Mboned T. Chown
Internet-Draft University of Southampton
Intended status: Informational L. Giuliano
Expires: September 13, 2012 Juniper Networks
March 12, 2012

Multicast Filtering Practices
draft-chown-mboned-multicast-filtering-02

Abstract

Operators of multicast networks may apply various filters to multicast traffic at boundary routers or on MSDP peerings. The aim of this text is to discuss appropriate filtering policies, as well as documenting existing filtering practices, with a view to generating some discussion towards producing guidance on best filtering practice.

Table of Contents

1. Introduction

Multicast filtering can be applied at a multicast boundary or on an MSDP peering as a means to prevent unintended leakage of multicast traffic beyond its desired scope. An informal discussion of filtering practices suggested that those practices vary from organisation to organisation. The aim of this text is to gather and document commonly used existing filtering practices. Whether it is then possible to draw up a definitive best practice is to be determined; it is quite possible that due to the shifting nature of the target that a point-in-time recommendation would quickly be overtaken by events. For example, the recent addition of unicast prefix-based IPv4 multicast addresses [RFC6034] meant that filtering of all of 234.0.0.0/8 became undesirable. However, general principles may remain valid over time.

The text begins with a discussion of appropriate policies in Section 2, followed by Section 3 in which we document a summary of reported practices at border, organisational and subnet scopes. We then draw some conclusions in Section 3.

For sites on academic research networks, some examples of filtering recommendations already exists, e.g. in documentation [I2multicast] from the Internet2 Multicast WG, and in the JANET IPv4 Multicast Guide [JANETmulticast]. There is also a more specific proposal for the Rutgers network [RutgersProposal], which includes a good discussion of organisational-local scope address usage within its network as a whole.

When determining filtering policies, one needs to consider how strict to be; some ranges are not supposed to be used, but there may be no harm per se in accepting them. There are certainly some ranges that should not be filtered, such as the newly assigned 234.0.0.0/8 range mentioned above, and the GLOP range under 233.0.0.0/8.

An additional resource is the registry of IPv4 multicast address space held by IANA [IANAmulticast]. This registry should be a definitive guide to the formal use of ranges of addresses within the overall IPv4 multicast address space. A similar registry is maintained for IPv6 multicast address space [IANA6multicast].

This text is still quite an early draft, aimed at soliciting feedback, both on content and whether the goal of the draft is actually worthwhile. Different sites may have different requirements. There may also be issues with handling scope boundaries that need to be considered. So there may be general principles that could be captured in a document such as this, even if specific filtering rules are not included.

2. General Discussion of Policies

In this section we discuss how we might believe filters to be applied with respect to various multicast protocol functions. This may include rate limiting and policing in addition to straight filtering. In the following section we summarise actual filtering practices that have been reported.

2.1. Multicast Addressing and Domain Borders

To date, IPv4 multicast addresses have been assigned from the following ranges for global usage: 224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 232/8, 233/8 and 234/8. All other IPv4 multicast addresses are therefore reserved, unassigned or scoped, and as such, have no legitimate reason for use on the Internet. Therefore we recommend operational filters that permit these address ranges and block all others at domain borders.

It should be noted that some networks do use multicast addresses outside of these ranges for internal purposes. For example, 239/8, which is administratively scoped for Organization-Local usage, is functionally analogous to RFC 1918 unicast IPv4 addresses, and is often used by networks to support internal customers or infrastructure services. As such, these types of addresses may be used within a domain, but should never be allowed to cross domain borders.

2.2. MSDP SA Policy and Policers

MSDP policies should include filters that allow SAs only from the assigned ASM IPv4 multicast address ranges: 224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 233/8 and 234/8. The 232/8 address range is reserved for SSM only, and thus, should never appear in MSDP.

The most common multicast attack vector that has appeared on the Internet has been MSDP SA storms. In nearly all cases, these were triggered by Internet worms that probed large blocks of addresses in an attempt to scan vulnerable ports. Typically, these worms were intended to scan only unicast addresses at random, however, the worm coders accidentally included multicast addresses in the random pools of destinations to scan. When port scans with destination addresses of multicast addresses occur on multicast-enabled network, these packets generate PIM register messages and, consequently, MSDP SA messages according to standard PIM and MSDP procedures, resulting in a flood of SA messages across the Internet that can put great strain on MSDP-speaking routers.

MSDP filters that allow SAs from only the assigned ranges do help to reduce the potential pool for these accidental attacks. However, the addition of MSDP policers provides much stronger protection from SA storms. MSDP policers can be applied on a per-peer and per-source basis. Per-peer policers are used to detect when the number of SAs received from an MSDP peer exceeds a certain configured threshold. These policers are functionally analogous to BGP max prefix limits applied on BGP peers, which alert an operator when the number of routes received from a BGP peer exceeds a certain configured threshold. Typically, this type of per-peer threshold is determined by observing the number of SAs that are normally advertised by a peer and then selecting some multiple of that to be a limit. For example, if a peer normally advertises 2k SAs, an operator may set the per-peer threshold to 5k.

The challenge with per-peer limits is that they are not granular enough to determine good SAs from bad ones. As such, these can actually lower the bar needed to launch a denial of service attack. For example, with a 5k per-peer threshold, an attacker need only generate 5k SAs to trigger the limit and potentially block all the legitimate SAs from propagating. As such, operators may want to use per-peer limits to trigger an alarm, rather than tear down the MSDP session. Additionally, per-source MSDP policers can be used to provide further granularity between good and bad SAs. Per-source SA policers define the maximum number of SAs that can be permitted from the same source (or source range). For example, an operator can use a per-source limit of 1k, which would prevent any source from generating more than 1k MSDP SAs. With per-source limits in place, a heavily distributed attack would be necessary to generate enough MSDP state to impact routers as no single source can generate enough state.

2.3. Multicast Scoping

Multicast scoping is used to block multicast packets based on group address. Scoping filters should be used on all domain borders to allow only the assigned IPv4 multicast addresses (224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 232/8, 233/8 and 234/8) and block all other multicast groups from ingress/egress these borders.

2.4. PIM Policy

While multicast scoping blocks traffic in the data plane, it does not affect the control plane. Thus, an illegitimate group could be joined, with traffic carried across a providers network only to be dropped at the border by a scoping filter. However, with PIM join filters, the join could be dropped at the border so that no state is created in the first place. Therefore, PIM join filters can be used

for additional protection on all domain borders to allow joins for only the assigned IPv4 multicast addresses (224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 232/8, 233/8 and 234/8) and block joins for all other multicast groups. Additionally, PIM join filters can be used to block joins based on the source address of the multicast source (ie, the S in the (S,G) join). So PIM join filters can also be configured to prevent joins for illegitimate interdomain sources, such as RFC 1918 addresses.

PIM register filters can also be used to protect an RP from creating register state for illegitimate groups by allowing PIM registers on an RP for only the assigned IPv4 ASM multicast addresses (224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 233/8 and 234/8) and blocking registers for all other multicast groups. PIM register filters can also be configured to prevent registers for illegitimate interdomain sources, such as RFC 1918 addresses.

2.5. IGMP Policy

As with PIM join filters, IGMP filters can be used for additional protection on all receiver LANs to allow IGMP reports for only the assigned IPv4 multicast addresses (224.0/16, 224.1/16, 224.2/16, 224.3/16, 224.4/16, 232/8, 233/8 and 234/8) and block reports for all other multicast groups. IGMP filters can also be configured to prevent IGMPv3 source-included reports for illegitimate interdomain sources, such as RFC 1918 addresses.

2.6. SAP

In the early years of multicast, routers were often configured to listen to the SAP group (224.2.127.254) and cache the SDP messages. This was used as a quick and dirty way of determining if multicast was working properly. Just by looking at the SAP/SDP cache on a router and guessing if the number of session entries looked about right, an operator could quickly determine if multicast flows were traversing the router. However, this imprecise method of management and troubleshooting eventually proved dangerous, as improperly formatted SDP messages occasionally crashed routers as well as improperly configured SAP sources accidently transmitted the multicast stream they intended to announce on the SAP group, causing harm to routers that cached this data. As such, routers should not be configured to listen to the SAP group and cache SDP messages.

3. Summary of Reported Filtering Practices

3.1. Border and MSDP Filtering

In this section we summarise IPv4 multicast addresses that are commonly filtered at site borders or on MSDP peerings. Based on responses we received from a couple of multicast community lists, it wasn't clear which filters are applied on border routers and which on MSDP SA messages. Some sites apply minimal traffic filters, but heavier MSDP filtering.

A site may choose to filter on addresses or on observed TTLs; it is now general practice to filter on addresses rather than the TTL filtering that was common a long time ago.

Some sites choose to route multicast around their unicast firewalls, for performance or other operational reasons, but this shouldn't alter the requirement to filter groups appropriately where necessary.

In this section we draw on the small number contributions so far; we hope to get more inputs in time. In general, many 224.0.0.* addresses that are used by infrastructure are typically blocked, as well as some addresses that are global scope but should not be, like Ghostcast.

The following list includes multicast IPv4 addresses that are being filtered based on the union of responses received so far (hence the apparent duplication of certain prefixes). The list of filters applied by all respondents would be somewhat shorter.

```
224.0.1.1       NTP
224.0.1.2       SGI-Dogfight
224.0.1.3       Rwhod
224.0.1.8       SUN NIS+
224.0.1.20      any private experiment
224.0.1.22      SVRLOC
224.0.1.24      microsoft-ds
224.0.1.25      nbc-pro
224.0.1.35      SVRLOC-DA
224.0.1.38      Retrospect
224.0.1.39      cisco-rp-announce
224.0.1.40      cisco-rp-discovery
224.0.1.41      gatekeeper
224.0.1.60      hp-device-disc
224.0.1.65      iapp
224.0.1.76      IAPP lucaent-avaya-ap
224.0.2.1       rwho
224.0.2.2       SUN RPC
224.0.2.3       EPSON-disc-set
224.0.23.1      Ricoh-device-ctrl
224.0.23.2      Ricoh-device-ctrl
224.1.0.1       Cisco Aironet
224.1.0.38      Retrospect
224.2.0.2       Altiris Rapideploy
224.2.0.3       Altiris Rapideploy
224.77.0.0/16   Norton Ghost
224.101.101.101 Sun Sunray
225.1.2.3       Altiris Development Server and Deployment Agent
226.77.0.0/16   Norton Ghost
229.55.150.208  Norton Ghost
231.0.0.0/8     ?
234.21.81.1     Limewire
234.42.42.0/30  ImageCast
234.42.42.32/31 ImageCast
234.42.42.40/30 ImageCast
234.142.142.42/31       ImageCast
234.142.142.44/30       ImageCast
234.142.142.48/28       ImageCast
234.142.142.64/26       ImageCast
234.142.142.128/29      ImageCast
234.142.142.136/30      ImageCast
234.142.142.140/31      ImageCast
234.142.142.142 ImageCast
239.0.0.0/8     Scoped groups
239.252.0.0/14  Scoped groups
239.234.5.6     ECopy ShareScan
```

One site gave figures for matches/hits on its filters; it may be

interesting to gather such statistics at other sites.

Different networks make different use of the scoped address space under 239.0.0.0/8, which may lead to different organisational filters in different scenarios. Organisation-local scope IPv4 multicast addressing is described in [RFC2365].

The SSM range 232.0.0.0/8 should not be carried in MSDP peerings; this is an example of different policy applied at the site border to an MSDP peering. Usually the filters are probably the same though.

As a general principle, multicast sourced from private address ranges [RFC1918] or from 169.254.0.0/16, 192.0.2.0/24 or 127.0.0.0/8 should be dropped, regardless of the multicast destination.

In certain cases, rate limiting may be desirable, where complete filtering might not, e.g. in mitigating against SAP [RFC2974] storms, or against unintended MSDP SA bursts.

Where BSR is deployed, a site should consider dropping BSR packets at its border, both BSR messages and C-RP messages. Except for Embedded-RP it probably makes sense to drop PIM register messages at the site border, unless a site's RP is external.

3.2. Organisational filtering

As described in [RutgersProposal], a site may use multiple organisational scopes within its site, which may use different blocks from 239.0.0.0/8, and thus require appropriate filtering at boundaries, e.g. between metropolitan campuses.

3.3. Subnet filtering

Two respondents are currently filtering uPNP between subnets, and one is filtering mDNS. One reason for the uPNP filtering was due to issues with errant Ricoh printers which flood announcements with too-large TTLs.

Subnet filtering may help protect against other forms of misconfigured client subnets. One site has networks that consist of multiple edge routers, where the outside 'LAN side' is strictly meant to be for local subnets only, and all intranet comms are to go through the 'WAN interface'. They have had cases where multihomed client networks were misconfigured, cross-connecting IP subnets with layer 2 boxes. To prevent multiplication of multicasts, they configure all edge routers in the intranet to accept multicast packets from the 'LAN side' only if the source IP of the multicast packets belongs to the IP subnet of that LAN. So it's a simple

filter, with no scaling issues.

At least one site is filtering multicast traffic from its wireless links; this is presumably streamed video or audio content. Multicast support is required on wireless links for IPv6 operation. At layer 2, multicast IPv6 Router Advertisements may be filtered on ports that do not have known routers attached.

One site is running per-subnet boundary filters on its wired multicast-enabled subnets. The list below reflects these. One could add local scope relative addresses, though in practice the latter would all fall under 239.255.0.0/16 if it is the smallest scope group applied to a subnet.

```
224.0.1.1       NTP
224.0.1.2       SGI-Dogfight
224.0.1.3       Rwhod
224.0.1.8       SUN NIS+
224.0.1.24      microsoft-ds
224.0.1.25      nbc-pro
224.0.1.60      hp-device-disc
224.0.1.76      IAPP
224.0.2.1       rwho
224.0.2.2       SUN RPC
234.21.81.1     Limewire
239.255.0.0/16  subnet scope
```

The importance of such subnet filtering may depend on TTLs used.

4. Conclusions

This document discusses filters and related mechanisms that should be applied in multicast deployments, and summarises the reported use of such multicast filtering practices in the wild. It includes discussion of various multicast protocols and of reported practices of applying filters at various scope boundaries. The next iteation of this draft will include more detailed conclusions.

Further feedback on the text, and the practices reported to date is welcomed.

5. Security Considerations

There are no extra security consideration for this document.

6. IANA Considerations

There are no extra IANA consideration for this document.

7. Acknowledgments

The author would like to thank the following people for their contributions to this text: Scott Bertilson, Alan Buxey, Bruce Curtis, Andy Gatward, Bob Gerdes, Jeffry J. Handal, Lonnie Leger, Bert Manfredi, Garry Peirce, William F. Maton Sotomayor, and Stig Venaas.

8. Informative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.

[RFC2974] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, October 2000.

[RFC6034] Thaler, D., "Unicast-Prefix-Based IPv4 Multicast Addresses", RFC 6034, October 2010.

[JANETmulticast] Price, D., "IPv4 Multicast on JANET", 2006, <http://www.ja.net/documents/publications/technical-guides/ipv4-multicast-web.pdf>.

[IANA6multicast] "IPv6 Multicast Address Space Registry", <http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xml>.

[IANAmulticast] "IPv4 Multicast Address Space Registry", <http://www.iana.org/assignments/multicast-addresses/multicast-addresses.xml>.

[RutgersProposal] "iDRAFT Proposal for RUNet Administratively Scoped Multicast", <http://www.td.rutgers.edu/documentation/Papers/Administratively_Scoped_Multicast_Proposal.pdf>.

[I2multicast]
"Enabling IP Multicast with Internet2", <http://noc.net.internet2.edu/i2network/multicast-cookbook.html>.

Authors' Addresses

Tim Chown
University of Southampton
Highfield
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

Leonard Giuliano
Juniper Networks

Email: lenny@juniper.net