivers, and in PIM-SM/BiDir, the Rendezvous Point (RP). They never create state in other parts of the network. State in a router is only built explicitly when it receives a join from a downstream router or receiver, hence the name “explicit join protocols”.

Both PIM-SM and PIM-BiDir employ “SHARED TREES”, which allow traffic from any source to be forwarded to a receiver. The multicast state on a shared tree is referred to as (*,G) state, where the * is a wild card for ANY SOURCE. Additionally, PIM-SM supports the creation of state that relates to traffic from a specific source. These are known as SOURCE TREES, and the associated state is referred to as (S,G) state.

Source-Specific Multicast

SSM is the model used when the receiver (or some proxy) sends (S,G) “joins” to indicate that it wants to receive traffic sent by source S to group G. This is possible with IGMPv3/MLDv2 “INCLUDE” mode membership reports. This model referred to as the Source-Specific Multicast (SSM) model. SSM mandates the use of an explicit-join protocol between routers. The standard protocol for this is PIM-SSM, which is simply the subset of PIM-SM used to create (S,G) trees. There are no shared trees (*,G) state in SSM.

Multicast receivers can thus “join” an ASM group G, or “join” (or more accurately “subscribe” to) an SSM (S,G) channel. To avoid repetition of the term “ASM group or SSM channel”, the term (multicast) flow is used, which implies that the flow could be an ASM group or an SSM channel.

Relevant Multicast Protocols / Packet Types

To secure a multicast network it is important to understand the packet types commonly encountered and how to protect against them. There are three main protocols with which to be concerned:

1. IGMP / MLD

2. PIM

3. MSDP

In the next section, each of these protocols are discussed and the issues that can arise with each, respectively.

IGMP / MLD Packets

IGMP / MLD is the protocol used by multicast receivers to signal to a router that they want to receive content for a particular multicast group. Internet Group Membership Protocol (IGMP) is the protocol used in IPv4, and Multicast Listener Discovery (MLD) is the protocol used in IPv6.

There are two versions of IGMP that are commonly deployed, IGMPv2 and IGMPv3. There are also two versions of MLD that are commonly deployed, MLDv1 and MLDv2.

IGMPv2 and MLDv1 are functionally equivalent, and IGMPv3 and MLDv2 are functionally equivalent.

These protocols are specified in these links:

IGMPv2: RFC 2236

MLDv1: RFC 3590

IGMPv3 and MLDv2: RFC 4604

IGMPv2 and IGMPv3 is not only a protocol but also an IPv4 IP protocol (specifically, protocol number 2). It is not only used as described in these RFCs to report multicast group membership, but also by other IPv4 multicast protocols such as DVMRP, PIM version 1, mtrace and mrinfo. This is important to remember when you attempt to filter IGMP (via Cisco IOS® ACLs, for example). In IPv6, MLD is not an IPv6 protocol; instead ICMPv6 is used to carry MLD packets. PIM version 2 is the same protocol type in IPv4 and IPv6 (protocol number 103).

PIM Control Packets

In this section, multicast and unicast PIM control packets are discussed. Auto-RP as well as Bootstrap Router (BSR), which are ways to elect Rendezvous Points and control Group-to-RP assignments in PIM-SM networks, are discussed.

Multicast PIM Control Packets

Multicast PIM Control Packets include:

- **PIM Hello** - The PIM Hello packet is a link-local scope IP multicast packet sent to a router attached to the same network to establish PIM neighbors.
- **PIM Join/Prune** - PIM Join/Prunes are link-local scope IP multicast packets sent to create / remove multicast state and are only sent to PIM neighbors. They are multicast within the LAN to facilitate assert, report suppression, and other PIM protocol details, but they are always directed to a specific neighbor.
- **PIM DF-elect** - PIM Designated Forwarder is the Bi-Dir PIM router responsible for (*,G) JOINS sent to the RP on behalf of attached receivers or downstream PIM neighbours. For cases where a PIM router detects another router that sends (*,G) JOINS on the same segment for the same group G, there is an election to determine the router with the best path to the RP.
- **PIM Assert** - PIM Asserts are link-local IP multicast packets sent when a PIM router attached to a network segment that actively forwards packets for a particular (S,G) out of a particular interface begins RECEIVING packets for that same (S,G) on the same interface on which are forwarded. This event indicates the presence of another router that thinks it is the Single Forwarder (SF) for this (S,G). The Assert mechanism elects a unique SF for that (S,G). The PIM SF router is elected to forward packets for a particular (S,G) stream. PIM allows for different routers to perform the role of the SF on behalf of different (S,G)s, ideally there is only one SF per (S,G). Do not confuse the SF with the Designated Router. The PIM Designated Router is the router responsible for JOIN / PRUNES or SOURCE REGISTERS that are sent to the RP in a PIM-SM network.
- **PIM Bootstrap** - PIM Bootstrap messages are sent in a PIMv2 network to facilitate the dynamic election of a Rendezvous Point for a particular group G.

Unicast PIM Control Packets

Unicast PIM Control Packets are directed to or from the RP and include:

- **Source Register Packet** - PIM Source Register Packets are sent to register a new multicast source with a Rendezvous Point. As soon as a Source starts to send multicast packets, the Designated Router that is attached to the source network sends a unicast register stream to the RP to indicate that there is an active source present for a multicast group for which the RP is responsible.
  Source Register packets are sent as a unicast encapsulation of the original multicast stream. PIM register messages are process-level switched and are sent only until the RP sends a register stop message. The performance impact of these packets is proportional to the rate of the source (per (S,G) flow).
- **Register Stop Packet** - PIM Register Stop Packets are sent from the Rendezvous Point to the PIM DR that sent the Register message. Register Stop messages are sent as soon as the RP starts to receive multicast packets natively from the source.
- **BSR Candidate-Rendezvous Point Advertisement Packet** - PIM BSR C-RP-Advertisement Packets are sent to the BSR to advertise a candidate RP once the BSR is elected.

Fig 1: PIM Unicast Packets

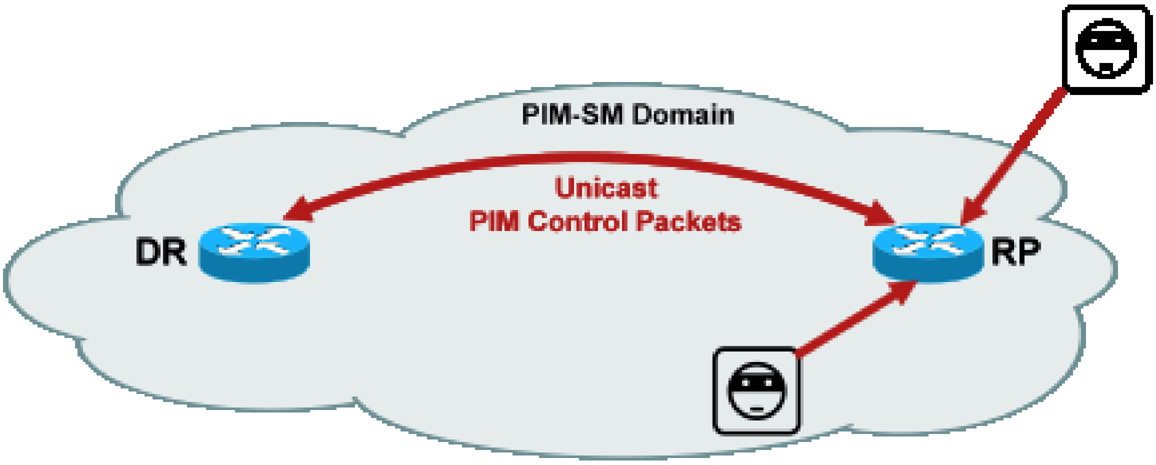

Fig1_PIM_unicast

Attacks that exploit such packets can originate from anywhere, as these packets are unicast.

Auto-RP packets

Auto-RP is a Cisco-developed protocol that serves the same purpose as PIMv2 BSR. Auto-RP was developed before BSR, and only supports IPv4. BSR supports IPv4 and IPv6. The Mapping Agent in Auto-RP serves the same function as the bootstrap router in BSR. In BSR, messages from the C-RP are unicast to the bootstrap router. In Auto-RP, messages are sent via multicast to the Mapping Agent, which allow easier filters at the boundary, as described later. Auto-RP is described in detail in this link: https://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/rps.html

In Cisco IOS, AutoRP/BSR packets are always forwarded, and currently not disabled. This can present a particular security exposure in the case of Auto-RP.

Fig 2: Auto-RP packets

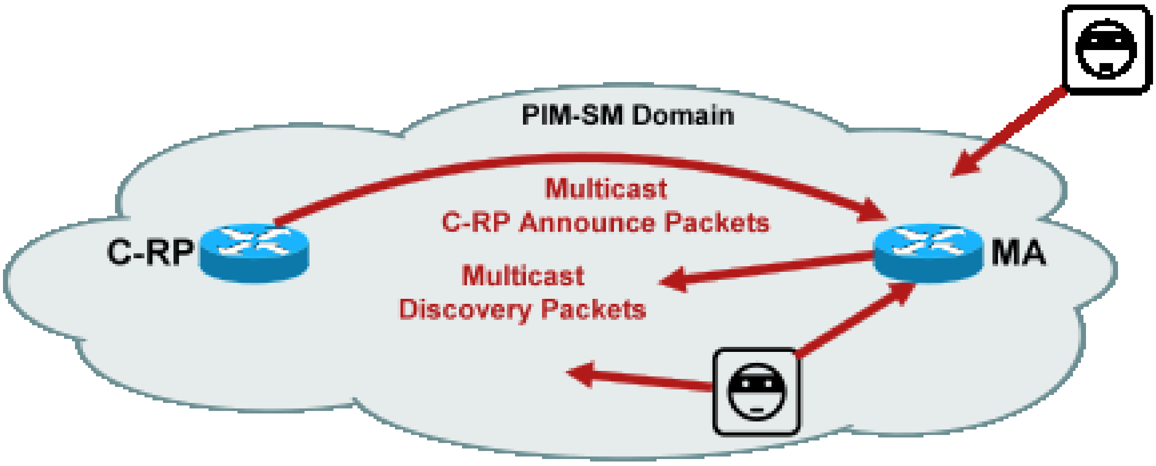

Fig2_A utoRP_packets

> **Note**: Even though Auto-RP is used as a mechanism for PIM-SM RP announcement and discovery, it does not use PIM packets (IP protocol 103); instead it uses User Datagram Protocol (UDP) port 496 packets with multicast addresses.

There are two packet types used by Auto-RP:

- C-RP-Announce packets: these packets are multicast to all Mapping Agents and uses an Internet Assigned Numbers Authority (IANA) reserved “well known” address (224.0.1.39). They are sent by a C-RP to announce the RP address and group range for which that RP is able to act as the RP.
- C-RP Discovery packets: these packets are multicast to all PIM routers and uses an IANA reserved “well known” address (224.0.1.40). They are sent by the Auto-RP Mapping Agent to announce the specific C-RP that is elected as the RP for a particular group range.

Each of these packet types are intended to be flooded through the network.

In Cisco IOS, both 224.0.1.39 and 224.0.1.40 are forwarded in PIM Dense Mode to avoid a problem of no prior knowledge of the RP for a group when that group is used to distribute RP information. This is the only recommended use of PIM Dense Mode.

In Cisco IOS XR, Auto-RP messages are Reverse Path Forwarding (RPF)-flooded hop by hop from neighbor to neighbor. Therefore there is no need to create a PIM DM mroute state to support Auto-RP in Cisco IOS XR. In fact, Cisco IOS XR does not support PIM-DM at all.

Multicast Servivce Discovery Protocol (MSDP) Packets

MSDP is the IPv4 protocol that allows a source in one domain to be announced to a receiver in another domain via their respective rendezvous points. MSDP is specified in RFC 3618.

In order to share information about active sources between PIM domains, MSDP is used. If a source becomes active in one domain then MSDP ensures that all peer domains learn about this new source in a timely manner, which allows receivers in other domains to rapidly make contact with this new source if it happens to have sent to a group in which receivers have an interest in. MSDP is needed for ASM / PIM-SM multicast communications, and runs over a unicast Transport Control Protocol (TCP) connection configured between Rendezvous Points in the respective domains.

Threats in a Multicast Environment

Zones of Trust and Trust Boundaries

This section of the document is organized by functional entities in the network. The threat model discussed is shaped around those entities. For example, this documents explains how a router in a multicast network can be secured (from a multicast point of view), independent of where the router is deployed. Similarly, there are considerations on how to deploy network-wide security measures, or measures on a designated router, rendezvous point, and so on

The threats described here also follow this logic, and are organized by logical function in the network.

Threat Overview

On an abstract level, any multicast deployment can be subject to a number of threats on various aspects of security. The key aspects of security are confidentiality, integrity and availability.

- **Threats against confidentiality**: In most applications, multicast traffic is not encrypted, and is therefore open to anyone to listen or capture on any line or network element in the path. In the section on GET VPN ways to encrypt multicast traffic to prevent such attacks are discussed.
- **Threats against traffic integrity**: Without application-level security or network-based security, such as GET VPN, multicast traffic is vulnerable to modification in transit. This is particularly important for control plane traffic that use multicast, such as OSPF, PIM, and many other protocols.
- **Threats against network integrity**: Without the security mechanisms described in this paper, unauthorized senders, receivers, or compromised network elements can access the multicast network, send and receive traffic without authorization (theft of service), or overload network resources.
- **Threats against availability**: There are a number of denial of service attack possibilities that can make resources unavailable to legitimate users.

The next sections discuss threats for each logical function in the network are discussed.

Basic Threats Against a Router

There are a number of fundamental threats against a router that are independent of whether the router supports multicast and whether the attack involves multicast traffic or protocols.

Denial of service (DoS) attacks are the most important generic attack vectors in a network. In principle, every network element can be targeted with a DoS attack, which can overload the element with potential subsequent loss or degradation of service for legitimate users. It is of paramount importance to follow the basic network security recommendations that apply to unicast.

It is noteworthy that multicast attacks are not always intentional, but often accidental. For instance, the Witty worm, first observed in March 2004, is one example of a worm that spread through random attacks on IP addresses. As a consequence of complete randomization of the address space, multicast IP destinations were also affected by the worm. In many organizations, a number of first-hop routers collapsed because the worm sent packets to many different multicast destination addresses. The routers, however, were not scoped for such a multicast traffic load with the associated state creation, and effectively experienced resource exhaustion. This illustrates the need to secure multicast traffic, even if multicast is not used in an enterprise.

Generic threats against routers include:

- Packet floods of any type; for example, against hardware paths such as slow (punt) paths, and software paths such as management or control plane ports, which includes Secure Shell (SSH), Telnet, Border Gateway Protocol (BGP), OSPF, Network Time Protocol (NTP), and so on
- Intrusions into the router, with subsequent exploitation of features on the router; weak Telnet or SSH passwords and weak Simple Network Management Protocol (SNMP) community strings are a common problem in modern networks.
- Operational issues such as misconfigurations or insider attacks can endanger the security of the entire network and its traffic.

When multicast is enabled on a router, it must be secured in addition to unicast. Usage of IP multicast does not change the fundamental threat model; however, it enables additional protocols (PIM, IGMP, MLD, MSDP) that could be subject to attacks, which need to be secured specifically. When unicast traffic is used in these protocols, the threat model is identical to other protocols run by the router.

It is important to note that multicast traffic cannot be used in the same way as unicast traffic to attack a router because multicast traffic is fundamentally “receiver driven” and cannot be targeted at a remote destination. An attack target needs to be explicitly “joined” to the multicast stream. In most cases (Auto-RP is the main exception), routers only listen to and receive “link local” multicast traffic. Link local traffic is never forwarded. Therefore, attacks on a router with multicast packets can only originate from directly attached attackers.

Threats From the Source Side

Multicast sources, whether PCs or video servers are sometimes not under the same administrative control as the network. Therefore, the sender is mostly treated as untrusted, from the network operator’s point of view. Given the powerful capabilities of PCs and servers, and their complex security settings, which are often incomplete, the senders pose a substantial threat against any network, which includes multicast. These threats include:

- **Layer 2 attacks:** There are a wide range of attack forms on layer 2 to carry out various types of attacks. These apply to unicast as well as multicast. Since these attack forms are not specific to multicast, they are not discussed in more detail in this document. For more information, see the Cisco Press book “LAN Switch Security”, ISBN-10: 1-58705-467-1.
- **Attacks with multicast traffic:** As described previously, it is difficult to conduct attacks with multicast traffic since the first-hop router does not forward multicast traffic unless there is a listener for the group. However, the first hop can be attacked in various ways with multicast packets:
- Network saturation attacks: An attacker can flood a segment with multicast packets, over utilization of the available bandwidth, which can lead to a DoS condition.
- Multicast state attacks: The first-hop router is flooded with multicast packets, which can create too much state, and a consequent DoS attack condition.
- A sender could attempt to become the PIM DR, through PIM hellos that are sent. In such cases no traffic would forward to or from the LAN.
- PIM DF election packets for a BiDir-PIM DF could be spoofed. In such cases no traffic would forward to or from the LAN.
- A sender could spoof AutoRP RP-discovery or BSR bootstrap messages. This would effectively announce a fake RP, and bring down or disrupt a PIM-SM/BiDir service.
- A sender could source unicast attacks, such as PIM source register/register-stop messages, or could send BSR announce packets and announce a fake BSR.
- A sender can send to any valid multicast group, unless this is filtered. If a source address is spoofed and not prevented at the edge, the sender can use the source IP address of a legitimate sender, and override content in parts of the network.
- Multicast attacks against control plane protocols: A number of protocols not associated with multicast, such as OSPF and Dynamic Host Configuration Protocol (DHCP), use multicast packets, which can be used to attack these protocols
- **Masquerading:** There are a number of attack forms where a sender can pretend to be another sender. Spoofed source IP addresses are one such attack form.
- **Theft of service:** Unless senders are controlled, it is possible to use the multicast service illegitimately from the sender side.

Note: Hosts normally do not send or receive PIM packets. Host that do this can likely attempt an attack.

Threats From the Receiver Side

The receiver is also typically a platform with significant CPU power and bandwidth, and allows for a number of attack forms. These are mostly identical to the threats on the sender side. Layer 2 attacks remain an important attack vector. Fake receivers and theft of service are also possible on the receiver side, except that the attack vector is typically IGMP (or layer-2 attacks, as mentioned).

Threats Against a Rendezvous Point and BSR

PIM-SM RPs and PIM-BSRs are critical points in a multicast network, and are therefore valuable targets for an attacker. When neither is the first-hop router, only unicast attack forms, which includes PIM unicast, can be targeted directly against those elements. The threats against RPs and BSRs include:

- All generic attack forms, as described in the section “Basic Threats Against a Router”.
- PIM unicast attacks, potentially with spoofed source IP addresses, allow for DoS attacks, though PIM register or register-stop messages that are sent by a malicious device.

Multicast and Unicast Security (compared)

State Considerations / Filters

Consider the topology in Figure 3, which shows a source, three receivers (A, B, C), a switch (S1), and two routers (R1 and R2). The blue line represents a unicast stream and the red line represents a multicast stream. All three receivers are members of the multicast flow.

Fig 3: Replication in Routers and Switches

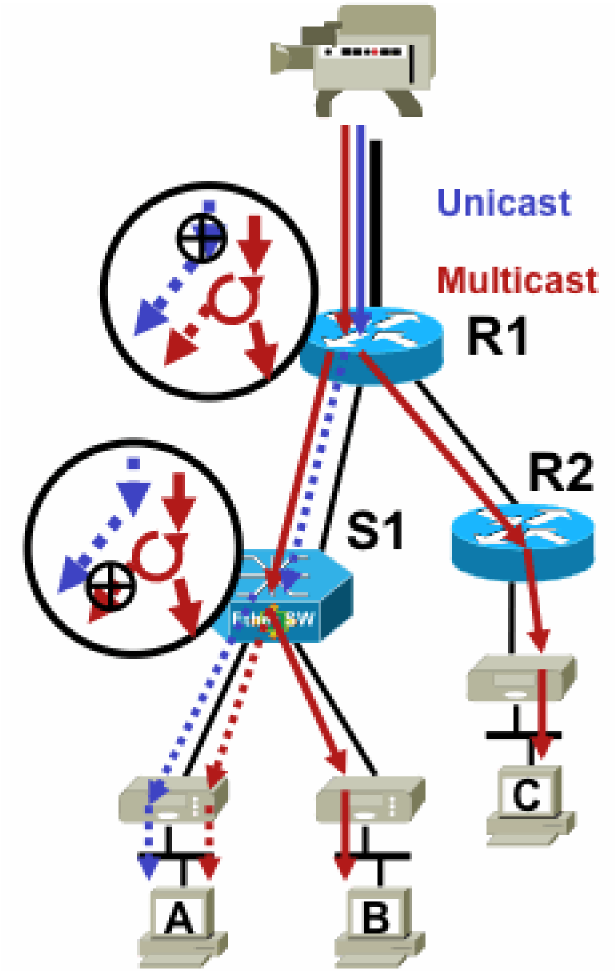

Fig3_replication_RS

To inhibit traffic flow from a specific source to a specific receiver:

- For the unicast stream, install a filter anywhere on the path from sender to receiver.
- For the multicast stream, however, administrators need to be more specific about where to install filters: at the receiver side filter after the last replication point before the receiver; at the source side filter before the first replication point after the source.

Attacks From Multicast Sources

This section applies to both the ASM and SSM service models, where traffic is forwarded based on the receipt of receiver-side explicit joins.

For unicast streams there is no implicit receiver protection. A unicast source can send traffic to a destination, even if this destination has not asked for the traffic. Therefore, defense mechanisms such as firewalls are typically used to protect end points. Multicast, on the other hand, has some implicit protection built into the protocols. Traffic ideally only reaches a receiver that has joined the flow in question.

With ASM, sources can launch traffic insertion or DoS attacks through multicast traffic transmission to any of the groups supported by an active RP. This traffic ideally does not reach a receiver, but can reach the first-hop router in the path at minimum, as well as the RP, which allows for limited attacks. If a malicious source, however, knows a group to which a target receiver is interested in, and if there are no appropriate filters in place, it can send traffic to that group. This traffic is received as long as receivers listen to the group.

With SSM, attacks by unwanted sources are only possible on the first-hop router where the traffic stops if no receiver has joined that (S,G) channel. This does not lead to any state attack on the first-hop router because it discards all SSM traffic for which no explicit join state exists from receivers. In this model it is not sufficient for a malicious source to know which group a target is interested in to because “joins” are source-specific. Here, IP source addresses that are spoofed plus potential routing attacks would be required to succeed.

State Attacks

Even without receivers present in a network, PIM-SM creates (S,G) and (*,G) state on the first-hop router closest to the source and also on the Rendezvous Point. Thus there exists the possibility of a state attack on the network at the source first-hop router and on the PIM-SM RP.

If a malicious source starts to send traffic to multiple groups, then for each of the groups that are detected, the routers in the network create state at the source and the RP, provided the groups in question are permitted by the RP configuration.

Therefore, PIM-SM is subject to state and traffic attacks by sources. The attack can be aggravated if the source is changes its source IP address randomly within the correct prefix, or in other words, only the host bits of the address are spoofed.

Fig 4: ASM RP attacks

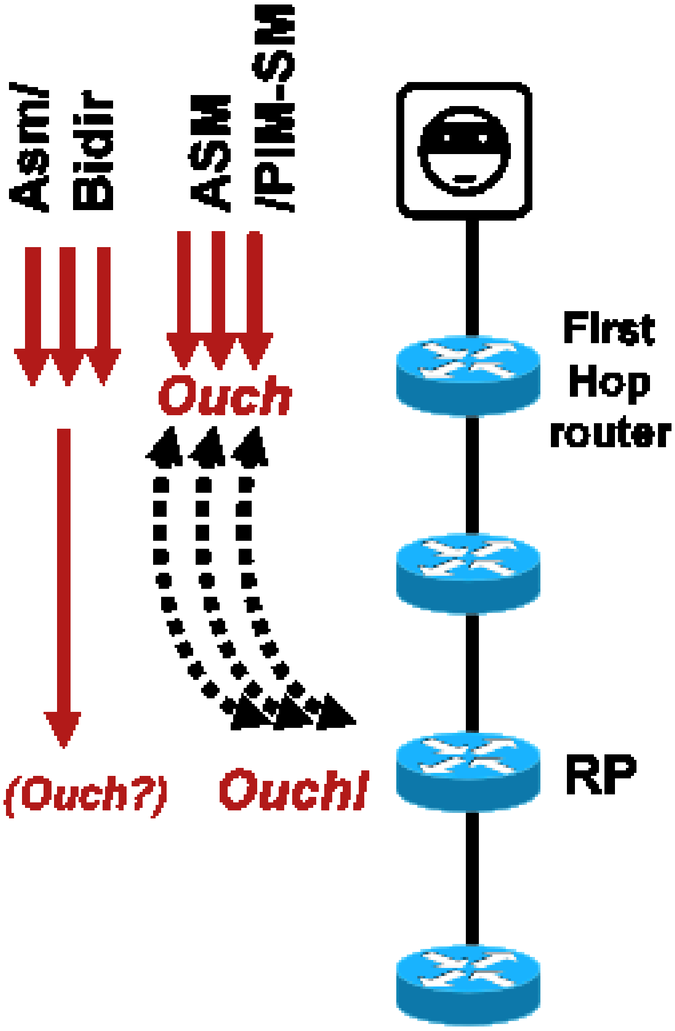

Fig4_ASM_RP_Attacks

As with PIM-SSM, PIM-BiDir state-creation attacks from sources are impossible. Traffic in PIM-BiDir is forwarded on state created by joins from receivers as well as on state forwarded traffic to the RP, such that it can reach receivers behind the RP, since the joins only go to the RP. State-to-forward traffic to the RP is called (*,G/M) state and is created by RP configuration (static, Auto-RP, BSR). It does not change in the presence of sources. Therefore, attackers can send multicast traffic to a PIM-BiDir RP, but unlike PIM-SM, a PIM-BiDir RP is not an “active” entity, and instead just forwards or discards traffic for PIM-BiDir groups.

> **Note**: On some Cisco IOS platforms (*,G/M) state is not be supported. In such cases sources can attack the router by multicast traffic transmission to multiple PIM-BiDir groups, which causes (*,G) state creation. For example the Catalyst 6500 switch does support (*,G/M) states).

Receiver-Initiated Attacks

Attacks can originate from multicast receivers. Any receiver that sends an IGMP/MLD reports typically create state on the first-hop router. There is no equivalent mechanism in unicast.

Fig 5: Receiver-Side Explicit Join-Based Traffic Forwarding

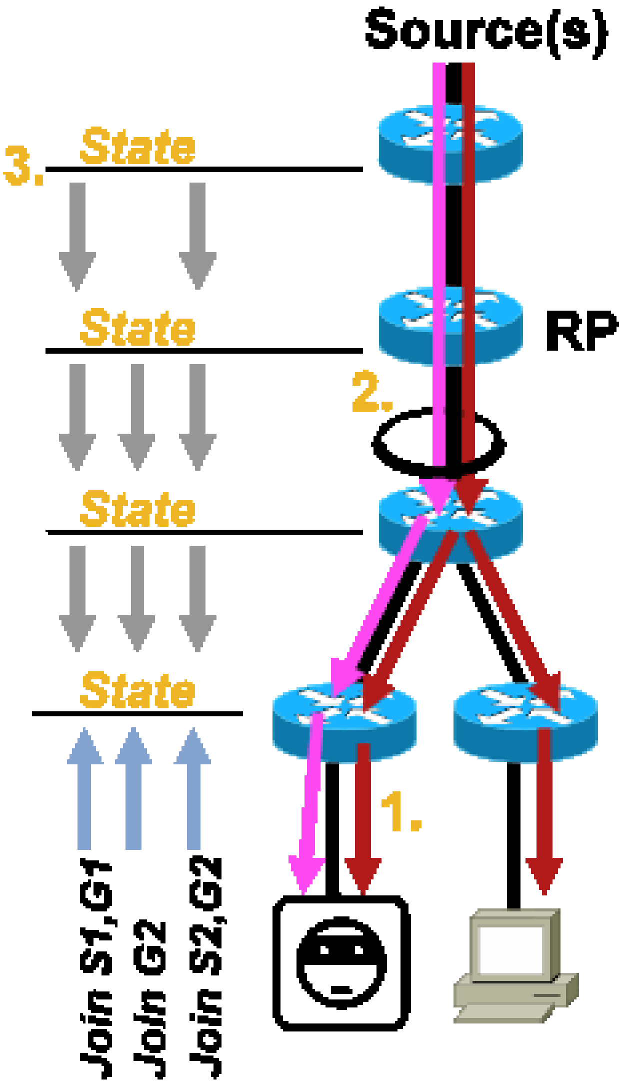

Fig5_Receiver_Explicit_Join

Receiver attacks can be of three types:

1. A multicast receiver can attempt to join a flow for which they are not authorized and attempt to receive content they are unauthorized to receive.
2. A multicast receiver can potentially overload available network bandwidth through interest in many groups or channels. This sort of attack becomes a shared bandwidth attack against other potential receivers of content.
3. A multicast receiver can attempt to launch an attack against routers or switches. A large number of IGMP reports can be generated, which can create a large amount of multicast tree state and potentially overload router capacity. This in turn can result in an increase in multicast convergence times, or in a DoS on the router.

Various ways to mitigate these sorts of attacks in the next section, Security within a Multicast Network.

Security within a Multicast Network

Network Element Security

Security is not a point feature, but an intrinsic part of every network design. As such, security must be considered at every point in the network. It is of paramount importance that each and every network element is appropriately secured. One possible attack scenario, applicable to any technology, is a router subverted by an intruder. Once an intruder has control of a router, the attacker can run a number of different attack scenarios. Each network element must therefore be appropriately secured against any form of basic attack, as well as against specific multicast attacks.

Control Plane Policing (CoPP)

CoPP is the evolution of Router ACLs (rACLs), and available on most platforms. The principle is the same: only traffic destined to the router is policed by CoPP.

The service policy uses the same syntax as any quality of service policy, with policy-maps and class-maps. Therefore, it extends the functionality of rACLs (permit/deny) with rate-limiters for certain traffic towards the control plane.

> **Note**: Certain platforms, such as the Catalyst 9000 series switches have CoPP enabled by default, and the protection is not superseded. See CoPP guide for additional information.

If you decide to adjust,modify, or create rACLS or CoPP in a live network, care must be taken. Since both features have the ability to filter all traffic to the control plane, all required control and management plane protocols must be explicitly permitted. The list of required protocols is large, and it can be be easy to overlook less obvious protocols such as Terminal Access Controller Access Control System (TACACS). All non-default rACL and CoPP configurations must always be tested in a lab environment before deployment to production networks. Furthermore, initial deployments need to start with a "permit" policy only. This allows validation of any unexpected hits with ACL hit counters.

In a multicast environment, the required multicast protocols (PIM, MSDP, IGMP, etc) must be permitted in rACL or CoPP for multicast to function properly. It is important to remember that the first packet in a multicast stream from the source in a PIM-SM scenario is used as a control plane packet, to help create multicast state, up at the control plane of the device. Therefore it is important to permit relevant multicast groups in rACL or CoPP. Since there are a number of platform-specific exceptions it is important to consult relevant documentation and test any planned configuration before deployment.

Local Packet Transport Service (LPTS)

On Cisco IOS XR, Local Packet Transport Service (LPTS) serves as a policer of traffic to the control plane of the router, similar to CoPP on Cisco IOS. Additionally, receive traffic, which includes unicast and multicast traffic, can be filtered and rate limited.

Multicast-Specific Security

In a multicast-enabled network, each network element needs to be secured with multicast-specific security features. These are outlined in this section, for generic router protection. Features that are not required on every router, but only in specific locations in the network, and features that require interaction between routers (such as PIM authentication) are discussed in the next section.

Mroute Limits

The mroute limit command limits the amount of multicast routes globally on a router, and helps to prevent DoS attacks.

```
ip multicast route-limit <mroute-limit> <warning-threshold>
```

Fig 6: Mroute Limits

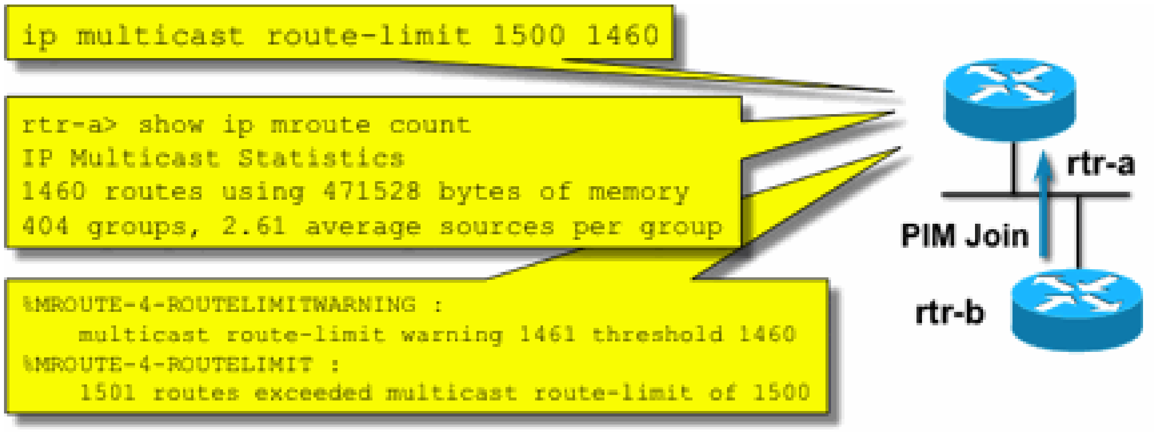

```
ip multicast route-limit 1500 1460
```

```
rtr-a> show ip mroute count
IP Multicast Statistics
1460 routes using 471528 bytes of memory
404 groups, 2.61 average sources per group
```

```
%MROUTE-4-ROUTELIMITWARNING :
    multicast route-limit warning 1461 threshold 1460
%MROUTE-4-ROUTELIMIT :
    1501 routes exceeded multicast route-limit of 1500
```

Fig6_Mroute_Limits

Mroute limits allow the setting of a threshold on the number of mroutes that are permitted into the multicast routing table. If a multicast route limit is enabled then no multicast state is created beyond the configured limit. There is also a warning threshold. When the number of mroutes exceeds the warning threshold, syslog warning messages are triggered. At the mroute limit any further packets that would trigger state are discarded.

The **ip multicast route-limit** command is also available per MVRF.

Disable SAP Listen: no ip sap listen

The **sap listen** command causes a router to receive Session Announcement Protocol/Session Description Protocol (SAP/SDP) messages. SAP/SDP is a legacy protocol that dates from the days of the multicast backbone (MBONE). These messages indicate directory information about multicast content that are available in the future or at present. This can be a source of a DoS against router CPU and memory resources, and therefore this feature needs to be disabled.

Control access to mrinfo information - the "ip multicast mrinfo-filter" command

The mrinfo command (available on Cisco IOS and also on some versions of Microsoft Windows and Linux) uses various messages to query a multicast router for information. The **ip multicast mrinfo-filter** global configuration command can be used to limit access to this information to a subset of sources, or disable it altogether.

This example denies queries sourced from 192.168.1.1, while queries are allowed from any other source:

```
ip multicast mrinfo-filter 51

access-list 51 deny 192.168.1.1
access-list 51 permit any
```

This example denies *mrinfo* requests from any source:

```
ip multicast mrinfo-filter 52

access-list 52 deny any
```

> **Note**: As expected with any ACL, a *deny* means the packet is filtered, while a *permit* means the packet is allowed.

If the **mrinfo** command is used for diagnostic purposes, it is highly recommended to configure the **ip multicast mrinfo-filter** command with an appropriate ACL to allow queries only from a subset of source addresses. The information provided by the *mrinfo* command can also be retrieved through SNMP. Complete blocks of mrinfo requests (block any source from queries of the device) is highly recommended.

Network Security

In this section various ways to secure PIM multicast and unicast control packets, as well as Auto-RP and BSR, are discussed..

Disable Multicast Groups

The **ip multicast group-range/ipv6 multicast group range** commands can be used to disable all operations for groups denied by the ACL:

```
ip multicast group-range <std-acl>
ipv6 multicast group-range <std-acl>
```

If packets appear for any of the groups denied by the ACL, they are dropped in all control protocols, which includes PIM, IGMP, MLD, MSDP, and are also dropped on the data plane. Therefore, no IGMP/MLD cache entries, PIM, Multicast Routing Information Base/Multicast Forwarding Information Base (MRIB/MFIB) state are ever created for these group ranges and all data packets are immediately dropped.

These commands are entered in global configuration of the device.

The recommendation is to deploy this command on all routers in the network, when and where available, so that all multicast traffic that originates outside the network is controlled. Note that these commands affect the data plane and control plane. Where available, this command provides more extensive coverage than standard ACLs, and is preferred.

PIM Security

PIM Neighbor Control

A PIM router must receive PIM Hellos to establish PIM Neighborship. PIM Neighborship is also the basis for Designated Router (DR) election, and DR failover as well as sent / received PIM Join/Prune/Assert messages.

Fig 7: PIM Neighbor Control

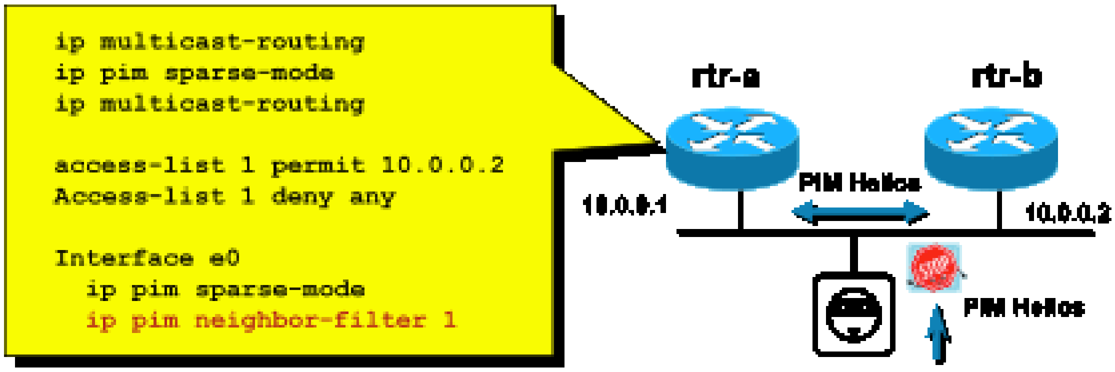

```
ip multicast-routing
ip pim sparse-mode
ip multicast-routing

access-list 1 permit 10.0.0.2
Access-list 1 deny any

Interface e0
  ip pim sparse-mode
  ip pim neighbor-filter 1
```

Fig7_PIM_neighbor_control

To inhibit unwanted neighbors use the **ip pim neighbor-filter** command illustrated in Figure 7. This command filters from all non-allowed neighbors PIM packets, which includes Hellos, Join/Prune packets, and BSR packets. Hosts on the segment can potentially spoof the source IP address to pretend to be the PIM neighbor. Layer 2 security mechanisms (namely IP source guard) are required to prevent source addresses from a spoof attempt on a segment or use a VLAN ACL in the access switch to prevent PIM packets from hosts. The keyword “log-input” can be used in ACLs to log packets that match the ACE.

The PIM Join/Prune packet is sent to a PIM neighbor to add or remove that neighbor from a particular (S,G) or (*,G) path. PIM multicast packets are link local multicast packets sent with a Time-To-Live (TTL)=1. All of these packets are multicast to the well known All-PIM-Routers address: 224.0.0.13 . This means that all such attacks must originate on the same subnet as the router that is attacked. Attacks can include forged Hello, Join/Prune, and Assert packets.

> **Note**: An artificial increase or adjustment of the TTL value in PIM multicast packets to a higher value than 1 does not create problems. The All-PIM-Routers address is always received and treated locally on a router. It is never directly forwarded by normal and legitimate routers.

To protect the RP against a potential flood of PIM-SM register messages, the DR needs to rate limit those messages. Use the **ip pim register-rate-limit** command:

```
ip pim register-rate-limit <count>
```

Fig 8: PIM-SM Register Tunnel Control

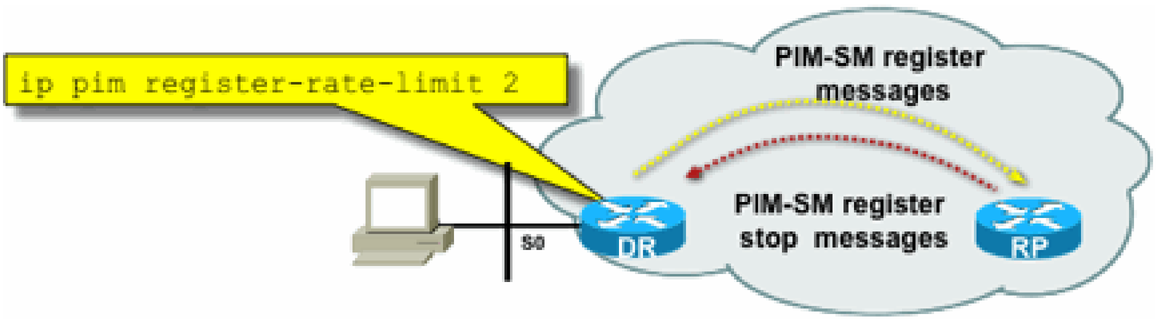

```
ip pim register-rate-limit 2
```

Fig8_PIMSM_Reg Tunnel

PIM unicast packets can be used to attack the RP. Therefore, the RP can be protected by infrastructure ACLs against such attacks. Remember, multicast senders and receivers never need to send PIM packets, so the PIM protocol (IP protocol 103) can usually be filtered at the subscriber edge.

Auto-RP Control - RP Announce Filter

The **ip pim rp-announce filter** command is an additional security measure that can be configured with Auto-RP where possible:

```
ip pim rp-announce-filter
```

This can be configured on the Mapping Agent to control which routers are accepted as Candidate RPs for which group ranges / group-mode.

Fig 9: Auto-RP – RP Announce Filter

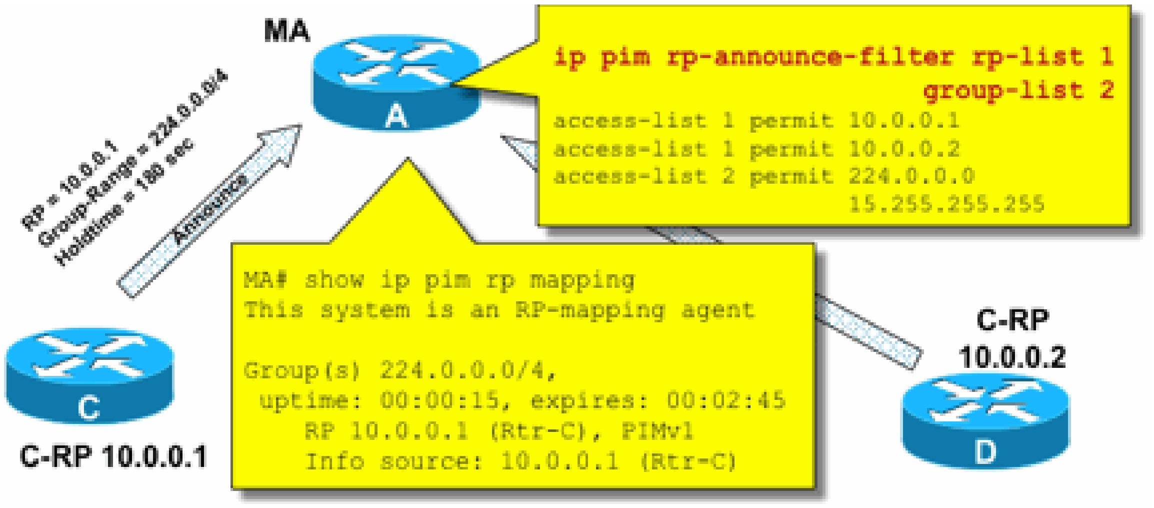

```
ip pim rp-announce-filter rp-list 1
                          group-list 2
access-list 1 permit 10.0.0.1
access-list 1 permit 10.0.0.2
access-list 2 permit 224.0.0.0
                     15.255.255.255
```

```
MA# show ip pim rp mapping
This system is an RP-mapping agent

Group(s) 224.0.0.0/4,
  uptime: 00:00:15, expires: 00:02:45
     RP 10.0.0.1 (Rtr-C), PIMv1
     Info source: 10.0.0.1 (Rtr-C)
```

Fig9_AutoRP_RP_Announce

Auto-RP Control - Constrain Auto-RP Messages

Use the multicast boundary command to constrain AutoRP packets, RP-announce (224.0.1.39) or RP-discover (224.0.1.40) to a particular PIM domain:

```
ip multicast boundary <std-acl>

access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40

224.0.1.39 (RP-announce) 224.0.1.40 (RP-discover)
```

Fig 10: Multicast Boundary Command

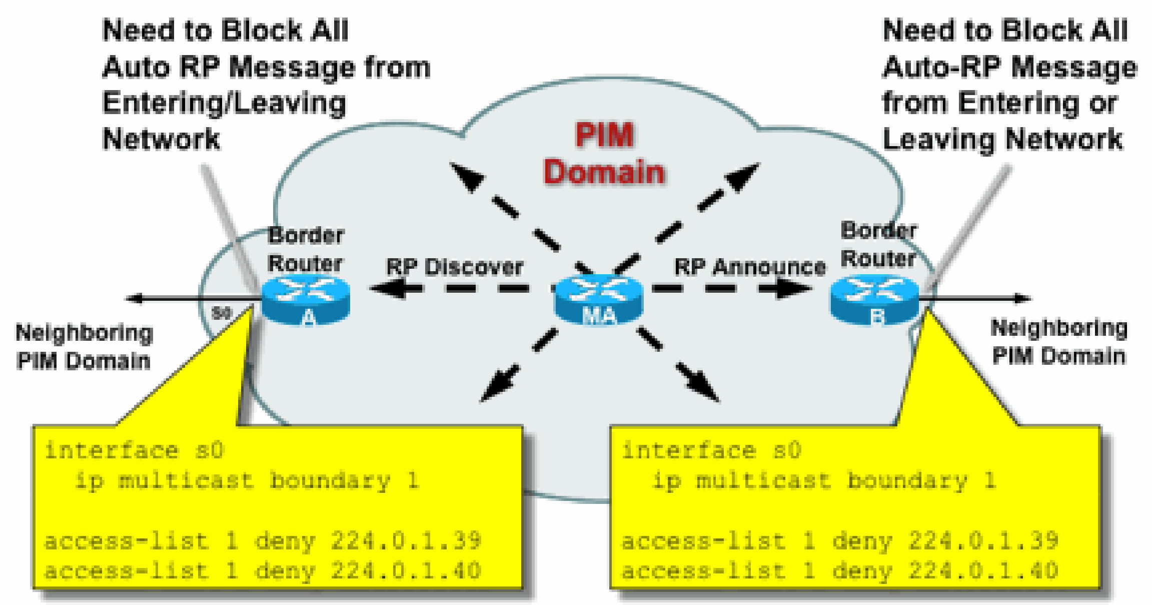

```
interface s0
  ip multicast boundary 1

access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40
```

```
interface s0
  ip multicast boundary 1

access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40
```

Fig10_Mcast_Boundary

BSR Control - Constrain BSR Messages

Use the **ip pim bsr-border** command to filter BSR messages at the border of a PIM domain. No ACL is necessary since BSR messages are hop-by-hop forwarded with link local multicast.

Fig 11: BSR Border

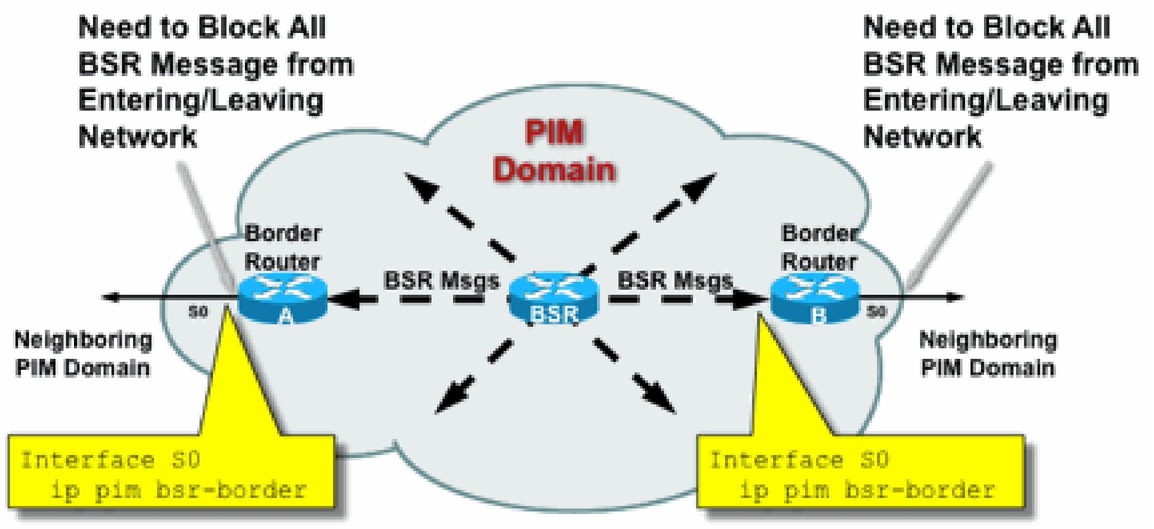
Need to Block All BSR Message from Entering/Leaving Network
Need to Block All BSR Message from Entering/Leaving Network
PIM Domain
Border Router
Border Router
BSR Msgs
BSR Msgs
A
BSR
B
S0
S0
Neighboring PIM Domain
Neighboring PIM Domain
Interface S0 ip pim bsr-border
Interface S0 ip pim bsr-border

Fig11_BSR_Router

RP / PIM-SM-related Filters

As part of this final section, filters against PIM-SP and RP control plane packets as well as Auto-RP, BSR, and MSDP messages are discussed.

Auto-RP Filters

Figure 12 shows an example of Auto-RP filters in conjunction with address scopes. Two different ways to bound a region are shown. The two ACLs are equivalent from an Auto-RP perspective.

Fig 12: Auto-RP Filters / Scopes

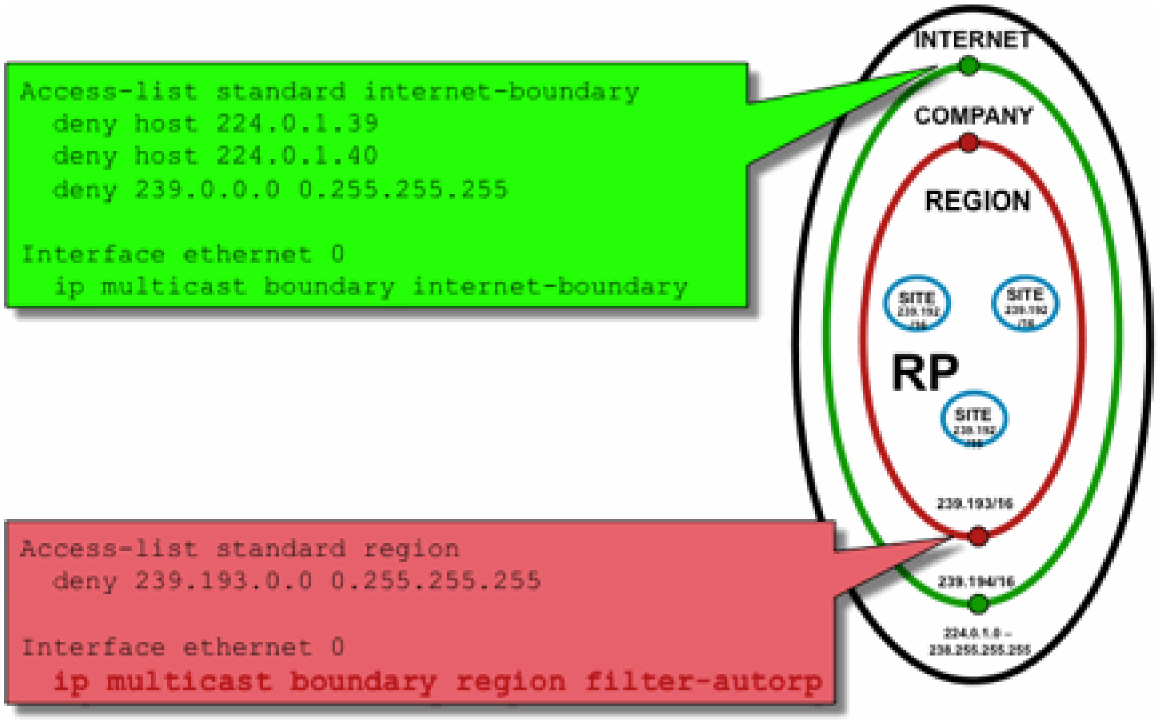

```
Access-list standard internet-boundary
  deny host 224.0.1.39
  deny host 224.0.1.40
  deny 239.0.0.0 0.255.255.255

Interface ethernet 0
  ip multicast boundary internet-boundary
```

```
Access-list standard region
  deny 239.193.0.0 0.255.255.255

Interface ethernet 0
  ip multicast boundary region filter-autorp
```

Fig12_AutoRP_Filtering_Scoping

The idea of the interface boundary filters for Auto-RP is to ensure that the auto-rp announcements only reach the regions they support. Regional, Company and Internet-wide scopes are defined, and in each case there are RPs and Auto-RP advertisements in each scope. Administrators only want the Regional RPs to be known to the Regional routers, the Company RPs to be known to the Regional and Company routers, and want any Internet RPs to be globally available. Further levels of scopes are possible.

As shown in the picture, there are two fundamentally different ways to filter Auto-RP packets: The Internet boundary explicitly calls out the auto-rp control groups (224.0.1.39 and 224.0.1.40), which results in filters against all Auto-RP packets. This method can be used at the edge of an administrative domain, where no Auto-RP packets are passed through. The Region boundary uses the filter-auto-rp keyword to cause an examination of the rp-to-group-range announcements within Auto-RP packets. When an announcement is explicitly denied by the ACL, it is removed from the Auto-RP packet before the packet is forwarded. In the example, this allows enterprise-wide RPs to be known within the regions, while the region-wide RPs are filtered at the boundary from the region to the rest of the enterprise.

Inter-Domain Filters and MSDP

In this example, ISP1 acts as a PIM-SM transit provider. They only support MSDP peering with neighbors and they only accept (S,G), but no (*,G) traffic on the border routers.

In inter-domain (usually between Autonomous Systems) there are two basic security measures to be taken:

1. Secure the data plane, through the **multicast boundary** command. This ensures that

multicast traffic is only accepted for defined groups (and potentially sources).

2. Secure the inter-domain control plane traffic (MSDP). This consists of a number of separate security measures: MSDP content control, state limitation, and neighbor authentication.

Figure 13 provides example configuration of an interface filter on one of ISP1's border routers.

To secure the data plane at the domain boundary inhibit (*,G) joins by filters against “host 0.0.0.0” and administratively scoped addresses via the **multicast boundary** command:

Fig 13: Interdomain (*,G) filter

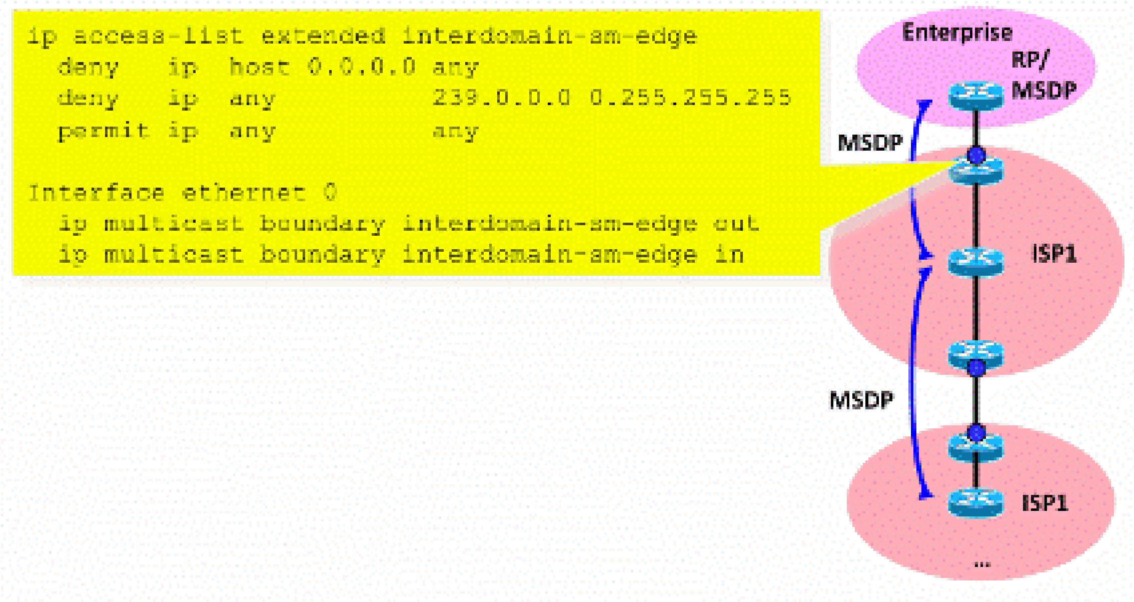

```
ip access-list extended interdomain-sm-edge
  deny   ip  host 0.0.0.0 any
  deny   ip  any          239.0.0.0 0.255.255.255
  permit ip  any          any

Interface ethernet 0
  ip multicast boundary interdomain-sm-edge out
  ip multicast boundary interdomain-sm-edge in
```

Fig13_Interdomain_Filter

To secure the control plane, harden MSDP via three basic security measures:

1) MSDP SA Filters

It is a “best common practice” to filter the content of MSDP messages via MSDP SA filters. The main idea of this filter is to avoid propagation of multicast state for applications and groups that are not Internet-wide applications and do not need to be forwarded beyond the source domain. Ideally, from a security point of view, the filters only allow known groups (and potentially senders), and deny any unknown senders and/or groups.

It is usually not possible to explicitly list all allowed senders and/or groups. it is recommended to use the default configuration filter for PIM-SM domains with a single RP for every group (no MSDP mesh-group):

```
!--- Filter MSDP SA-messages.
          !--- Replicate the following two rules for every external MSDP peer.
          !
```

```
        ip msdp sa-filter in <peer_address> list 111
        ip msdp sa-filter out <peer_address> list 111
        !
        !--- The redistribution rule is independent of peers.
        !
        ip msdp redistribute list 111
        !
        !--- ACL to control SA-messages originated, forwarded.
        !
        !--- Domain-local applications.
        access-list 111 deny   ip any host 224.0.2.2     !
        access-list 111 deny   ip any host 224.0.1.3     ! Rwhod
        access-list 111 deny   ip any host 224.0.1.24    ! Microsoft-ds
        access-list 111 deny   ip any host 224.0.1.22    ! SVRLOC
        access-list 111 deny   ip any host 224.0.1.2     ! SGI-Dogfight
        access-list 111 deny   ip any host 224.0.1.35    ! SVRLOC-DA
        access-list 111 deny   ip any host 224.0.1.60    ! hp-device-disc
        !--- Auto-RP groups.
        access-list 111 deny   ip any host 224.0.1.39
        access-list 111 deny   ip any host 224.0.1.40
        !--- Scoped groups.
        access-list 111 deny   ip any 239.0.0.0 0.255.255.255
        !--- Loopback, private addresses (RFC 6761). access-list 111 deny ip 10.0.0.0
0.255.255.255 any access-list 111 deny ip 127.0.0.0 0.255.255.255 any access-list 111 deny ip
172.16.0.0 0.15.255.255 any access-list 111 deny ip 192.168.0.0 0.0.255.255 any !--- Default
SSM-range. Do not do MSDP in this range. access-list 111 deny ip any 232.0.0.0 0.255.255.255
access-list 111 permit ip any any !
```

It is recommended to filter as strictly as possible, and in both directions, inbound and outbound.

Please use for more details on MSDP SA filter recommendations: https://www.cisco.com/c/en/us/support/docs/ip/ip-multicast/13717-49.html

2) MSDP State Limitation

When MSDP is enabled between multiple autonomous systems (AS) it is recommended to limit the amount of state that is built in the router due to “Source-Active” (SA) messages received from neighbors. You can use the **ip msdp sa-limit** command:

```
ip msdp sa-limit <peer> <limit>
```

Fig 14: MSDP Control Plane

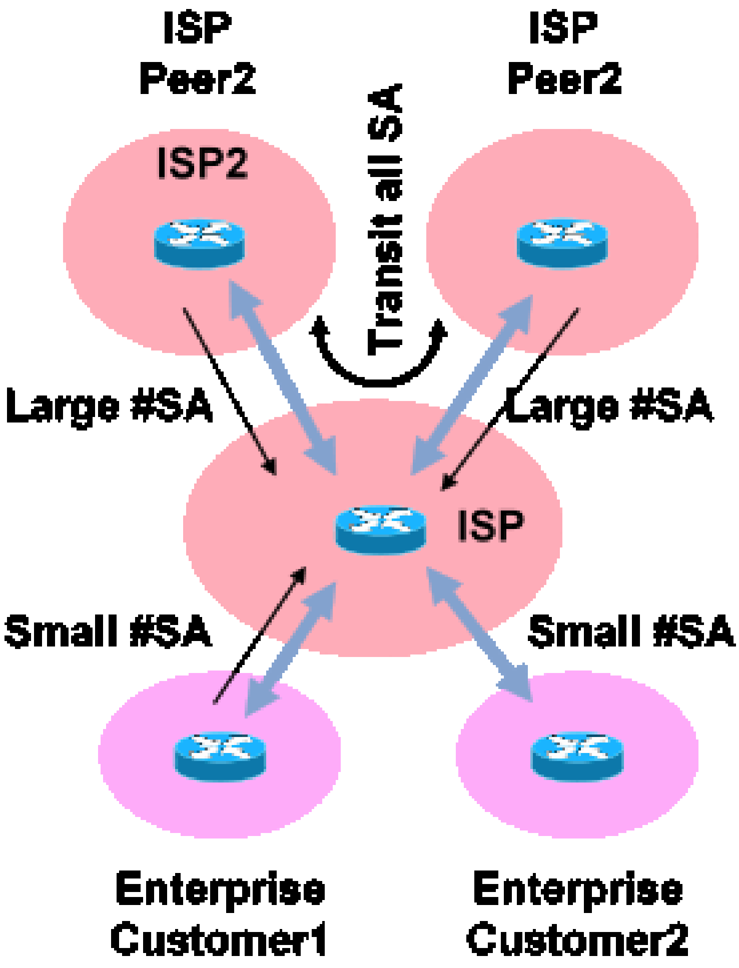

Fig14_MSDP_ControlPlane

With the **ip msdp sa-limit** command you can limit the number of SA states created due to SA messages accepted from an MSDP peer. Some simple rule-of-thumb recommendations include:

- Small limit from stub-neighbor
- Large limit from transit neighbor (for example maximum #SAs in Internet)
- Transit ISP - configure maximum #SAs your platform can support

3) MSDP MD5 Neighbor Authentication

It is recommended to use of Message-Digest Algorithm (MD5) password authentication on MSDP peers. This uses the TCP MD5 signature option, equivalent to the use described in RFC 6691 to secure BGP.

Fig 15: MSDP MD5 Neighbor Authentication

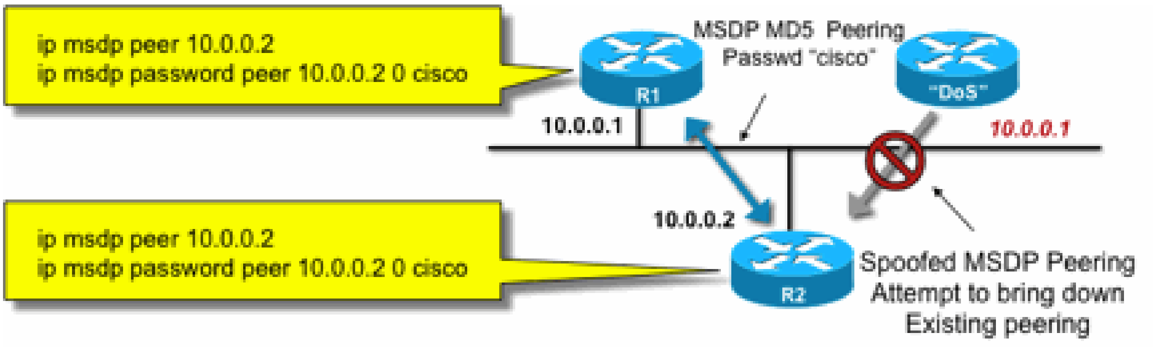

Fig15_MSDP_MD5Auth

These three MSDP security recommendations pursue different goals:

- Neighbor authentication (with MD5) ensures that only trusted MSDP peers can send messages.
- The SA filters ensure that even a trusted MSDP peer can only send SA announcements that are in line with pre-agreed source/group policy.
- The SA limit further ensures that even with legitimate (S,G) announcements from legitimate peers, the available memory cannot be exhausted.

Sender / Source issues

Many multicast security issues that originate at the sender can be mitigated with appropriate unicast security mechanisms. A number of unicast security mechanisms are recommended best practices here:

- **Source address spoof protection** (Unicast Reverse Path Forwarding, uRPF or ACL and IP source guard for the access layer)
- **Infrastructure ACLs** (deny ip any (to) <core address space>)

Such measures can be used to block directed attacks on the core. This would, for example, also solve issues like attacks that use PIM unicast packets to the RP, which is "inside" the network and would therefore be protected by the infrastructure ACL.

Packet Filter-Based Access Control – Control Sources

In the example shown in Figure 16, the filter is configured on the LAN interface (E0) of the first-hop multicast router (Designated Router). The filter is defined by an Extended Access Control List called “source”. This ACL is applied to the source-facing interface of the Designated Router connected to the source LAN. In fact, because of the nature of multicast traffic, there could need to be a similar filter configured on all LAN-facing interfaces on which sources could become active. Since it is not possible in all cases to know exactly where source activity occurs, it is recommended to apply such filters on all ingress points into the network.

Fig 16: Control Sources

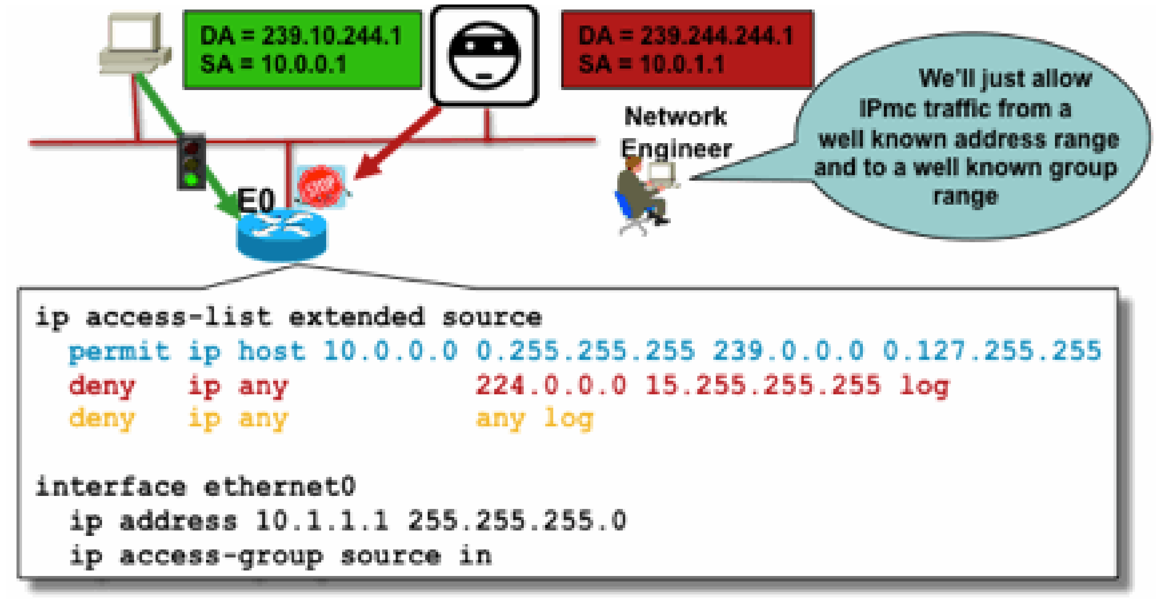
DA = 239.10.244.1
SA = 10.0.0.1
DA = 239.244.244.1
SA = 10.0.1.1
Network Engineer
We'll just allow IPmc traffic from a well known address range and to a well known group range
E0

```
ip access-list extended source
  permit ip host 10.0.0.0 0.255.255.255 239.0.0.0 0.127.255.255
  deny   ip any            224.0.0.0 15.255.255.255 log
  deny   ip any            any log

interface ethernet0
  ip address 10.1.1.1 255.255.255.0
  ip access-group source in
```

Fig16_Controlling _Sources

The purpose of this filter is to prevent traffic from a specific source or range of source addresses to a specific group or range of group addresses. This filter acts before PIM creates any mroutes and helps limit state.

This is a standard data plane ACL. This is implemented on ASICs on high-end platforms, and no performance penalty is incurred. Data plane ACLs are recommended and preferred over control plane for directly connected sources because they minimize any control plane impact of unwanted traffic. It is also very effective to limit the destination (IP multicast group addresses) to which packets can be sent. As this is a router command, it cannot overcome an source-IP address that is spoofed (see earlier part of this section). Therefore, it is recommended to either provide additional layer 2 (L2) mechanisms or a consistent policy for all devices that can connect to a particular local area network/virtual local area network (LAN/VLAN).

> **Note**: The “log” keyword in an ACL is very useful to understand hits against a specific ACL entry; however, this consumes CPU resources, and need to be handled with care. Also, on hardware-based platforms, ACL log messages are produced by a CPU, and therefore the CPU impact must be considered.

PIM-SM Source Control

One of the actual advantages of the ASM / PIM-SM architecture from a security standpoint is the fact that the Rendezvous Point gives a single point of control for all sources in the network for any group range. This can be leveraged with a device called the accept-register filter. The command for this filter is as follows:

```
ip pim accept-register / ipv6 pim accept-register
```

Fig 17: PIM-SM Source Control

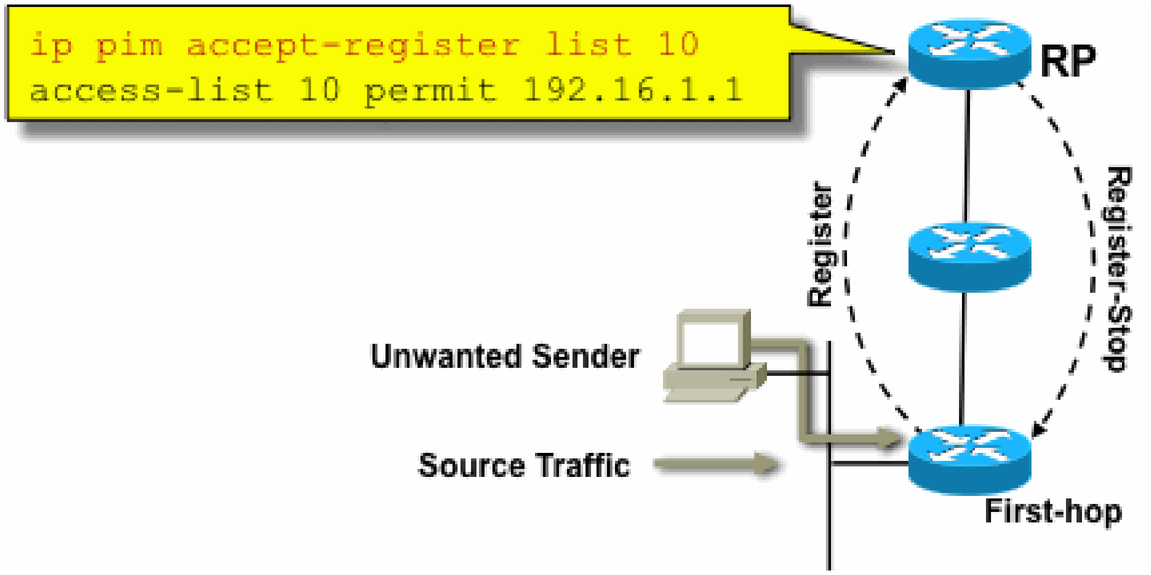

```
ip pim accept-register list 10
access-list 10 permit 192.16.1.1
```

Fig17_PIMSM_ Control

In a PIM-SM network, an unwanted traffic source can be controlled with this command. When the source traffic hits the first-hop router, the first-hop router (DR) creates (S,G) state and sends a PIM Source Register message to the RP. If the source is not listed in the accept-register filter list (configured on the RP), then the RP rejects the Register and sends back an immediate Register-Stop message to the DR.

In the example shown, a simple ACL has been applied to the RP, which filters only on the source address. It is also possible to filter the source AND the group with the use of an extended ACL on the RP.

There are drawbacks with source filters because with the **pim accept-register** command on the RP, PIM-SM (S,G) state is still created on the source's first-hop router. This can result in traffic at receivers local to the source and located between the source and the RP. Furthermore, the **pim accept-register** command works on the control plane of the RP. This could be used to overload the RP with fake register messages, and possibly cause a DoS condition.

It is recommended to apply the pim accept-register command on the RP in addition to other methods, such as application of simple data plane ACLs on all DRs, on all ingress points into the network. While ingress ACLs on the DR would be sufficient in a perfectly configured and operated network, it is recommended to configure the **pim accept-register** command on the RP as a secondary security mechanism in case of misconfigurations on the edge routers. Layered security mechanisms with the same goal is called "defense in depth", and is a common design principle in security.

Receiver Issues – Control IGMP/MLD

Most receiver issues fall in the domain of IGMP/MLD receiver protocol interactions.

Fig 18: Control IGMP

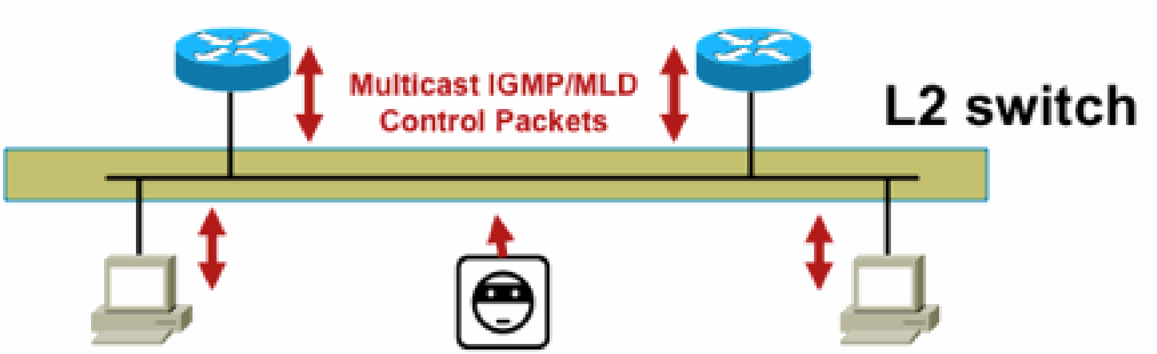

Fig18_Controlling_IGMP

When IGMP or MLD packets are filtered, remember these points:

- IPv4: IGMP is an IPv4 protocol type (IPv4 procotol 2)
- IPv6: MLD is carried in ICMPv6 protocol type packets

The IGMP process is enabled by default as soon as IP Multicast is enabled. IGMP packets also carry these protocols, and therefore all of these protocols are enabled whenever multicast is enabled:

- PIMv1 – PIMv1 was the first version of PIM and is always enabled in Cisco IOS for migration purposes. Current deployments all use PIMv2.
- Mrinfo - Mrinfo is a Unix command that Cisco IOS inherited to display multicast neighbors. Cisco recommends the use of SNMP instead of the mrinfo command.
- DVMRP - DVMRP is a legacy dense mode distance vector protocol with very limited scaling characteristics. Cisco IOS support for DVMRP is retired or already deprecated.
- Mtrace - Mtrace is the multicast equivalent of unicast “traceroute” and is a useful tool

For more information, See IANA's Internet Group Management Protocol (IGMP) Type Numbers

```
Router> mtrace 172.16.0.0 172.16.0.10 239.254.254.254

Type escape sequence to abort.
Mtrace from 172.16.0.0 to 172.16.0.10 via group 239.254.254.254
From source (?) to destination (?)
Querying full reverse path...
 0  172.16.0.10
-1  172.16.0.8 PIM  thresh^ 0  0 ms
-2  172.16.0.6 PIM  thresh^ 0  2 ms
-3  172.16.0.5 PIM  thresh^ 0  894 ms
-4  172.16.0.3 PIM  thresh^ 0  893 ms
-5  172.16.0.2 PIM  thresh^ 0  894 ms
-6  172.16.0.1 PIM  thresh^ 0  893 ms
```

Unicast IGMP packets (for IGMP/UDLR) can be filtered, as these are most likely attack packets and not valid IGMP protocol packets. Unicast IGMP packets are supported by Cisco IOS in support of unidirectional links and other exception conditions.

Forged IGMP/MLD query packets can result in a lower IGMP version than expected.

In particular, hosts ideally never send IGMP queries because a query sent with a lower IGMP version can cause all hosts that receive this query to revert to the lower version. In the presence of IGMPv3 / SSM hosts, this can “attack” the SSM streams. In the case of IGMPv2, this can result in longer leave latencies.

If a non-redundant LAN with a single IGMP querier is present, the router needs to drop IGMP queries received.

If a redundant / common passive LAN exists, then a switch capable of IGMP snooping is required. There are 2 specific features that can help in this case:

- Router Guard
- IGMP Minimum Version command

Router Guard

Any switch port can become a multicast router port if the switch receives a multicast router control packet (IGMP general query, PIM Hello, or CGMP Hello) on that port. When a switch port becomes a multicast router port, all multicast traffic is sent to that port. This can be prevented with “Router Guard”. The Router Guard feature does not require IGMP snooping to be enabled.

The Router Guard feature allows a specified port to be designated a multicast host port. The port cannot become a router port, even if multicast router control packets are received.

These packet types are discarded if they are received on a port that has Router Guard enabled:

- IGMP query messages
- IPv4 PIMv2 messages
- IGMP PIM messages (PIMv1)
- IGMP DVMRP messages
- Router-port Group Management Protocol (RGMP) messages
- Cisco Group Management Protocol (CGMP) messages

When these packets are discarded, statistics are updated which indicate that packets are dropped due to Router Guard.

IGMP Minimum Version

It is possible to configure the minimum version of IGMP hosts allowed. For example, you can disallow all IGMPv1 hosts or all IGMPv1 and IGMPv2 hosts. This filter applies only to membership reports.

If the hosts are attached to a common "passive" LAN (for example, a switch that does not support IGMP Snooping, or is not configured for it), there is also nothing a router can do about such false queries other than ignore the “old version” membership reports that are then triggered, and not fall back itself.

Since IGMP queries must be visible to all hosts, it is not possible to use a Hash-based message authentication (HMAC) mechanism with a pre-shared key, such as static key IPSec, to authenticate IGMP queries from “valid routers”. If two or more routers are attached to a common LAN segment, an IGMP querier election is required. In that case, the only filter that could be used is an ip access-group filter based on the source IP address of the other IGMP router that sends queries.

“Normal” multicast IGMP packets must be permitted.

This filter can be used on receiver ports to allow only “good” IGMP packets, and to filter known “bad” ones:

```
ip access-list extended igmp-control
   <snip>
   deny   igmp any any  pim         ! No PIMv1
   deny   igmp any any  dvmrp       ! No DVMRP packets
   deny   igmp any any  host-query ! Do not use this command with redundant routers.
                                    ! In that case this packet type is required !
   permit igmp any host 224.0.0.22 ! IGMPv3 membership reports
   permit igmp any any  14          ! Mtrace responses
   permit igmp any any  15          ! Mtrace queries
   permit igmp any 224.0.0.0 10.255.255.255 host-query  ! IGMPv1/v2/v3 queries
   permit igmp any 224.0.0.0 10.255.255.255 host-report ! IGMPv1/v2 reports
   permit igmp any 224.0.0.0 10.255.255.255 7           ! IGMPv2 leave messages
   deny   igmp any any              ! Implicitly deny unicast IGMP here!
<snip> permit ip any any ! Permit other packets interface ethernet 0 ip access-group igmp-
control in
```

> **Note**: This type of IGMP filter can be used in receive ACLs or CoPP. In both applications, it needs to be combined with filters for other traffic handled, such as routing and management plane protocols.

Fig 19: Host Receiver-Side Access Control

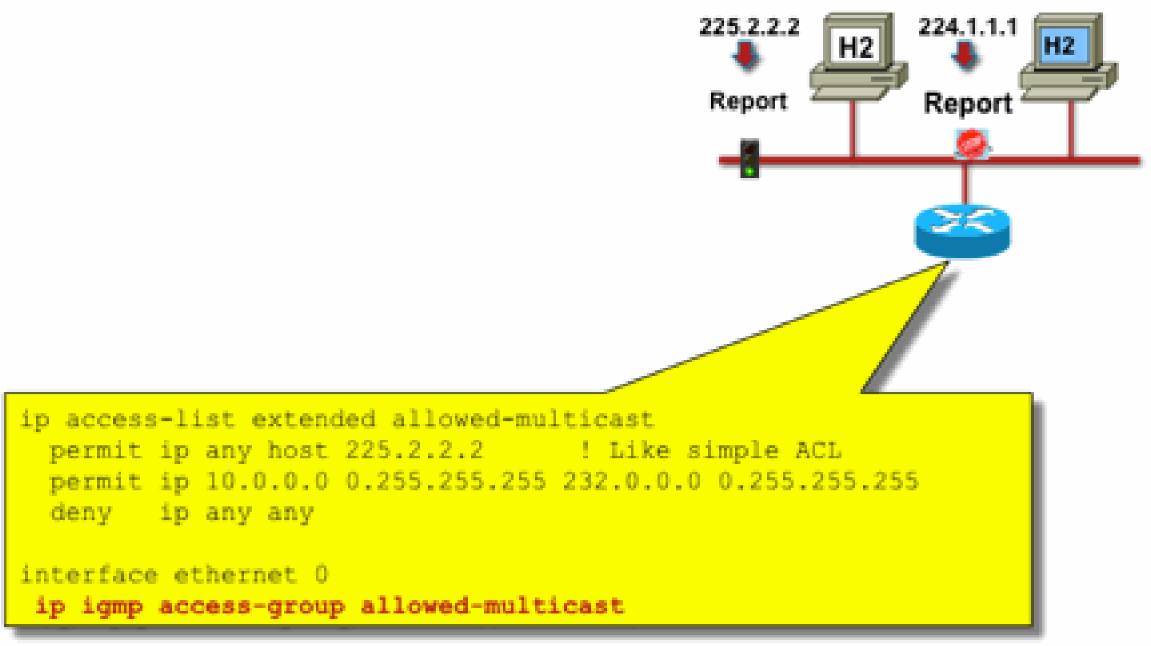

```
ip access-list extended allowed-multicast
  permit ip any host 225.2.2.2        ! Like simple ACL
  permit ip 10.0.0.0 0.255.255.255 232.0.0.0 0.255.255.255
  deny   ip any any

interface ethernet 0
 ip igmp access-group allowed-multicast
```

Fig19_Host_Receiver_Access

To filter traffic to a receiver do not filter data plane traffic, but rather the control plane protocol IGMP. Since IGMP is a necessary prerequisite to receive multicast traffic, no data plane filters are required.

In particular, you could restrict which multicast flows receivers can join (attached to the interface that the command is configured on). In this case, use the **ip igmp access-group / ipv6 mld access-group** command:

```
ip igmp access-group / ipv6 mld access-group
```

For ASM groups, this command only filters based on the destination address. The source IP address in the ACL is then ignored. For SSM groups that use IGMPv3 / MLDv2, it filters on source and destination IP.

This example filters a given group for all IGMP speakers:

```
access-list 1 deny 226.1.0.0 0.0.255.255
access-list 1 permit any log
! interface ethernet 1/3 ip igmp access-group 1
```

This example filters specific IGMP speakers (hence, specific multicast receivers) for a given group:

```
ip access-list extended test5
 deny igmp host 10.4.4.4 host 232.2.30.30
 permit igmp any any
!
interface Ethernet0/3
 ip igmp access-group test5
```

> **Note**: Remember for ASM groups, the source is ignored.

Admission Control

Access control delivers either a binary, yes or no answer for certain flows, independently of the state of the network. Admission control by contrast limits the number of resources that a sender / receiver can use, assume they passed the access control mechanisms. Various devices are available to help with admission control in a multicast environment.

Global and Per Interface IGMP Limits

At the router closest to interested multicast receivers, there is the possibility to limit the number of IGMP groups joined both globally and per interface. You can utilize the **ip igmp limit/ipv6 mld limit** commands:

```
ip igmp limit <n> [ except <ext-acl> ]
ipv6 mld limit <n> [ except <ext-acl> ]
```

It is recommended that this limit always be configured per interface and also globally. In each case, the limit refers to counts of entries in the IGMP cache.

The next two examples show how this command can be used to help limit the number of groups at the edge of a residential broadband network.

Example 1 - Restrict received groups to only the SDR announcements plus one received channel

Session Directory (SDR) acts as a channel guide to some muticast receivers. See RFC 2327 for more details.

A common requirement is to restrict receivers to receive the SD group plus one channel. This example configuration can be used:

```
ip access-list extended channel-guides
  permit ip any host 239.255.255.254 ! SDR announcements
  deny   ip any any

ip igmp limit 1 except channel-guides

interface ethernet 0
  ip igmp limit 2 except channel-guides
```

The access list in this example specifies the channel guide only; the global **ip igmp limit** command limits each IGMP source to a single (1) channel, but does not include the channel guide, which can always be received. The interface command overrides the global command and allows two (2) channels to be received, in addition to the channel guide, on this interface.

Example 2 - Admission Control on Aggregation-DSLAM Link

This command can also be used to provide a form of bandwidth admission control. For example, if it were necessary to distribute 300 SDTV channels, which are 4Mbps each, and there is a 1Gbps link to the Digital-Subscriber-Line-Access-Multiplexer (DSLAM), you can take a policy decision to limit the TV bandwidth to 500 Mbps and leave the rest for Internet and other uses. In that case, you can limit the IGMP states to 500 Mbps/4 Mbps = 125 IGMP states.

This configuration can be used in this case:

Fig 20: Use of Per-Interface IGMP Limits; Admission Control on Agg-DSLAM Link

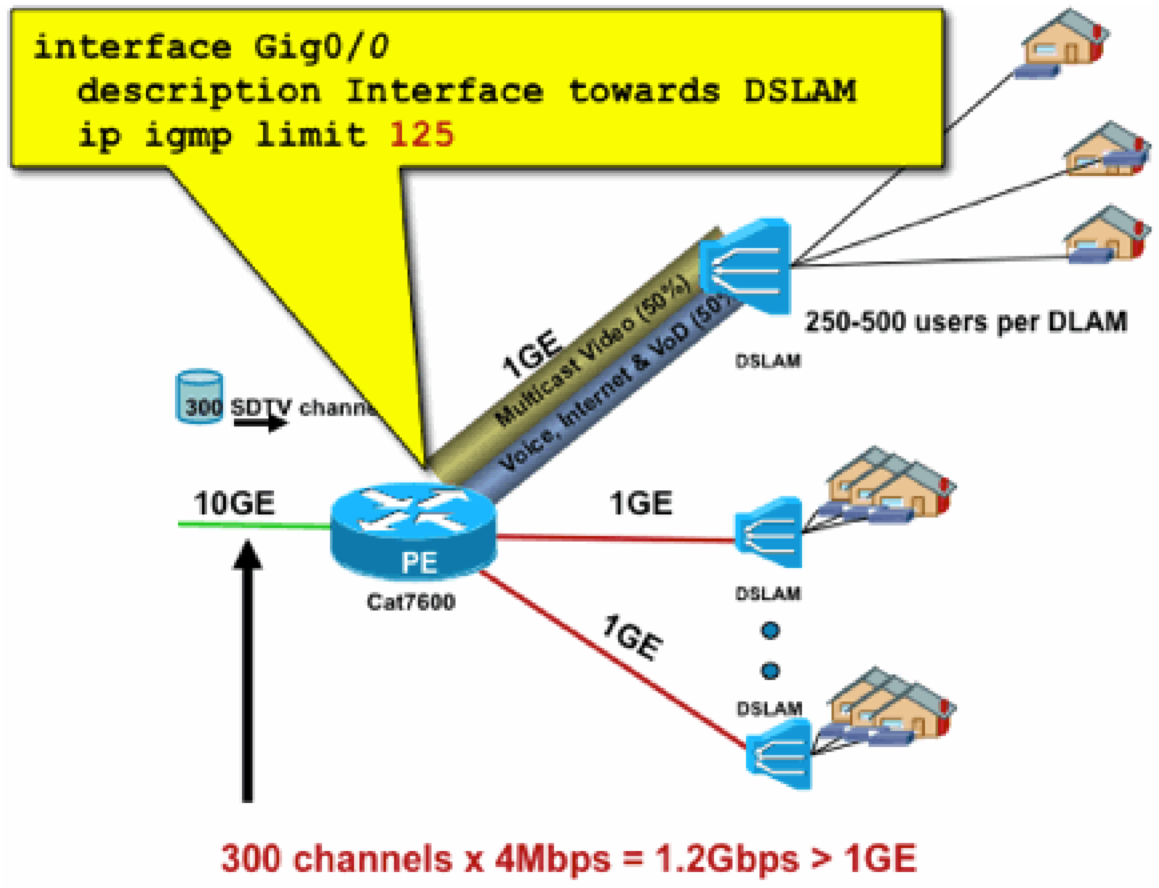

```
interface Gig0/0
  description Interface towards DSLAM
  ip igmp limit 125
```

Fig20_PerInterface_IGMP

Per-Interface mroute Limits

Enablement of per-interface mroute state limits is a more generic form of admission control. It not only limits IGMP and PIM state on an outgoing interface, but also provides a way of state limits on incoming interfaces.

Use the **ip multicast limit** command:

```
ip multicast limit [ rpf | out | connected ] <ext-acl> <max>
```

State can be separately limited on input and output interfaces. Directly attached source state can also be limited with the use of the “connected” key word. Examples illustrate the use of this command:

Example 1 – Egress Admission Control on Agg-DSLAM Link

In this example, there are 300 SD TV channels. Assume that each SD channel needs 4 Mbps, with a total of no more than 500 Mbps. Finally, also assume that there is a need for support Basic, Extended, and Premium bundles. Example bandwidth allocations:

- 60% / 300 Mbps Basic
- 20% / 100 Mbps Extended
- 20% / 100 Mbps Premium

Then use 4 Mbps per channel, limit the DSLAM uplink to:

- Basic 75 states
- Extended 25 states
- Premium 25 states

Configure the limit on the outbound interface which faces the DSLAM from the PEAgg:

Fig 21: Use of Per-Interface mroute Limits; Admission Control on Agg-DSLAM Link

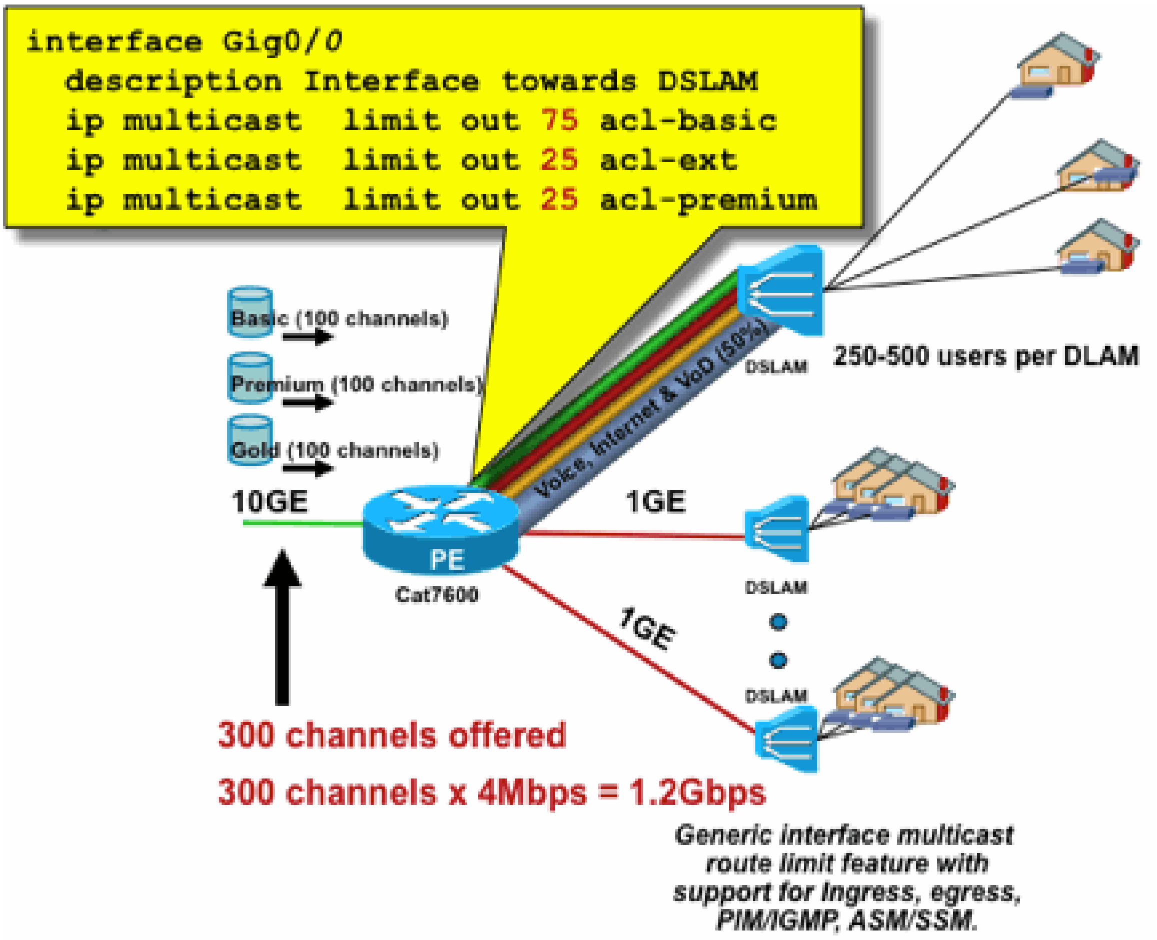

```
interface Gig0/0
  description Interface towards DSLAM
  ip multicast  limit out 75 acl-basic
  ip multicast  limit out 25 acl-ext
  ip multicast  limit out 25 acl-premium
```

Basic (100 channels)
Premium (100 channels)
Gold (100 channels)
10GE
PE
Cat7600
Voice, Internet & VoD (50%)
DSLAM
250-500 users per DLAM
1GE
DSLAM
1GE
DSLAM
300 channels offered
300 channels x 4Mbps = 1.2Gbps
Generic interface multicast route limit feature with support for Ingress, egress, PIM/IGMP, ASM/SSM.

Fig21_P

erInterface_Mroute

Example 2 – Ingress Admission Control on Agg-DSLAM Link

Instead of the “out” limit on the upstream device’s outbound interface, it is possible to use RPF limits on the downstream device’s RPF interface. This effectively has the same result as the previous example, and could be useful if the downstream device is not an Cisco IOS device.

Fig 22: Use of Per-Interface mroute Limits; Input Admission Control

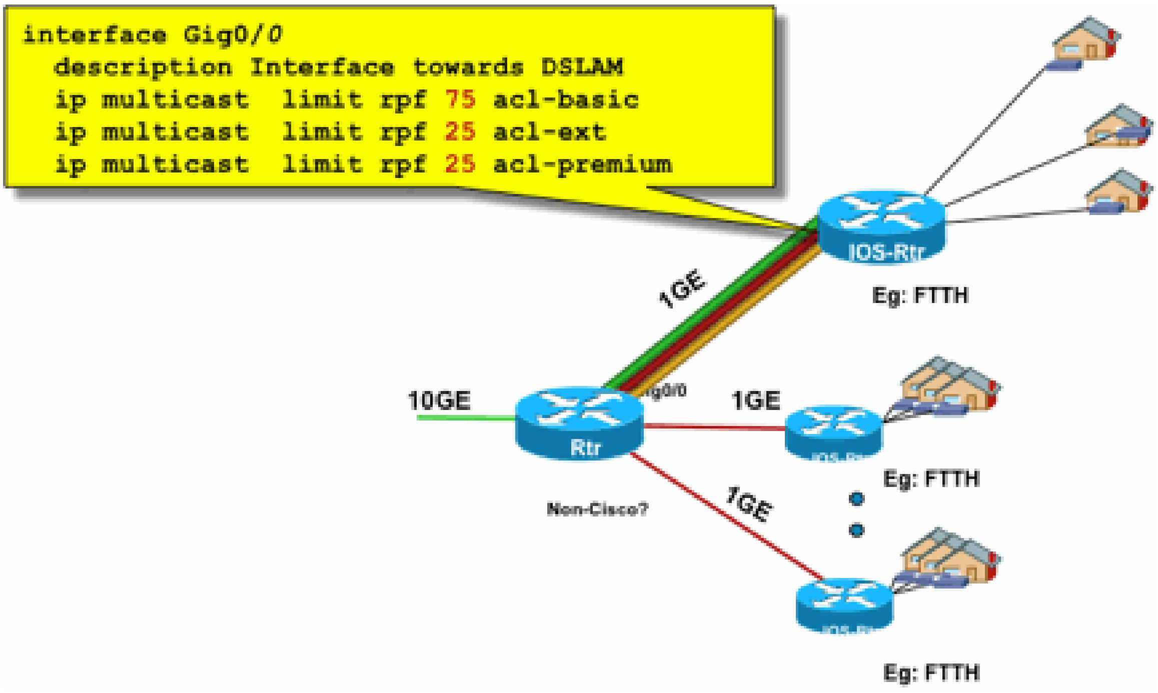

```
interface Gig0/0
  description Interface towards DSLAM
  ip multicast  limit rpf 75 acl-basic
  ip multicast  limit rpf 25 acl-ext
  ip multicast  limit rpf 25 acl-premium
```

Fig22_PerInterface_Mroute_inputControl

Example 3 - Bandwidth-Based limits

You can make a further subdivision of access bandwidth between multiple content providers and offer each content provider a fair share of the bandwidth on the uplink to the DSLAM. In that case, use the **ip multicast limit cost command**:

```
ip multicast limit cost <ext-acl> <multiplier>
```

With this command, it is possible to attribute a “cost” (use the value specified in “multiplier”) to any states that match the extended ACL in the ip multicast limit.

This command is a global command and multiple simultaneous costs can be configured.

In this example, it is necessary to support three different content providers with fair access to each into the network. Additionally, in this example it is a requirement to support Moving Picture Experts Group (MPEG) streams of various types:

MPEG2 SDTV: 4Mbps
MPEG2 HDTV: 18Mbps
MPEG4 SDTV: 1.6Mbps
MPEG4 HDTV: 6Mbps

In such a case you could allocate bandwidth costs to each stream type and share the remainder of the 750Mbps between the three content providers with this configuration:

```
ip multicast limit cost acl-MP2SD-channels   4000 ! from any provider ip multicast limit cost
acl-MP2HD-channels 18000 ! from any provider ip multicast limit cost acl-MP4SD-channels 1600 !
from any provider ip multicast limit cost acl-MP4HD-channels 6000 ! from any provider !
interface Gig0/0 description --- Interface towards DSLAM --- <snip> ! CAC ip multicast limit out
250000 acl-CP1-channels ip multicast limit out 250000 acl-CP2-channels ip multicast limit out
```

```
250000 acl-CP3-channels
```

Fig 23: Cost Factor for Per-Interface Mroute State Limits

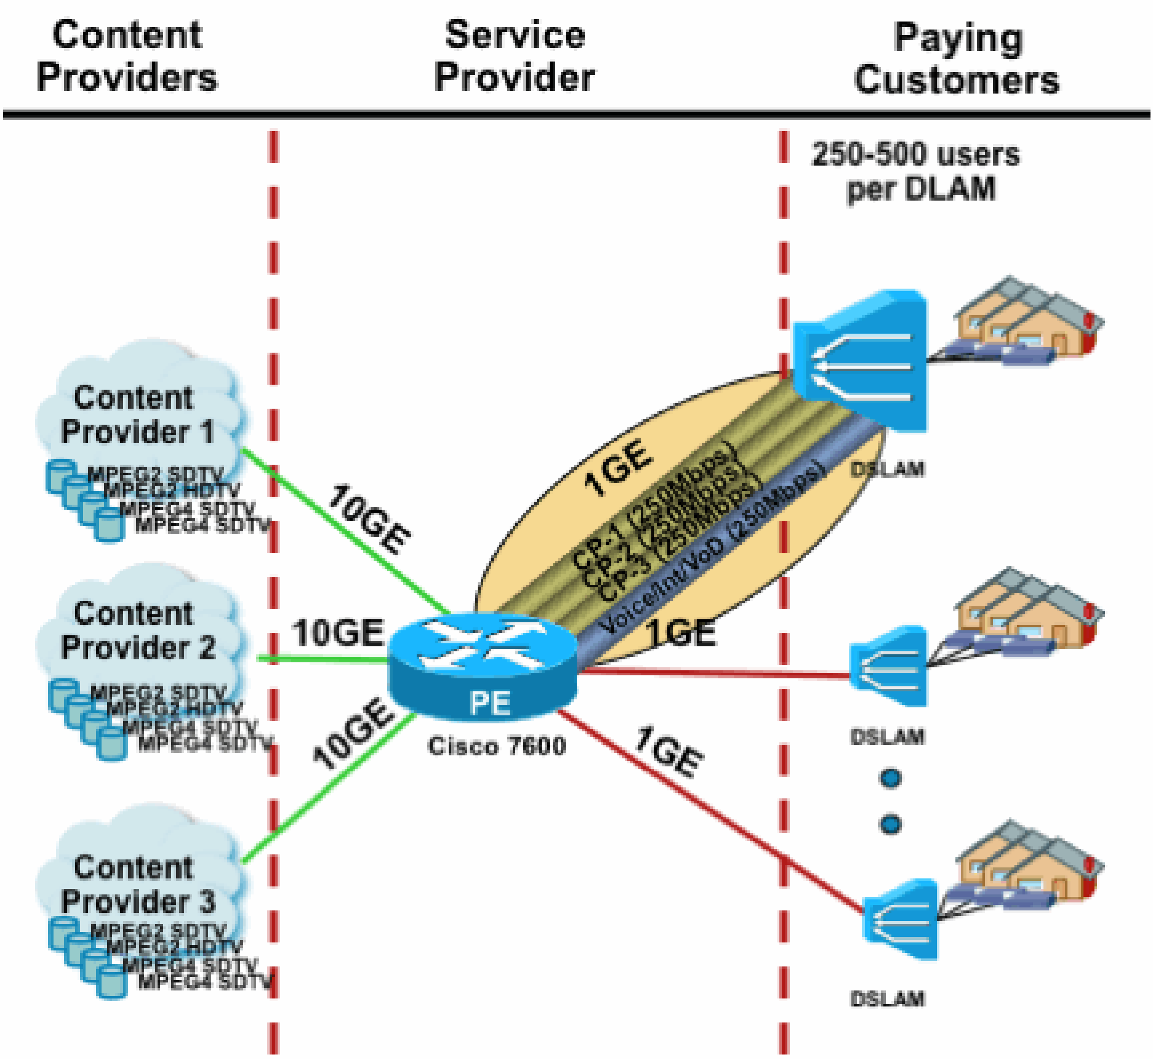

Fig23_Cost_ PerInterface

Multicast and IPSec

Introduction to GET VPN

As with unicast, multicast traffic also sometimes needs to be secured to provide confidentiality or integrity protection. There are two major areas where such services could be required:

- Encryption of multicast streams (for example in banking applications that stream confidential data to a large set of receivers that use multicast) - this is data plane security.
- Encryption of control plane protocols that use multicast, OSPF or PIM, for example - this is control plane security.

IPSec as a protocol [RFCs 6040, 7619, 4302, 4303, 5282] is specifically limited to unicast traffic (by RFC). There, a “security association” (SA) is established between two unicast peers. In order to apply IPSec to multicast traffic, one option is to encapsulate multicast traffic within a GRE tunnel and then to apply IPSec to the GRE tunnel, which is unicast. A newer approach uses a single security association established between all members of the group. The Group Domain of Interpretation (GDOI) [RFC 6407] defines how this is achieved.

Based on GDOI, Cisco developed a technology called Group Encryption Transport (GET) VPN. This technology uses “Tunnel Mode with Address Preservation,” as defined in the document “draft-ietf-msec-ipsec-extensions”. In GET VPN, first a group security association is established between all members of the group. Subsequently the traffic is protected, either with ESP (encapsulated security payload) or AH (authentication header), that uses tunnel mode with address preservation.

In summary, GET VPN encapsulates a multicast packet that uses the address information of the original header, and then protects the inner packet in relation to the group policy, with an ESP, for example.

The advantage of GET VPN is that multicast traffic is not affected at all by the security encapsulation mechanisms. The routed IP header addresses remain the same as the original IP header. Multicast traffic can be secured the same way with or without GET VPN.

The policy that is applied to the GET VPN nodes is centrally defined on a Group Key Server and distributed to all group nodes. Therefore, all group nodes have the same policy, and the same security settings applied to group traffic. Similar to standard IPSec, the crypto policy defines what type of traffic needs to be protected in which way. This allows GET VPN to be used for various purposes.

Use GET VPN to Encrypt Multicast Data Plane Traffic

The network-wide crypto policy is set on the group key server, and distributed to the GET VPN endpoints. The policy contains the IPSec policy (IPSec mode - here: tunnel mode with header preservation), and security algorithms to be used (for example AES). It also contains a policy that describes which traffic can be secured, as defined by an ACL.

GET VPN can be used for multicast as well as unicast traffic. A policy to secure unicast traffic could be defined by an ACL:

```
permit ip 10.0.0.0   0.255.255.255   10.0.0.0   0.255.255.255
```

This would encrypt all traffic with a source IP from 10/8 and a destination IP 10/8. All other traffic, for example, traffic from 10/8 to another address, would be ignored by GET VPN.

The application of GET VPN for multicast traffic is technically the same. For example, this access-control entry (ACE) can be used to secure traffic from any source to respective multicast groups:

```
permit ip any 239.192.0.0 0.0.255.255
```

This policy matches all sources (“any”) and all multicast groups that start with 239.192. Traffic to other multicast groups are not secured.

> **Note**: Great attention must be paid to the construction of the crypto ACL. Management traffic, or traffic that originates outside the GET VPN domain but terminates inside (that is traffic that passes only one crypto endpoint), must be excluded from the GDOI policy.

Common mistakes include:

- permit ip any 224.0.0.0 0.255.255.255: This also encrypts OSPF traffic and other control plane traffic, which is destined to a peer router, for example.
- The management traffic is not excluded from the crypto policy, which terminates inside the network. This includes GDOI traffic itself.

Use GET VPN to Authenticate Control Plane Traffic

It is generally a best practice to authenticate control plane traffic, such as routing protocols, to ensure that messages come from a trusted peer. This is comparatively simple for control plane protocols that use unicast, such as BGP. However, many control plane protocols use multicast traffic. Examples are OSPF, RIP, and PIM. See IANA's IPv4 Multicast Address Space Registry for the full list.

Some of these protocols have built-in authentication, like Routing Information Protocol (RIP) or Enhanced Interior Group Routing Protocol (EIGRP), others rely on IPSec to provide this authentication (for example OSPFv3, PIM). For the latter case, GET VPN provides a scalable way to secure these protocols. In most cases, the requirement is protocol message authentication, or in other words, verification that a message was sent by a trusted peer. However, GET VPN also allows encryption of such messages.

To secure (typically authenticate only) such control plane traffic, the traffic needs to be described with an ACL and included in the GET VPN policy. The details depend on the protocol to be secured, where attention needs to be paid to whether the ACL includes traffic that passes only an ingress GET VPN node (that is is encapsulated), or also an egress node.

There are two fundamental ways to secure PIM protocols:

- **permit ip any 224.0.0.13 0.0.0.0**: This is the “All PIM Routers” multicast group. However, this does not secure unicast PIM messages
- **permit pim any any**: This secures the PIM protocol, independent of whether multicast or unicast is used

> **Note**: The commands are provided as examples to help explain a concept. For example, it is necessary to exclude certain PIM protocols used to bootstrap PIM, such as BSR or Auto-RP. Noth methods have certain advantages and inconveniences which are dependent on the deployment. Please refer to specific literature on how to secure PIM with GET VPN for details.

Conclusions

Multicast is increasingly common service in networks. The emergence of IPTV services in residential / home broadband networks, and the move towards electronic trading applications in many of the world’s financial markets are just two examples of requirements that make multicast an absolute requirement. Multicast comes with a variety of different configuration, operation and management challenges. One of the key challenges is security.

This document examined a variety of ways in which multicast can be secured:

- First, look at the overall multicast control and data planes, an explanation of how the

differences from unicast present new security challenges.
- Next, an examination of the key protocols that are encountered in a multicastnetwork, in particular IGMP, PIM, and MSDP were examined in some detail. In each case, a description of security threats and recommended best practices for mitigation against these threats were provided.
- Additionally, some specific examples of how multicast can be secured in some specific applications, such as broadband edge networks where bandwidth can be limited in comparison to the amount of bandwidth that specific video flows could require.
- Finally, GET VPN architecture was described as a means of integrated multicast with IPSec for delivery of secure VPNs.

With multicast security in mind, remember how it is different to unicast. Multicast transmission is based upon the creation of dynamic state, multicast involves dynamic packet replication, and multicast builds unidirectional trees in response to PIM JOIN / PRUNE messages. Security of this whole environment involves the understanding and deployment of a rich framework of Cisco IOS commands. These commands are largely centered around protection of protocol operations, states (multicast), or policers placed against packets like CoPP. With correct use of these commands it is possible to provide a robust protected service for IP multicast.

In summary, there are multiple approaches that are promoted and described in this paper:

1. Widespread use of SSM – this is the most simple PIM mode that also allows the use of (S,G) forwarding.
2. If ASM services are needed, ensure a robust service can be provided – use of statically defined RPs provides a more secure control plane than dynamic RP announcements. Auto-RP and BSR are more flexible
3. If PIM-SM is enabled, look at areas of particular vulnerability, like the register tunnel to the RP, and ensure that the DR is always well protected. CoPP is very helpful in these areas.
4. If inter-domain ASM services are needed, consider whether BiDir PIM can be deployed.
5. Use global mroute/igmp state limits – understand the capabilities of your platforms together with the expected maximum amount of state you need under normal circumstances and in the worst case scenario. Configure limits within your platform’s capabilities that allow your network to operate to its maximal limits.
6. Fundamental filters – rACL/CoPP and infrastructure ACLs, which blocks PIM at the access layer

IP Multicast is an exciting and scalable means to deliver a variety of application services. Like unicast, it needs to be secured in a variety of different areas. This paper provides the basic building blocks that can be used to secure an IP multicast network.

Related Information

- Guidelines for Enterprise IP Multicast Address Allocation
- ConfigureIPv4 IGMP Filters
- Group Encrypted Transport VPN