

Internet Draft
Expires: August 2001

Rick Boivie
Nancy Feldman
IBM Watson Research Center

February 2001

Small Group Multicast
<[draft-boivie-sgm-02.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

Multicast has become increasingly important with the emergence of network-based applications such as IP telephony and video conferencing. The Internet community has done a significant amount of work on IP multicast over the last decade [1-10] and as a result, there are a number of multicast applications that are used today on the Mbone, the multicast-capable virtual network that is layered on top of (portions of) the Internet [10]. However, while today's multicast schemes are scaleable in the sense that they can support very large multicast groups, there are scalability issues when a network needs to support a very large number of distinct multicast groups. This document describes a new scheme for multicast that complements the existing schemes. Whereas the existing schemes can support a limited number of very large multicast groups, the scheme described here can support a very

large number of small multicast groups.

1. Introduction

Multicast, the ability to efficiently send data to a group of destinations, is becoming increasingly important for applications such as IP telephony, video-conferencing, video distribution, collaborative work environments, multiparty networked games, etc.

There seem to be two kinds of multicasts that are important: a broadcast-like multicast that sends data to a very large number of destinations and a "narrowcast" multicast that sends data to a fairly small group. An example of the first is the audio & video multicasting of a working group session from an IETF meeting to sites all around the world. An example of the second is a videoconference involving 3 or 4 parties. For reasons described below, it seems prudent to use different mechanisms for these two cases. As the reliable multicast transport group has stated: "it is believed that a 'one size fits all' protocol will be unable to meet the requirements of all applications" [[11](#)].

1.1. Current Multicast Schemes

Current multicast schemes [[1,3,4,6,8](#)] were designed to handle very large multicast groups. These work well if one is trying to distribute broadcast-like channels all around the world but they have scalability problems when there is a very large number of groups.

In some of these schemes, the nodes in the network build a multicast distribution tree for each <source, multicast group> pair and they disseminate this multicast routing information to places where it isn't necessarily needed, which leads to scaling problems if there are a large number of multicast groups.

Some other schemes try to limit the amount of multicast routing information that needs to be disseminated, processed and stored throughout the network. These schemes use a "shared distribution tree" that is shared by all the members of a multicast group and they try to limit the distribution of multicast routing information to just those nodes that "really need it". But these schemes also have problems. Because of the shared tree, they use less than optimal paths in routing packets to their destinations and they tend to concentrate traffic in small portions of a network. They also require that all of the routers in a multicast tree "signal", process and store multicast routing information. And they require that multicast routing information for the

various multicast groups be communicated across inter-AS administrative boundaries. These requirements cause scalability

problems and increase administrative complexity if there are a large number of multicast groups.

2. Small Group Multicast (SGM)

2.1. Overview

The "Small Group Multicast" (SGM) scheme described here attempts to eliminate these problems for the case of small groups. In SGM, the source node keeps track of the destinations that it wants to send packets to, and creates packet headers that contain the list of destination addresses. SGM-capable routers receive these packets, parse the SGM headers and use the ordinary unicast route table to determine how to route the packet to each destination. This eliminates the need for class D addresses, as well as the need for network routers to store state for the various multicast groups. This makes SGM very scaleable in terms of the number of groups that can be supported since the nodes in the network do not need to disseminate or store any multicast routing information for these groups. And since it doesn't use any multicast routing protocol, there are no inter-AS multicast routing "peering" issues to contend with. SGM has the additional benefit that packets always take the "right" path as determined by the ordinary unicast route protocols. Unlike the shared tree schemes, SGM minimizes network latency and maximizes network "efficiency". This removes some important obstacles that have, to this point, prevented the widespread acceptance and adoption of multicast. Thus, SGM makes multicast practical for very large numbers of small groups, which as suggested above is a very important case. Note that while SGM is not suitable for large multicast groups, such as the broadcast of an IETF meeting, it does provide an important complement to existing multicast schemes in that it can support very large numbers of small groups.

SGM takes advantage of one of the fundamental tenets of the Internet "philosophy", namely that one should move complexity to the edges of the network and keep the middle of the network simple. This is the principle that guided the design of IP and TCP and it's the principle that has made the incredible growth of the Internet possible. For example, one reason that the Internet has been able to scale so well is that the routers in the core of the network deal with large CIDR blocks as opposed to individual hosts or individual "connections". The routers in the core don't need to keep track of the individual TCP connections that are passing through them. Similarly, the IETF's diffserv effort is based on the idea that the routers shouldn't have to keep track of a large

number of individual RSVP flows that might be passing through them. It's the authors' belief that the routers in the core shouldn't have to keep track of a large number of individual

multicast flows either.

2.2. Mechanism

In Small Group Multicast, the source node keeps track of the destinations that it wants to send packets to. The source encodes a list of unicast destinations in an SGM header that follows the L3 header, and then sends the packet to a router. Each router along the way parses the SGM header, partitions the destinations based on each destination's next hop, and forwards an appropriate SGM packet to each of the next hops. Each "final" hop removes the SGM encoding, and forwards the data as a standard unicast packet to its destination.

For example, suppose that A is trying to get packets distributed to B, C & D in figure 1 below:




Figure 1

This is accomplished as follows: A sends an SGM packet to its default router, R1, that includes the list of destinations for the packet. Ignoring some details, the packet that A sends to R1 might look like this:

Level 3 header:

```

< dest = R1 >
< src = A >
< protocol = small group multicast > (new protocol type)

```

Level "3.5" header:

```

< dest = B C D >
< src = A >

```

followed by the payload that A wants delivered to B, C and D.

(Note that the SGM layer is said to be at "level 3.5" since it's between IP and UDP in the IP protocol stack. SGM is carried within

an IP datagram and it carries a UDP payload.)

When R1 receives this packet, it needs to properly process the SGM header. The processing that a router does on receiving one of these small group multicast packets is as follows:

- o Perform a route table lookup to determine the next hop for each of the destinations listed in the packet.
- o Partition the set of destinations based on their next hops.
- o Replicate the packet so that there's one copy of the packet for each of the next hops found in the previous steps.
- o Modify the list of destinations in each of the copies so that the list in the copy for a given next hop includes just the destinations that ought to be routed through that next hop.
- o Send the modified copies of the packet on to the next hops.
- o Optimization: If there is only one destination for a particular SGM next hop, send the packet as a standard unicast packet to the destination, as there is no multicast gain by formatting it as an SGM packet.

So, in the example above, R1 will send a single packet on to R2 with a destination list of < B C D > and R2 will send a single packet to R3 with the same destination list.

When R3 receives the packet, it will, by the algorithm above, send one copy of the packet to destination R5 with an SGM list of < C D >, and one ordinary unicast packet addressed to < B >. R4 will receive a standard unicast packet and forward it on to < B >. R5 will forward the SGM packet that it receives on to R6 which will pass it on to R7. When the packet reaches R7, R7 will transmit ordinary unicast packets addressed to < C > and < D > respectively. R8 and R9 will receive standard unicast packets, and forward the packets on to < C > and < D > respectively.

It's important that the SGM packet that is sent to a given next hop only includes destinations for which that next hop is the next hop listed in the route table. If the list of destinations in the packet sent to R4, for example, had also included C and D, R4 would send extra packets on to those nodes on a less than optimum path. This could waste a lot of bandwidth if one is, for example, multicasting a videoconference. And this could cause serious problems when route loops occur since a multicast packet could "spray" large numbers of packets in a number of different directions as it travels around a loop. Since the SGM packet that is sent to a given next hop only includes the destinations that are supposed to be reached through that next hop, these problems are eliminated.

Note that when routing topology changes, the routing for a

multicast flow will automatically adapt to the new topology since the path a multicast packet takes to a given destination always follows the ordinary, unicast routing for that destination.

Note also that since the SGM header is added to the data portion of the packet, if the sender wishes to avoid IP fragmentation, it must take the size of the SGM header into account.

See the Appendix for the SGM encoding formats.

3. Interoperability with Today's Routers

In the description above all of the routers in the network were "SGM-capable". But the scheme can also be used in an environment that includes routers that have no knowledge of SGM. This is important since it allows SGM to be used without upgrading all the routers in a network. There are two methods that can be used: the first uses SGM "tunnels"; the second, icmp unreachable messages.

3.1. SGM Tunnels

One way to deploy SGM in a network that has routers that have no knowledge of SGM is to setup "tunnels" between SGM peers. This enables the creation of a virtual network layered on top of an existing network. The SGM routers exchange and maintain SGM routing information via any standard unicast routing protocol (e.g. RIP, OSPF, ISIS). The SGM routing table that is created is simply a standard unicast routing table that contains the destinations that have SGM connectivity, along with their corresponding SGM next hops. In this way, packets may be forwarded hop-by-hop to other SGM routers, or may be "tunneled" through non-SGM routers in the network.

For example, suppose that A is trying to get packets distributed to B, C & D in figure 2 below, where "S" routers are SGM-capable, and "R" routers are not. Figure 3 shows the routing tables created via the SGM tunnels:




Figure 2

Router S1 establishes a tunnel to SGM peer S3.

Router S3 establishes a tunnel to SGM peers S1 and S7.
 Router S7 establishes a tunnel to SGM peer S3.

S1 routing table:	S3 routing table:	S7 routing table:
Dest NextHop	Dest NextHop	Dest NextHop
-----+-----	-----+-----	-----+-----
B S3	A S1	A S3
C S3	C S7	B S3
D S3	D S7	

Figure 3

The source A will send an SGM packet to its default SGM router, S1, that includes the list of destinations for the packet. As described in [section 2.2](#), the packet that A sends to S1 might look like this:

Level 3 header:

```
< dest = S1 >
< src = A >
< protocol = small group multicast > (new protocol type)
```

Level "3.5" header:

```
< dest = B C D >
< src = A >
```

followed by the payload that A wants delivered to B, C and D.

When S1 receives this packet it needs to properly process the multicast. The processing that this router does on receiving one of these small group multicast packets is as follows:

- o Perform a route table lookup in the SGM routing table to determine the SGM next hop for each of the destinations listed in the packet.
- o If no SGM next hop is found, replicate the packet and send a standard unicast to the destination.
- o For those destinations for which an SGM next hop is found, partition destinations based on their next hops.
- o Replicate the packet so that there's one copy of the packet for each of the SGM next hops found in the previous steps.
- o Modify the list of destinations in each of the copies so that the list in the copy for a given next hop includes just the destinations that ought to be routed through that next hop.
- o Send the modified copies of the packet on to the next hops.
- o Optimization: If there is only one destination for a particular SGM next hop, send the packet as a standard unicast packet to the

destination, as there is no multicast gain by formatting it as an SGM packet.

So, in the example above, S1 will send a single packet on to S3 with a destination list of < B C D >. This packet will be received by R2 as a unicast packet with destination S3, and R2 will forward it on, having no knowledge of SGM. When S3 receives the packet, it will, by the algorithm above, send one copy of the packet to destination < B > as an ordinary unicast packet, and 1 copy of the packet to S7 with a destination list of < C D >. R4, R5, and R6 will behave as standard routers with no knowledge of SGM. When S7 receives the packet, it will parse the packet and transmit ordinary unicast packets addressed to < C > and < D > respectively.

By using tunnels, packets follow the "best possible" path to the various destinations (as determined by the unicast routing protocols), while at the same minimizing the total number of packets that must be transmitted on the network.

3.2. ICMP

A second method for the gradual deployment of SGM in IP networks is based on the use of ICMP. In this case, forwarding follows as described in figure 1 (see [section 2.2](#)). However, an assumption is made that any router receiving an SGM packet that does not understand this new protocol will send an icmp message back to the source. This is a router requirement as specified in [RFC1812](#), "Requirements for IP Version 4 Routers" [12]. [Section 5.2.7.1 of RFC1812](#) says that a router should send an ICMP Destination Unreachable message with a code of 2, signifying Protocol Unreachable, if the transport protocol designated in a datagram is not supported in the transport layer of the final destination.

Thus, a router that doesn't understand this new protocol should send an icmp "destination unreachable, protocol unreachable" message back to the source. (If the icmp message is lost for some reason, a subsequent small group multicast packet will cause another icmp to be sent.) This enables the source to know when a router doesn't understand the new protocol. Furthermore, since the icmp message will include the initial part of the original packet, the source will also know the destinations that are not reachable via SGM, so the source can use unicast packets to reach those destinations. When routing topology changes, additional icmp "destination unreachable, protocol unreachable" messages may be generated and the source may use unicasts for additional destinations. The source can also periodically send an SGM packet to the destinations that are on its "unicast list" i.e. the list of nodes that it is reaching via unicast. Destinations that become

reachable via SGM (i.e. those do not appear in subsequent icmp "destination unreachable, protocol unreachable" messages) can then be removed from the unicast list. (Note that another possibility

would be to send an SGM "ping" periodically to the set of destinations and then use unicast to reach those destinations that don't respond to the ping).

Thus, Small Group Multicast can perform some multicasting in an environment that includes "legacy" routers which do not understand SGM. It won't work particularly well if there are many routers that don't understand SGM but this backwards compatibility may be important since it makes some of the benefits of multicast possible before all the routers in a network have been upgraded. This can be very useful since it may take some time to upgrade all the routers in a large network.

The use of the above method requires the bitmap style SGM packet encoding (see [Appendix A.1.1](#)).

4. Host Considerations

4.1. Control Plane

Unlike today's multicast schemes, SGM has no "control plane". There is no IGMP, and as mentioned above, there are no intradomain or interdomain multicast routing protocols or source discovery protocols. With SGM, the means by which multicast groups are defined is an application level issue and applications are not confined to the model in which hosts use IGMP to join a multicast group.

This allows applications such as voice and multimedia conferencing to be programmed in a very natural way, i.e. the user can place the call to the parties he or she wants to talk to as he or she would using the conferencing buttons on his or her telephone. The application developer is not limited to the receiver-initiated joins of the IGMP model.

Although there is no SGM "control plane", several control planes have been defined for the establishment of voice and multimedia conferences over IP networks, including SIP[13] and H.323[14].

These protocols use two mechanisms in establishing sessions among multiple participants: multicast using today's IP multicast schemes and "multi-unicasting". In "multi-unicasting", the application keeps track of the participants' unicast addresses and sends a unicast to each of those addresses. For reasons described in the introductory section of this paper, multi-unicasting rather than multicast is the prevalent solution in use today.

It's a simple matter to replace multi-unicast code with SGM
multicast. All that the developer has to do is replace a loop that

sends a unicast to each of the participants by a single "sgm_send" that sends the data to the participants. Thus it's a simple matter to incorporate SGM into real conferencing applications.

It's also worth mentioning that although SGM doesn't depend on a receiver-initiated join, the receiver-initiated join model can be implemented, if desired, by introducing a server that hosts can talk to join a conference.

4.2. SGM API

As mentioned earlier, each of the final destinations of an SGM packet receives an ordinary UDP packet. This means that hosts can receive an SGM multicast with a standard, unmodified TCP/IP stack. Hosts can also transmit SGM packets with a standard TCP/IP stack with a small SGM library that sends SGM packets on a raw socket. This has been used to implement SGM based applications on both Unix and Windows platforms without any kernel changes.

Another possibility is to modify the sockets interface slightly. For example, one might add an "sgm_sendto" function that works like sendto but that uses a list of destination addresses in place of the single address that sendto uses.

5. Summary

In summary, the disadvantages of Small Group Multicast are:

- o the extra bytes that are sent in a multicast packet for the list of destinations.
- o the need to use unicast packets in some cases to reach destinations that are behind "legacy" routers.
- o the need for a new IP stack in hosts (unless raw sockets are available in which case SGM packets can be sent over a raw socket).
- o the fact that SGM is not suitable for huge broadcast-like multicasts. It's targeted for "small" conferences.

The key advantages are:

- o it's very scaleable; it can handle a very large number of small groups.
- o the work involved is limited to just the nodes that are on the multicast tree.
- o no per flow state information is stored on the routers.
- o no multicast route protocol messages are communicated or processed; no intra-AS or inter-AS route protocols.
- o minimum administrative complexity; no need for complicated inter-AS peering agreements; it's just as easy for a network

administrator to support multicast as it is to support unicast,
and it will be just as easy to support multicast across the

Internet as it is to support unicast.

- o traffic follows the correct paths; traffic is not concentrated in a small part of the network; minimum network latency; maximum network "efficiency".
- o no need for class D addresses which means:
 - no need for a server that hands out class D addresses which can be a bottleneck or a point of failure.
 - no one can join the class D group and "eavesdrop" on the class D address; the source knows who he's sending to.
- o SGM can be easily adapted to provide a "reliable multicast".

The advantages of Small Group Multicast suggest that this scheme can be a very useful complement to the existing multicast schemes. Whereas the existing schemes can support a limited number of very large multicast groups, SGM can support a huge number (i.e. virtually an unlimited number) of small multicast groups and thus can play an important role in supporting applications such as conferencing applications on the Internet.

6. Security Considerations

An "eavesdropper" cannot join a multicast "group", as the source controls the membership of the multicast transmissions. Also, there is no need for a server to hand out class D addresses which could be subject to "denial of service" or other forms of attack.

7. Acknowledgements

The authors would like to thank Brian Carpenter, Dirk Ooms and Yuji Imai for their feedback and ideas. The authors would also like to thank Orit Levin of RADVision, and Pat Galvin and Jeff Durham of DataBeam, for providing an application developer's perspective. We also thank Joe Mambretti and Ralph Demuth of iCAIR, as well as Micah Beck and Bert Dempsey, co-leads of the Internet2 Distributed Storage Infrastructure Initiative, for their support in deploying SGM in the Internet2 environment.

8. References

- [1] [RFC 1075](#), Distance Vector Multicast Routing Protocol, D. Waitzman, C. Partridge, S.E. Deering, Nov. 1988
- [2] S.E. Deering. Multicast Routing in a Datagram Internetwork. Ph.D. thesis, Electrical Engineering Dept., Stanford University, Dec. 1991

[3] [RFC 1584](#), Multicast Extensions to OSPF, J. Moy, March 1994

- [4] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei. The Pim Architecture for Wide-area Multicast Routing, ACM Transactions on Networks, April 1996
- [5] [RFC 2189](#), Core Based Trees (CBT version 2) Multicast Routing Protocol Specification, A. Ballardie, Sept., 1997
- [6] [RFC 2201](#), Core Based Trees (CBT) Multicast Routing Architecture, A. Ballardie, Sept. 1997
- [7] [RFC 2236](#), Internet Group Management Protocol, Version 2, W. Fenner, Nov. 1997
- [8] [RFC 2362](#), Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, D. Estrin et al, June 1998
- [9] D. Estrin, D. Farinacci, V. Jacobson, C. Liu, L. Wei, P. Sharma, and A. Helmy, "Protocol Independent Multicast-dense Mode (pim-dm): Protocol Specification", Work in Progress.
- [10] Frequently Asked Questions (FAQ) on the Multicast Backbone (MBONE), <ftp://venera.isi.edu/mbone/faq.txt>
- [11] Reliable Multicast Transport Working Group web site, <http://www.ietf.org/html.charters/rmt-charter.html>, June 15, 1999
- [12] [RFC 1812](#), Requirements for IP Version 4 Routers, F. Baker, June 1995
- [13] [RFC 2543](#), SIP: Session Initiation Protocol, M. Handley et al, March 1999
- [14] ITU-T Recommendation H.323 (2000), Packet-Based Multimedia Communications Systems

9. Authors

Rick Boivie
IBM T. J. Watson Research Center
30 Saw Mill River Rd.
Hawthorne, NY 10532
Phone: 914-784-3251
Email: rhboivie@us.ibm.com

Nancy Feldman
IBM T. J. Watson Research Center

30 Saw Mill River Rd.
Hawthorne, NY 10532
Phone: 914-784-3254
Email: nkf@us.ibm.com

Boivie et al.

Expires August 2001

[Page 12]

[A.](#) **Appendix: SGM Encoding**

This appendix describes the packet formats for SGM layered on IP. An SGM header is comprised of a fixed header followed by one or more destination lists.

SGM packets are assigned IP protocol type:

IPPROTO_SGM 0xTBA

[A.1.](#) **Fixed Header**

The SGM fixed header may be either a bitmap-style header or a vanilla-style header, as determined by the bitmap bit within the fixed header.

[A.1.1](#) **Bitmap Fixed Header**

The bitmap fixed header allows for efficient SGM header processing. Each destination corresponds to a bit in the bitmap; if set, then a packet should be forwarded to the corresponding destination; if clear, then the corresponding destination is ignored. A modification to the destination list is achieved by simply overwriting the appropriate bits in the bitmap.

The bitmap fixed header encoding is as follows:

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Version|B|S|A|r|                               Bitmap . . .
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     |          GroupId          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Checksum          |   TTL   |   Reserved   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Version

Four-bit field indicating the format of the SGM header. This document describes version 1.

B-bit (Bitmap header)

If this bit is set, the bitmap-style header is in use, else the vanilla header is in use. See section A.1.2.

S-bit(SGM header persistence)

If this bit is set, an SGM router MUST NOT convert the SGM

packet to unicast packet(s), i.e. the packet MUST stay an SGM packet end-to-end.

A-bit (Anonymity)

If this bit is set, the destination address for which the corresponding bit in the bitmap is zero **MUST** be overwritten with zeros.

r-bit (Reserved)

Reserved bit. It must be zero on transmission and must be ignored on receipt.

Bitmap

This five octet field is a bitmap which corresponds to the SGM destinations in the header, such that the first SGM destination corresponds to the first bit, the second destination corresponds to the second bit, and so on. The *i*-th bit is set if the packet should be forwarded to the *i*-th destination.

Note that when an ICMP error message is returned to the originating sender, it is guaranteed to contain only the first 64-bits of the original IP datagram payload. Due to this limitation in the size of the ICMP error message, a size of five bytes was chosen for the bitmap. This is the maximum size that can be accommodated which still returns sufficient information in the ICMP error message, such that the sender can determine which destinations are not reachable via SGM (see [section 3.2](#)). Five bytes enforces a limit of at most forty destinations included within an SGM packet. If more than forty destinations are required, the vanilla style SGM header may be used (see section A.1.2).

GroupID

Two octet field that uniquely identifies a group of SGM destinations. This is returned in an ICMP error message, and may be used to correlate the error with the originating SGM group.

Checksum

Two octet checksum on the SGM header only. This is verified and recomputed at each point that the SGM header is modified. The checksum field is the 16 bit one's complement of the one's complement sum of all the bytes in the header. For purposes of computing the checksum, the value of the checksum field is zero.

TTL (Time-to-Live)

One octet field indicating the maximum time the SGM packet is allowed to remain in the SGM network. Each SGM router **MUST** decrement this field by one; if it is decremented to zero, the

packet MUST be discarded.

Boivie et al.

Expires August 2001

[Page 14]

Reserved

One octet reserved field. It must be zero on transmission and must be ignored on receipt.

A.1.2 Vanilla Fixed Header

The vanilla header does not contain a bitmap; SGM routers remove destinations when they are no longer on the forwarding path. This type of header gives the sender the opportunity to have a virtually unlimited number of destinations. However, this requires more per-router processing overhead than the bitmap header, as a modification to destination list requires the building of a new SGM header. Also note that this header style is not compatible with the ICMP deployment method (see [section 3.2](#))

The vanilla fixed header encoding is as follows:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Version|B|S|res|      TTL      |      Checksum      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Version

Four-bit field indicating the format of the SGM header. This document describes version 1.

B-bit (Bitmap header)

If this bit is set, the bitmap-style header is in use, else the vanilla header is in use. See section A.1.1.

S-bit(SGM header persistence)

If this bit is set, an SGM router **MUST NOT** convert the SGM packet to unicast packet(s), i.e. the packet **MUST** stay an SGM packet end-to-end.

res (reserved)

Reserved two bits. They must be zero on transmission and must be ignored on receipt.

TTL (Time-to-Live)

One octet field indicating the maximum time the SGM packet is allowed to remain in the SGM network. Each SGM router **MUST** decrement this field by one; if it is decremented to zero, the packet **MUST** be discarded.

Checksum

Two octet checksum on the SGM header. This is verified and recomputed at each point that the SGM header is processed. The


```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|               DestAddress(es) ...
~
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               PortNum(s)           |               ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

DestAddress(es)

List of destination address(es). Each address length is four bytes for IPv4 and sixteen bytes for IPv6.

PortNum(s)

List of two octet destination port number(s), where each port corresponds in placement to the preceding DestAddress(es).

