Internet Engineering Task Force
Internet Draft
draft-ietf-mboned-mzap-00.txt
December 15, 1997
Expires: June 1998

MboneD WG
M. Handley
ISI

Multicast-Scope Zone Announcement Protocol (MZAP)

STATUS OF THIS MEMO

ABSTRACT

This document defines a protocol, the Multicast-Scope Zone Announcement Protocol (MZAP), for discovering the multicast administrative scope zones that are relevant at a particular location. MZAP also provides mechanisms whereby two common misconfigurations of administrative scope zones can be discovered.

1 Status

This is a strawman proposal. It has not been subject to any peer review or implementation.

2 Introduction

IP Multicast groups can be global scope, or they can be restricted in scope using a scoping mechanism. In this document, we only consider administrative scoping, as defined in [1]. An administrative scope zone is defined by a set of border routers to a region of the network. These border routers are configured to not pass multicast traffic destined for a particular range of multicast addresses to or from links leaving the scope zone.

Administrative scope zones may be of any size, and a particular host may be within many administrative scope zones. The only zones a host can assume that it is within are the global zone, and local scope Local scope is defined as being the smallest administrative scope zone encompassing a host, and the border is configured for addresses in the range 239.255.0.0 to 239.255.255.255 inclusive. [1] specifies: 239.255.0.0/16 is defined to be the IPv4 Local Scope. The Local Scope is the minimal enclosing scope, and hence is not further divisible. Although the exact extent of a Local Scope is site dependent, locally scoped regions must obey certain topological constraints. In particular, a Local Scope must not span any other scope boundary. Further, a Local Scope must be completely contained within or equal to any larger scope. In the event that scope regions overlap in area, the area of overlap must be in its own local scope. This implies that any scope boundary is also a boundary for the Local Scope.

Two problems make administrative scoping difficult to deploy and difficult to use:

- o Misconfiguration is easy. It is difficult to detect scope zones that have been configured so as to not be convex (the shortest path between two nodes within the zone passes outside the zone) or to leak (one or more border routers was not configured correctly).

- o Applications have no way to discover the scope zones that are relevant to them. This makes it difficult to use admin scope zones, and hence reduced the incentive to deploy them.

This document defines the Multicast Scope Zone Announcement Protocol (MZAP) which will provide applications with information about the scope zones they are within, and also provide diagnostic information to detect misconfigured scope zones.

3 Specification

A multicast scope Zone Border Router (ZBR) is a router that is configured to be a zone border on one or more of its interfaces. Any interface that is configured to be a border for any admin scope zone MUST also be a border for the local scope zone, as defined in [1].

Routers SHOULD be configured so as the router itself is within the scope zone. This is should in figure 1A, where router 1 is inside the scope zone and has the border configuration. It is possible for the first router outside the scope zone to be configured with the border, as illustrated in figure 1B where routers 2 and 3 are outside the zone and have the border configuration, but this is NOT RECOMMENDED.

```
         ............                     ................
        .            .   +O+-->          .                *O+-->
       .              . /  2            .                /. 2
      .              1 *               .              1 +  .
      .          <---+O*---+O+->       .          <---+O+---*O+->
      .              + .     3         .              +     . 3
      .             /  .               .             /      .
       . zone X  <--  .                 . zone X  <--      .
        ..............                   ..................

       O - router       * - border interface      + - interface

     A. Correct zone border           B. Incorrect zone border
```

Figure 1: Correct admin scope zone border placement

This rule does not apply for local scope borders, but applies for all other admin scope border routers.

When a ZBR router is configured correctly, it can deduce which side of the boundary is inside the scope zone and which side is outside it. It can also send messages into the scope zone, which it SHOULD NOT be able to do if the router itself is considered outside the scope zone.

Such a ZBR router should then send periodic Zone Announcement Messages (ZAMs) for the zone for which it is configured as a border from each of its interfaces that go into that scope zone. These messages are multicast to the address 239.255.255.254, which is a locally scoped address.

Each ZBR router should also listen for ZAM messages from other ZBRs for the same border. The ZBR router with the lowest interface IP address within the zone from those ZBRs forming the zone border becomes the zone-id router for the zone. The combination of this IP address and the base address of the scoping range server to uniquely identify the scope zone.

Every local scope ZBR that receives any ZAM for a scope zone other than local scope SHOULD then forward the ZAM out of all the other interfaces that are in different local scope zones except ones that form a border for the zone described in the ZAM. It adds the zone-id of the local scope zone that the message came from to the ZAM path list before doing so. A local scope ZBR receiving a ZAM with a non-null path list MUST NOT forward that ZAM back into a local scope zone that is contained in the path list. This process is illustrated in figure 2.

in

Figure 2: ZAM Message Flooding

The packet also contains a Zones Traveled Limit. If the number of Local Zone IDs in the ZAM path becomes equal to the Zones Traveled Limit, the packet should be dropped. Zones Traveled Limit is set when the packet is first sent, and defaults to 32, but can be set to a lower value if a network administrator knows the expected size of the zone.

Addition messages called Zone Convexity Messages (ZCM) SHOULD also be sent to the second highest address in the scope zone range itself (For example, if the scope zone border is for 239.1.0.0 to 239.1.0.255, then these messages should be sent to 239.1.0.254.) As these are not locally scoped packets, they are simply multicast across the scope zone itself, and require no path to be built up, or forwarding by local scope zone ZBRs. These messages are used to detect non-convex admin scope zones, as illustrated in figure 3. Here Router B and Router C originates ZCM messages, each reporting each other's presence. Router D cannot see Router B's messages, but sees Router C's report of Router B, and so concludes the zone is not convex.

in

Figure 3: ZAM Message Flooding

3.1 Packet Formats

Zone Announcement Message (IPv4)

The format of a Zone Announcement Message is shown in figure 4.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | V=0 |R| PTYPE |      ZT       |     ZTL       | IP  |  MLEN   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                       Zone Base Address                       |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Message Origin                         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Zone ID Address                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                     Local Zone ID Address 0                   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Router Address 0                       |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                .....
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                     Local Zone ID Address N                   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Router Address N                       |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                           Zone Name                           |
   |                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Zone Announcement Message Format

The fields are defined as follows:

Version (V): 3 bits The version defined in this document is version 0.

Respond (R): 1 bit When set to 1, this bit indicates that a router MAY generate a Zone Limit Exceeded message in response to this ZAM message. When set to zero, a router MUST NOT generate a Zone Limit Exceeded message in response to this message.

Packet Type (PTYPE): 4 bits A Zone Announcement Message has PTYPE=0.

Zone Traveled (ZT): 8 bits This gives the number of Local Zone IDs contained in this message path.

Zones Traveled Limit (ZTL): 8 bits This gives the limit on number of local zones that the packet can traverse before it MUST be

dropped.

IP Protocol Version (IP): 3 bits This indicates the format of the following packet. The following values are defined:

0: IPv4

1: IPv6

Mask length (MLEN): 5 bits This gives the mask length which together with the zone base address defines the range of addresses that form the border to this zone. For example, if the zone is a border for 239.1.0.0 to 239.1.0.255, then MLEN has the value 24. A value of zero means all the bits are included in the mask, and the zone is a border for a single address.

Zone Base Address: 32 bits This gives the base address for the scope zone border. For example, if the zone is a border for 239.1.0.0 to 239.1.0.255, then Zone Base Address is 239.1.0.0.

Message Origin: 32 bits This gives the IP address of the interface that originated the ZAM message.

Zone ID Address: 32 bits This gives the lowest IP address that has been observed in the zone sending ZAM messages. Together with Zone Base Address and MLEN it forms a unique ID for the zone.

Zone Path: multiple of 64 bits The zone path is a list of 32 bit Local Zone ID Addresses (the Zone ID Address of a local zone) through which the ZAM message has passed, and 32 bit IP addresses of the router that forwarded the packet. Every local scope router that forwards the ZAM across a local scope boundary MUST add the Local Zone ID Address of the local zone that the packet was received from and its own IP address to the end of this list, and increments ZT accordingly. The zone path is empty which the ZAM message is first sent.

Zone Name: multiple of 8 bits The Zone Name is an ISO 10646 character string in UTF-8 encoding indicating the name given to the scope zone (eg: "ISI-West Site"). It should be relatively short and MUST be less that 256 bytes in length. All the border routers to the same region SHOULD be configured to give the same Zone Name, or a zero length string MAY be given. A zero length string is taken to mean that another router is expected to be configured with the zone name. Having ALL the ZBR routers for a scope zone announce zero length names should be considered an error.

3.1.1 Zone Limit Exceeded (ZLE)

This packet is sent by a local-zone border router that would have exceeded the Zone Traveled Limit if it had forwarded a ZAM packet. It is only sent if the "Respond" bit in the ZAM packet is set, and it is unicast to the Message Origin given in the ZAM packet.

The format is the same as a ZAM message, and is shown in figure 5:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  V=0  | PTYPE |      ZT       |     ZTL       | IP  |  MLEN   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Zone Base Address                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Message Origin                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Zone ID Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Local Zone ID Address 0                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                             .....
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Local Zone ID Address N                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Zone Name                           |
|                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Zone Announcement Message Format

All fields are copied from the ZAM message, except PTYPE which is set to one.

A router receiving ZLE messages SHOULD log them and attempt to alert the network administrator that the scope zone is misconfigured.

Zone Convexity Message

Unlike Zone Announcement Messages which are sent to the locally scoped address 239.255.255.254, Zone Convexity Messages are sent to the second highest address in the scope zone itself. The format of a

ZCM is shown in figure 6 and is similar to a ZAM, expect PTYPE take the value two.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  V=0  | PTYPE |     ZNUM      |    unused     | IP  |  MLEN   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                       Zone Base Address                       |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Message Origin                         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                        Zone ID Address                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                          ZBR Address 0                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                .....
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                          ZBR Address N                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                           Zone Name                           |
   |                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: Zone Convexity Message Format

The fields are as follows:

Number of ZBR addresses (ZNUM): 8 bits This field gives the number of ZBR Addresses contained in this message.

ZBR Address (32 bits) These fields give the addresses of ALL the other ZBR routers that the Message Origin ZBR has received ZCM messages from during the time that it has taken the Message Origin ZBR to send ten ZCM messages.

4 Using Zone Announcement Messages

Any host or application may listen to Zone Announcement Messages to build up a list of the scope zones that are relevant locally. However, listening to Zone Announcement Messages is not the recommended method for regular applications to discover this

information. These applications will normally query a local Multicast Address Allocation Server[2], which in turn will listen to Zone Announcement Messages.

A ZBR can discover misconfiguration of scope boundaries in one of several ways:

- o It receives ZLE messages indicating that the scope zone is leaking.

- o It receives ZAM messages originating inside the scope boundary on an interface that points outside the zone boundary. Such a ZAM message must have escaped the zone through a leak, and flooded back around behind the boundary. This is illustrated in figure 7.

- o Other ZBR routers in the scope zone are announcing that they are seeing a different set of routers than our router observes from arriving ZCM messages. If this is a persistent condition, it indicates that the scope zone is probably not convex, as illustrated in figure 3.

in

Figure 7: ZAM Message Leaking

All these conditions should be considered errors and the router should attempt to alert the network administrator to the nature of the misconfiguration.

Zone Convexity Messages can also be sent and received by correctly configured ordinary hosts within a scope region, which may be a useful diagnostic facility that does not require privileged access.

5 Message Timing

Each ZBR router should send a Zone Announcement Message for each scope zone for which it is a boundary every ZAM-INTERVAL seconds, +/- 30% of ZAM-INTERVAL each time to avoid message synchronisation.

Each ZBR router should send a Zone Convexity Message for each scope zone for which it is a boundary every ZCM-INTERVAL seconds, +/- 30% of ZCM-INTERVAL each time to avoid message synchronisation.

A router SHOULD NOT send more than one Zone Limit Exceeded message every ZLE-MIN-INTERVAL regardless of destination.

Default values are:

ZAM-INTERVAL 300 seconds.

ZCM-INTERVAL 300 seconds.

ZLE-MIN-INTERVAL 300 seconds.

6 Security Considerations

MZAP does not include authentication in its messages. Thus it is open to misbehaving hosts sending spoof ZAM or ZCM messages.

In the case of ZCM messages, these spoof messages can cause false logging of convexity problems. It is likely that is would be purely an annoyance, and not cause any significant problem.

In the case of ZAM messages, spoof messages can also cause false logging of configuration problems. This is also considered to not be a significant problem.

Spoof zone announcements however might cause applications to believe that a scope zone exists when it does not. If these were believed, then applications may choose to use this non-existent admin scope zone for their uses. Such applications would be able to communicate successfully, but would be unaware that their traffic may be traveling further than they expected. As a result, applications MUST only take scope names as a guideline, and SHOULD assume that their traffic sent to non-local scope zones might travel anywhere. The confidentiality of such traffic CANNOT be assumed from the fact that it was sent to a scoped address that was discovered using MZAP.

7 Bibliography

[1] D. Meyer, "Administratively Scoped IP Multicast", Internet Draft, draft-ietf-mboned-admin-ip-space-04.txt [2] M. Handley, D. Thaler, D. Estrin, "The Internet Multicast Address Allocation Architecture", Internet Draft, Dec 1997.