MULTICAST SECURITY RECOMMENDATIONS

Introduction

IP multicast differs from IP unicast in many ways. One major difference between multicast and unicast is with security implications. The method by which data transfer is initiated and received is unique to Multicast. The possible distribution trees and their creation, also introduces a number of factors in how to secure Multicast.

In this document we will define the security considerations that relate to IP multicast and discuss possible methodologies to mitigate security threats. We will make generic recommendations based on related deployment experience. However we urge the reader to decide for themselves what the requirements are for their own deployment as it is impossible to create a 'one size fits all' recommendation on such a wide topic.

This is not an introduction to IP Multicast. If you require information on the details of IP Multicast, and its related protocols, please use the reference section at the end of this document.

Definitions

Multicast security can and frequently does mean different things to many people. In order to discuss this topic we need to define terms of reference. We can separate the subject into four main areas: Access Control, Admission Control, Control Plane security and Data Plane security. We describe below the recommended methodologies for securing each area.

Access control

This defines the **permissibility** of a device to **transmit** and **receive** multicast data.

Access control can be divided into two areas:

- Source constraint
- Receiver constraint

Receivers use IGMP to signal interest in a particular group or channel, source's have no equivalent protocol.

Source Constraint

Source constraint requires the use of the data plane as well as the control plane. We can use ACLs and, with PIM SM, filter at the RP. MSDP is used to share information across PIM SM domains and we can also filter MSDP SA messages. We can also apply filters at configured multicast boundaries.

Receiver Constraint

It is possible use extended ACLs to define what data can be forwarded by an interface. There are also protocol specific commands for IGMP which operate on a per interface basis. If the interface is also the boundary to a PIM domain it's possible to constrain data via the Multicast Boundary command. Used with TACACS, Radius, etc we can make these commands dynamic based on a user profile

Admission control

This defines the ability to **transmit** or **receive** based on available **network resources.**

```
We can utilize both RSVP and QOS to ensure that the network can provide
the resources required to service a multicast stream.
```

Control plane security

This defines the ability to secure the **protocols** used in the **distribution** and **reception** of Multicast data

It is possible to use extended ACLs to define certain sources for each protocol a router will be permitted to process. There are also a number of protocol specific commands such as the PIM Neighbor Filters, IGMP access group, RP announce filter (Auto RP specific), MSDP filters both in and out, Multicast Boundary and BSR Border. Cisco also

keeps an up to date list of recommended MSDP SA filters available online. See references.

Data Plane security

This defines the ability to secure the **content** of the Multicast transmission.

This normally requires an application to encrypt the data to ensure only validated/permitted users have access to the content.

Threat on multicast case study: The Sasser Worm

Internet worms (or other DOS type attacks) occasionally spread throughout the internet. While mostly targeting unicast addresses, they do occasionally target multicast addresses as well. The most recent worm began propagating on vulnerable Microsoft windows systems on April 29th, 2004 and is dubbed the Sasser worm, due to the exploitation of the LSASS vulnerability. To propagate, this worm sends a specially-crafted packet to TCP port 445 with random IP destination addresses, **including multicast addresses**. The packet causes a buffer overrun on vulnerable systems, which results in the execution of a remote shell that opens port 9996. This worm commands the remote shell to download its copy from the original infected source via port 5554 using FTP. Because it will scan multicast addresses, there have been reports that some routers which route multicast traffic have become unstable as a result of Sasser infections.

The recommended method to help stop the effects of this particular Sasser worm is to **filter ICMP and TCP packets destined to multicast packets. There is no use for the reception of multicast packets within TCP**. MSDP uses TCP but MSDP uses the TCP unicast endpoints for communication. **The only need for ICMP packets destined for a multicast address is for testing multicast through PING**. **ICMP packets destined for multicast addresses can therefore also be filtered. Cisco routers will automatically drop packets with a multicast source address.** This is required per IETF multicast RFCs. The syntax to deny ICMP and TCP packets destined to multicast addresses is as follows:

```
access-list 115 deny icmp any 224.0.0.0 15.255.255.255
access-list 115 deny tcp any 224.0.0.0 15.255.255.255
```

Specific Commands to Assist in Securing a Multicast Network

Below are brief portions of the command reference for security related commands. For the complete information on each command please refer to the detailed information online.

ip dvmrp accept-filter

This command configures an acceptance filter for incoming DVMRP Reports. Any destinations that match <access-list> received in DVMRP reports from neighbors are stored in the DVMRP routing table with <distance>. The distance is used to compare with the same destination in the unicast routing table. The lower distance route (either from the unicast routing table or DVMRP routing table) will take precedence when computing the RPF interface for a source of a multicast packet. When no filters are configured on an interface, all destinations are accepted with distance configured from the "ip dvmrp distance" global command. An <access-list> value of 0, accepts all destinations. <access-list> can be a simple or extended access-list. An neighbor-list value of 0 means accept from all neighbors on interface.

The following example will permit all incoming DVMRP reports to be installed in the DVMRP routing table with a distance of 100:

```
ip dvmrp accept-filter 1 100
access-list 1 permit 0.0.0.0 255.255.255.255
```

First released in 10.0. Neighbor filter added in 11.2

ip dvmrp route-limit

The following example changes the limit to 5000 DVMRP routes allowed to be advertised:

```
ip dvmrp route-limit 5000
```

First released in 11.0

ip multicast route-limit

This command limits the total number of (*,G) and (S,G) multicast routing table entries as shown in "show ip mroute" to <routes>. Use this command to limit the impact of Denial of Service attacks based on creating useless IP multicast routing state. Valid arguments are 1... 2,147,483,646. "no ip multicast route-limit" establishes a multicast route-limit of 2,147,483,647 (the maximum 32-bit integer value). This is also the default configuration and it will not show up in the configuration.

If the router needs to create a new multicast routing table entry but has exceeded the number of configured <routes>, a warning level log message will be emitted:

"<current-routes> routes exceeded multicast route-limit of <routes>"

<current-routes> can be larger than <routes> if you configured "ip multicast route-limit <routes>" when the router already had more routes than <routes> installed. In that case the router will not remove already existing routes (you can force deletion of routes with "clear ip mroute"). The currently configured value <routes> is also displayed in the "show ip mroute count" command.

The following example shows how to limit the total number of multicast routing table entries to 1,000:

```
ip multicast route-limit 1000
```

First release in: 12.1(2), 12.0(11)S

ip igmp limit (global)

The following example shows how to limit the number of IGMP states on a router to 300:

```
ip igmp limit 300
```

First released in 12.2(15)T

ip igmp limit (interface)

The following example shows how to limit the number of IGMP membership reports on Ethernet interface 0:

```
interface ethernet 0
 ip igmp limit 100
```

First released in 12.2(15)T

ip igmp access-group

In the following example, hosts serviced by Ethernet interface 0 can join the group 225.2.2.2 only:

```
access-list 1 225.2.2.2 0.0.0.0
interface ethernet 0
 ip igmp access-group 1
```

First released in 10.0 Extended ACL added in 12.3(7)T

ip msdp sa-limit

The following example configures the SA message limit to 100 for the MSDP peer with IP address 172.16.0.0:

```
ip msdp sa-limit 172.16.0.0 100
```

First released in 12.1(7)

ip msdp sa-filter in

The following example configures the router to filter all SA messages from the peer named router.cisco.com:

```
ip msdp peer router.cisco.com connect-source ethernet 0
ip msdp sa-filter in router.cisco.com
```

First released in 12.0(7)T with VRF Keyword added in 12.0(23)S & 12.2(13)T.

ip msdp sa-filter out

The following example allows only (S, G) pairs that pass access list 100 to be forwarded in an SA message to the peer named router.cisco.com:

```
ip msdp peer router.cisco.com connect-source ethernet 0
ip msdp sa-filter out router.cisco.com list 100
access-list 100 permit ip 224.69.0.0 0.0.255.255 224.2.0.0 0.0.255.255
```

First released in 12.0(7)T with VRF Keyword added in 12.0(23)S & 12.2(13)T

ip multicast rate-limit

In the following example, packets to any group from sources in network 172.16.0.0 will have their packets rate-limited to 64 kbps:

```
interface serial 0
 ip multicast rate-limit out group-list 1 source-list 2 64
access-list 1 permit 0.0.0.0 255.255.255.255
access-list 2 permit 172.16.0.0 0.0.255.255
```

First released in 11.0

ip multicast boundary

The following example sets up a boundary for all administratively scoped addresses:

```
access-list 1 deny 239.0.0.0 0.255.255.255
access-list 1 permit 224.0.0.0 15.255.255.255
interface ethernet 0
 ip multicast boundary 1
```

First released in 11.1 with Filter Auto-RP Keyword added in 12.0(22)S, 12.1(12c)E, 12.2(11) & 12.2(13)T

ip pim accept-register

The following example shows how to restrict the RP from allowing sources in the Source Specific Multicast (SSM) range of addresses to register with the RP. These statements need to be configured only on the RP.

```
ip pim accept-register list no-ssm-range

ip access-list extended no-ssm-range
 deny ip any 232.0.0.0 0.255.255.255
 permit ip any any
```

First released in 12.0(7)T with VRF Keyword added in 12.0(23)S & 12.2(13)T

ip pim accept-rp

The following example states that the router will accept join or prune messages destined for the RP at address 172.17.1.1 for the multicast group 224.2.2.2:

```
ip pim accept-rp 172.17.1.1 3
```

```
access-list 3 permit 224.2.2.2
```

First released in 12.0(7)T with VRF Keyword added in 12.0(23)S & 12.2(13)T

ip pim border / ip pim bsr-border

The following example configures the interface to be the PIM domain border:

```
interface ethernet 1
ip pim bsr-border
```

First released in 11.3T command became bsr-border 12.0(8)

ip pim neighbor-filter

The following example enables stub multicast routing on Router A, which has an outgoing interface with IP address 10.0.0.1. Router B is a central router with an incoming interface with address 10.0.0.2. Access list 1 filters PIM messages from the source (stub Router A).

Router A Configuration

```
ip multicast-routing
ip pim dense-mode
ip igmp helper-address 10.0.0.2
```

Router B Configuration

```
ip multicast-routing
 ip pim dense-mode : or ip pim sparse-mode
 ip pim neighbor-filter 1
access-list 1 deny 10.0.0.1
```

First released in 11.3

ip pim register-rate-limit

The following example shows how to configure the ‘ip pim register-rate-limit’ command with a maximum rate of two register messages per second:

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

First released in 11.3T with VRF Keyword added in 12.0(23)S & 12.2(13)T

ip pim rp-announce-filter

The following example configures the Mapping Agent router(s) to accept RP announcements from candidate RPs in access list 1 for group ranges described in access list 2:

```
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 224.0.0.0 192.168.255.255
```

First released in 11.1 with VRF Keyword added in 12.0(23)S & 12.2(13)T

References

Introductory Material

http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_sol/mcst_ovr.htm
http://www.cisco.com/en/US/tech/tk828/tk363/tech_brief09186a00800e9952.html
http://www.cisco.com/warp/public/732/Tech/multicast/multicast_preso.shtml

IOS Command Reference Vol3 IP Multicast

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_command_reference_book09186a00801a7ec5.html

IGMP state Limit
http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a00801541c2.html

MSDP filters
ftp://ftpeng/ipmulticast/config-notes/msdp-sa-filter.txt