

PIM
Internet-Draft
Intended status: Standards Track
Expires: January 22, 2015

W. Atwood
B. Li
Concordia University/CSE
July 21, 2014

Group Security Association Management Protocol
draft-atwood-pim-gsam-00

Abstract

This document specifies a Group Security Association Management (GSAM) protocol, which manages the IPsec Group Security Associations that are used to protect some packets of Secure IGMP (SIGMP) and Secure MLD (SMLD). In GSAM, one router is elected as the group controller / key server to create group security associations for all the interesting secure groups and distribute them to authorized users and other routers.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 22, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	3
2. Assumption	4
3. GSAM Overview	4
4. Phase 1: Registration	4
4.1. Message Exchanges	5
4.1.1. GSAM_INIT Exchange	5
4.1.2. GSAM_AUTH Exchange	5
4.2. EU Operations	6
4.3. Non-Querier Operations	7
4.4. Querier Operations	8
5. Phase 2: GSA Distribution	9
5.1. Message Exchanges	9
5.1.1. GSAM_PUSH Exchange	9
5.1.2. GSAM_RE_DISTRIBUTION Exchange	10
5.2. Querier Operations	11
5.3. GM Operations	12
6. Handover of Q	13
7. IANA Considerations	13
8. References	13
8.1. Normative References	14
8.2. Informative References	14
Authors' Addresses	14

1. Introduction

This document specifies a Group Security Association Management (GSAM) protocol, which manages the IPsec Group Security Associations (GSAs) that are used to protect some packets of Secure IGMP (SIGMP) [I-D.atwood-pim-sigmp] and Secure MLD (SMLD) (not yet issued). GSAM is implemented in the multicast enabled segment. The Querier on this segment is responsible for distributing GSAs to all the authorized users and other routers. Negotiation of certain parameters of the GSA may be triggered if necessary.

GSAM is similar to GDOI [RFC6407] and g-ikev2 [I-D.yeung-g-ikev2], although it is different from these protocols in important ways. First, GDOI and g-ikev2 deliver only the necessary keys for IPsec, while all the parameters of the GSAs of the IPsec system are distributed in GSAM. The GSAs include not only keys, but also security parameter indexes (SPIs) of the IPsec system [RFC4301]. Second, there is a super group, 224.0.0.22 in IPv4 system or FF02:0:0:0:0:0:0:16 in IPv6 system, in GSAM. All the group members

registered in the super group are also registered in all other active groups on this network segment. Third, GSAM is a link-local protocol while GDOI and g-ikev2 are group domain protocols. In GSAM, the TTL of all the messages is equal to 1.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

In addition, the following terms are used in this document

Querier (Q):

A Querier is an edge router that has won in the querier election in SIGMP or SMLD. In GSAM, it takes the role of group controller / key server (GCKS).

Non-Querier (NQ):

A Non-Querier is an edge router that has lost in the querier election in SIGMP or SMLD.

Group Member (GM):

Group Member is an end user or an edge router that has registered in Q.

Secure Group Table (SGT):

Secure Group Table is a table in Q that records the secure groups and the GMs in the secure groups. It consists of two fields: multicast address (MA) and group member set (GMS). MA is an index of SGT and its value is a secure multicast group address. GMS contains the unicast addresses of GMs in a group identified by the value of MA. The initial SGT only has one record whose MA field is 224.0.0.22 in IPv4 system or FF02:0:0:0:0:0:0:16 in IPv6 system and whose GMS field is empty.

GSAM_TEK_SA:

GSAM_TEK_SA is a pair of GSAs, including GSA_q and GSA_r. GSA_q is a GSA of IPsec system used to protect a secure group query in SIGMP or SMLD. GSA_r is a GSA of IPsec system used to protect the secure group report in SIGMP or SMLD.

GSAM_KEY_SA:

GSAM_KEK_SA is a pair of SAs, including KEK_USA and KEK_GSA. KEK_USA is a unicast SA whose direction is from a GM to Q. It is used to protect the messages in Phase 2 sent by GMs. KEK_GSA is a GSA whose direction is from Q to a secure group. It is used to protect the messages in Phase 2 sent by Q.

2. Assumption

The protocol GSAM is based on two assumptions as follows:

The end users have been authenticated and authorized at the application layer. The authorized EU and its Q have shared the same secret key, call MSSK, as a pre-shared key. The details of how to authenticate and authorize users is not specified in this document. It is implemented based on PANA [RFC5191] as shown in draft-atwood-mboned-mrac-pana (not yet issued).

Instead of IGMP or MLD, SIGMP or SMLD is used by users (or routers) to report (or learn) IP multicast group memberships to neighboring multicast routers (or from the users that are only one IP hop away) in an IPv4 network or an IPv6 network.

3. GSAM Overview

GSAM is a protocol that manages group security associations (GSAs) of the IPsec system used to protect some packets of SIGMP or SMLD. The network entities mentioned in GSAM are the same as those in SIGMP or SMLD, including edge routers (ERs) and end users (EUs). In GSAM, an ER (called Querier) plays the role of GCKS. It accepts the registration from EUs and NQs and grants them the status of GMs in secure groups. It creates or updates GSAs of IPsec system for secure groups and distributes them to all GMs in the secure group.

Security parameter index (SPI) as a parameter of GSAs must be paid specific attention. Different from a unicast SA that is used by only one receiver, a GSA is shared by multiple receivers. As a result, instead of one receiver to determine the SPI value, all the GMs in the same secure group should negotiate the SPI value together in order to avoid SPI collisions at GMs. In GSAM, Q suggests SPI values first. If any GM rejects the offered suggestion, a negotiation will be triggered to determine suitable SPI values.

4. Phase 1: Registration

In Phase 1, both NQs and EUs should register themselves in Q in order to become GMs in a group. A pair of SAs, named GSAM_KEK_SA is distributed to GMs.

4.1. Message Exchanges

The registration involves two message exchanges: GSAM_INIT exchange and GSAM_AUTH exchange. An EU / NQ performs GSAM_INIT exchange only once as long as no new Q is elected in SIGMP or SMLD. However, an EU may perform a GSAM_AUTH exchange many times. The number of GSAM_AUTH exchanges for an EU is equal to the number of secure groups that an EU is authorized to join at the application layer.

4.1.1. GSAM_INIT Exchange

GSAM_INIT exchange is identical to IKE_SA_INIT of IKE v2 defined in [RFC5996]. An EU / NQ takes the role of an initiator and Q takes the role of a responder.

4.1.2. GSAM_AUTH Exchange

GSAM_AUTH exchange as shown in Figure 1 is similar to IKE_AUTH of IKE v2. In this exchange, an EU / NQ and Q mutual authenticate for a secure group. However, instead of being negotiated between two peers as in IKE v2, an SA pair, named GSAM_KEK_SA, is downloaded from Q to an EU / NQ.

```
EU / NQ -> Q: HDR, SK{ IDg, IDh, AUTH }
Q -> EU / NQ: HDR, SK{ IDg, SA, KD, AUTH }
```

Figure 1: GSAM_AUTH Exchange

HDR is a header payload whose format is identical to that in IKE v2. The notation SK { ... } indicates that all the payloads in "{}" are encrypted and integrity protected using an SA, called GSAM_INIT_SA, which is negotiated in the GSAM_INIT Exchange. The message exchange is explained as follows:

In the first message, an EU / NQ asserts its identification and the identification of a secure group (i.e., for an EU, it is the group that an EU requests to join in SIGMP or SMLD; for an NQ, it is the group 224.0.0.22 in IPv4 system or FF02:0:0:0:0:0:0:16 in IPv6 system listened to by all ERs) in the payload of IDh and IDg respectively. Moreover, an EU / NQ also declares a message authentication code (MAC) or its signature in the AUTH payload. The AUTH payload is used by the message receiver (Q) to authenticate the two identifications in IDh and IDg and to protect the integrity of the first message in the GSAM_INIT exchange.

In the second message, Q asserts its identification in payload IDq and distributes an SA pair, called GSAM_KEK_SA, (and more KEK_GSAs

sometimes) in payloads SA and KD. Moreover, Q also declares a MAC or its signature in AUTH payload. The AUTH payload is used by the message receiver (an EU / NQ) to authenticate the identification in payload IDq and protect the integrity of the second message in the GSAM_INIT exchange.

4.2. EU Operations

An EU initiates a GSAM_INIT exchange when an EU requests GSAs to secure SIGMP packets or SMLD packets for the first time or when an EU discovers a new Q. The EU operations in a GSAM_INIT exchange are identical to the initiator operations in the IKE_SA_INIT exchange of IKE v2.

After the GSAM_INIT exchange, a new security association, named GSAM_INIT_SA, has been negotiated. It will be used to protect the GSAM_AUTH exchange and achieve private communication between an EU and Q. Moreover, GSAM_INIT_SA will be maintained as a long-term security association. No new GSAM_INIT exchange between an EU and Q will be required for the subsequent request for GSAs as long as an EU does not discover a new Q.

An EU initiates a GSAM_AUTH exchange when a request for GSAs is received from SIGMP and GSAM_INIT_SA has been negotiated between an EU and Q. An EU must use the pre-shared key authentication method to finish the registration in the GSAM_AUTH exchange.

An EU calculates a MAC and encapsulates it in the AUTH payload of the first message of GSAM_AUTH. The calculation of the MAC is the same as that in IKE v2. The secret key used in the MAC is the MSSK for the secure group calculated at the network layer. It has been independently derived by the EU and the Q as a pre-shared key when an EU has been authorized to join in the secure group at the application layer.

Upon receiving the second message of GSAM_AUTH, an EU verifies the value in the received AUTH payload using the MSSK to authenticate Q. If verification fails, the EU will discard the received message. Otherwise, verification succeeds and the EU will accept the GSAM_KEK_SA specified in the SA and KD payloads. Moreover, an EU marks itself as a GM in the requested secure group. The EU updates its local GSPD [RFC5374] as shown in Table 1. G_IP is the IP address of the group identified in the IDg payload. Q_IP and H_IP are the IP addresses of the Q and the EU. The updated records in the GSPD indicate that the SIGMP/SMLD packets that are sent from a GM / Q to the group that a GM wants to join must be protected by IPsec.

Destination address	Source address	Protocol number	Action
G_IP	Q_IP	SIGMP (2)	IPsec protect
G_IP	H_IP	SIGMP (2)	IPsec protect
G_IP	*	SIGMP (2)	Discard
G_IP	*	*	Bypass

Table 1: Updated Records in local GSPD

Finally, the EU must update the SAD, to record the SA parameters that have been given to it.

4.3. Non-Querier Operations

An NQ initiates a GSAM_INIT exchange when an ER has just lost in the querier election for SIGMP/SMLD and has become an NQ. NQ operations in the GSAM_INIT exchange are identical to the initiator operations in IKE_SA_INIT of IKE v2.

After the GSAM_INIT exchange, GSAM_INIT_SA has been negotiated. It will be used to protect the GSAM_AUTH exchange and achieve private communication between an NQ and Q. Moreover, the GSAM_INIT_SA will be maintained as a long-term security association. No new GSAM_INIT exchange between an NQ and Q is necessary as long as an NQ does not discover a new Q.

An NQ initiates a GSAM_AUTH exchange when an ER where an NQ is located has just lost in a querier election in SIGMP / SMLD and a GSAM_INIT_SA has been negotiated between an NQ and Q. An NQ could use any authentication method configured by the network administrator to finish registration in GSAM_AUTH.

An NQ calculates a MAC or a signature according to the assigned authentication method and encapsulates it into the AUTH payload of the first message. Here the authentication method depends on the configuration of the network administrator.

Upon receiving the second message of GSAM_AUTH, an NQ verifies the value in the received AUTH payload to authenticate Q using the assigned method. If verification fails, an NQ will discard the received message. Otherwise, verification succeeds and an NQ will accept the GSAM_KEK_SA (and more KEK_GSAs if existing) specified in the SA and KD payloads. Moreover, an NQ marks itself as a GM in the

group 224.0.0.22 for IPv4 system or FF02:0:0:0:0:0:0:16 for IPv6 system (and also a GM in all the groups mentioned in KEK_GSAs). If additional KEK_GSAs are specified in SA and KD payloads, NQ also updates its local GSPD as shown in Table 1 and G_IP indicates all the IP addresses of the groups mentioned in additional KEK_GSAs.

4.4. Querier Operations

Q operations in the GSAM_INIT exchange are identical to the responder operations in IKE_SA_INIT of IKE v2.

Upon receiving the first message of GSAM_AUTH exchange, Q parses the payload of IDg. If IDg identifies a super group (224.0.0.22 for IPv4 system or FF02:0:0:0:0:0:0:16 for IPv6 system), the sender of the message is considered to be an NQ. Otherwise, the sender is considered to be an EU.

If the sender is an EU, Q retrieves the pre-shared key MSSK shared with the EU identified in the received IDh payload for a secure group identified in the received IDg payload. Similarly, if the sender is an NQ, Q retrieves a certification or a secret key of an NQ identified in IDh payload. Then Q uses the retrieved key or certification to verify the received AUTH payload. If retrieval or verification fails, Q will discard the received message and terminate the GSAM_AUTH exchange. Otherwise, it indicates that an EU has been authorized to join the secure group at the application layer or an NQ has been authorized by the network administrator in its configuration file. In this case, Q starts the "registration" to an EU for the secure group or an NQ for all secure groups.

The registration is based on a secure group table (SGT). For an NQ, Q updates all the records in SGT: the source address of the received message is added into GMS field of all the records. It means an NQ becomes a GM in all the groups that Q is maintaining. For an EU, Q searches its SGT to look for a record whose MA is the address of the group identified in the received IDg payload. If the record is found, the source address of the received message is added into GMS of the found record. It means an EU becomes a GM in the group identified in the received IDg payload. Otherwise, Q creates a new record in SGT. In the new record, the value of the MA field is the address of the group identified in the received IDg payload. The addresses showing in the GMS field of the record indexed by 224.0.0.22 (or FF02:0:0:0:0:0:0:16) are copied into the GMS field of the new record. Moreover, the source address of the received message is also filled in the GMS of the new record. It means the EU and all the registered NQs become GMs of the group identified in the received IDg payloads. After the registration of an EU, Q updates its local GSPD as Table 1.

After registration, Q creates an SA pair, named GSAM_KEK_SA, which consists of two SAs: 1) KEK_GSA and 2) KEK_USA. KEK_GSA is a group security association whose direction is from Q to a secure group identified in the received IDg payload. In detail, when the exchange is between an EU and Q, the direction of KEK_GSA is from Q to a secure group that an EU requests to join. When the exchange is between an NQ and Q, the direction is from Q to the group 224.0.0.22 or FF02:0:0:0:0:0:0:16. It is used to protect the messages in Phase 2 sent by Q. KEK_USA is a unicast security association whose direction is from the new GM (an EU/NQ) to Q. It is used to protect the message in Phase 2 sent by GMs.

Moreover, Q also calculates a new MAC or a signature according to the negotiated authentication method. If the exchange is between an EU and Q, the authentication method must be pre-shared key. Q uses the retrieved MSSK as the secret key to calculate a MAC value. If the exchange is between an NQ and Q, the authentication method depends on the network configuration of an NQ. Q may calculate a MAC or a signature for NQ.

After that, Q sends to an EU / NQ the second message as a response. In the response to an EU, SA and KD payloads specify the newly created GSAM_KEK_SA. In the response to an NQ, SA and KD payloads specified not only the newly created GSAM_KEK_SA, but also all other KEK_GSAs that Q is maintaining. the AUTH payload contains the new MAC or signature.

5. Phase 2: GSA Distribution

In Phase 2, Q suggests GSAM_TEK_SA to GMs. If any GM rejects the suggestion due to SPI collisions, a negotiation will be required among GMs.

5.1. Message Exchanges

There are two exchanges in GSA distribution: GSAM_PUSH and GSAM_RE_DISTRIBUTION.

5.1.1. GSAM_PUSH Exchange

GSAM_PUSH exchange is shown in Figure 2 .

Q -> GMs: HDR, SK{ SA, KD, AUTH}

Figure 2: GSAM_PUSH Exchange

In this message, Q distributes an SA pair (i.e., GSAM_TEK_SA, but sometimes more GSAM_TEK_SAs) or an SA (i.e., KEK_GSAs) in the payload SA and KD. Moreover, Q declares a signature in payload AUTH. The notation SK {...} indicates that all the payloads in "{}" are encrypted and integrity protected using a KEK_GSA.

5.1.2. GSAM_RE_DISTRIBUTION Exchange

The GSAM_RE_DISTRIBUTION exchange is triggered when any GM detects an SPI collision and refuses to accept the GSAM_TEK_SA received in the GSAM_PUSH message. In other words, it is just optional: there is no GSAM_RE_DISTRIBUTION exchange if no SPI collisions are detected by any GM. The GSAM_RE_DISTRIBUTION exchange is shown in Figure 3. All the messages are protected by GSAM_KEK_SA. It is explained as follows:

```
GM -> Q : HDR, SK{ IDg, REJ, AUTH }
Q -> GMs: HDR, SK{ S_REQ, AUTH }
GMs -> Q: HDR, SK{ SPI_LIST, AUTH }
Q -> GMs: HDR, SK{ SATf, KDtf, AUTH }
```

Figure 3: GSAM_RE_DISTRIBUTION Exchange

In the first message, a GM that detects an SPI collision asserts the identification of the secure group that is the destination address of the rejected GSAM_TEK_GSA in payload IDg and shows its rejection to the suggested GSAM_TEK_GSA in payload REJ. Moreover, GM declares a MAC in payload AUTH.

Q multicasts the second message into a secure group identified by the received IDg. In this message, Q requests a list, called spi_list, in payload S_REQ and shows its signature in payload AUTH.

All GMs in the secure group will send the third message to respond to the request from Q. In the third message, a GM reports its spi_list in payload SPI_LIST and declares its MAC in the payload AUTH.

The fourth message is the same as the GSAM_PUSH message. In the fourth message, Q multicasts an SA pair (i.e., GSAM_TEK_SA) in payload SA and KD and declares a signature in the AUTH payload. However, the SPI parameter in GSAM_TEK_SA has been negotiated with all GMs in the secure group and therefore it cannot cause any collision.

5.2. Querier Operations

When an EU is registered as the first GM of a secure group in the segment, Q will multicast the GSAM_PUSH exchange message in two groups in order: (1) the group 224.0.0.22 or FF02:0:0:0:0:0:0:16 and (2) the secure group that an EU requests to join.

Q multicasts the first multicast GSAM_PUSH message into group 224.0.0.22 or FF02:0:0:0:0:0:0:16. In this message, the KEK_GSA that is just distributed to an EU using GSAM_AUTH exchange is specified in the payloads of SA and KD.

Q creates a new SA pair, called GSAM_TEK_SA, which consists of two GSAs: (1) GSA_q whose direction is from Q to the secure group that an EU (the new GM) requests to join and (2) GSA_r whose direction is from an EU to the secure group that an EU requests to join. The values of important parameter of SPI in GSA_q and GSA_r are suggested ones since they are assigned by Q with no negotiation with other GMs.

After making sure that all the NQs have received the previous GSAM_PUSH message, Q starts to multicast the other GSAM_PUSH message into the group that the EU has requested to join. In the second message, the payloads SA and KD specify the parameter and key material of the new SA pair (GSAM_TEK_SA).

After that, Q starts two timers, called q-timer and r-timer respectively. When q-timer / r-timer expires, Q will update its local SAD [RFC5374] according to GSA_q / GSA_r. The initial value of q-timer should be large enough to make sure all GMs have updated their local SADs according to the distributed GSA_q.

There must be an interval between the first GSAM_PUSH message and the second one. The interval should be large enough to make sure the first message has been received by GMs in 224.0.0.22 or FF02:0:0:0:0:0:0:16 before the second one is sent.

If the registered EU is not the first GM of a secure group, Q multicasts the second GSAM_PUSH message directly without the first message.

When an NQ is registered as a GM in all the groups, Q will directly multicast GSAM_PUSH exchange message in the group of 224.0.0.22 for IPv4 system or FF02:0:0:0:0:0:0:16 for IPv6 system. In this message, GSAM_TEK_GSAs for all the groups are specified in the payloads SA and KD. If the only group Q is maintaining is the super group, no GSAM_PUSH exchange is needed.

Upon receiving the first message of GSAM_RE_DISTRIBUTION, Q verifies the value in the payload AUTH to authenticate a GM. If authentication fails, Q discards the received message directly. Otherwise, Q deletes the q_timer and r_timer if they exist. It multicasts the second message of GSAM_RE_DISTRIBUTION to negotiate SPI values with all the GMs in the secure group.

Upon receiving the third message, Q verifies the AUTH payload to authenticate a GM. If authentication fails, Q discards the received message directly. Otherwise, Q searches its local SGT and looks for a record that is indexed by a secure group identified in the received IDg. The GMS of the found record contains the addresses of all the GMs in the secure group. Q compares the source addresses of the received third messages with the values in the GMS until it has received the third message from all the GMs in the secure group.

After that, Q starts to calculate a list, called spi_list_all, which is a union of spi_lists received from all GMs in the secure group. Then Q resets the values of SPI in GSAM_TEK_SA. The new SPI values must not be in the spi_list_all to effectively avoid SPI collisions at any GM. Then Q multicasts the fourth message of GSAM_RE_DISTRIBUTION, whose payloads SA and KD specify the revised GSAM_TEK_SA. Finally, Q re-starts q-timer and r-timer. When q-timer / r-timer expires, Q updates its local SAD according to GSA_q / GSA_r whose SPI value has been negotiated among GMs.

5.3. GM Operations

Upon receiving the GSAM_PUSH message, a GM verifies the value in the payload AUTH to authenticate Q. If authentication fails, a GM discards the received message directly. Otherwise, a GM parses the received payloads SA and KD. If payloads SA and KD specify KEK_GSAs and a GM is an NQ, a GM will accept the KEK_GSA directly and wait for receiving the following GSAM_PUSH message protected by KEK_GSA. Otherwise payloads SA and KD specify a GSAM_TEK_SA. In this case, a GM checks SPI, an important parameter of GSAM_TEK_SA. If SPI values in GSAM_TEK_SA have not used in its local SAD, a GM will start q-timer and r-timer and no other exchange is needed. When q-timer / r-timer expires, a GM updates its local SAD according to GSA_q / GSA_r. If the source address of received GSA_r is the same as a local address, the initial value of r-timer should be large enough to make sure all other GMs and Q have updated their local SADs according to GSA_r. If the suggested SPI values in GSAM_TEK_SA have collided with the used SPI values in local SAD, a GM must start GSAM_RE_DISTRIBUTION exchange as follows.

A GM calculates a MAC and encapsulates it in AUTH payload. Then it sends the first message of GSAM_RE_DISTRIBUTION to Q to show its rejection.

Upon receiving the second message in GSAM_RE_DISTRIBUTION, a GM verifies the received AUTH payload. If the verification fails, a GM discards the received message. Otherwise, a GM deletes the pending q-timer and r-timer at once if they exist. It accesses its local SAD to obtain the all the used SPI values in the SAD and saves them in an spi_list. After that, the status of local SADB is set as "read_only" to prevent any modification from any other processes. The GM encapsulates an spi_list in the payload SPI_LIST. Moreover, a MAC value is calculated and encapsulated in the AUTH payload. After that, the GM sends the third message with the payload SPI_LIST and AUTH payload.

Upon receiving the fourth message in GSAM_RE_DISTRIBUTION, the GM verifies the value in the payload AUTH to authenticate Q. If authentication fails, the GM discards the received message directly. Otherwise, the GM is forced to accept GSAM_TEK_SA specified in the received payload SA and KD. It re-starts q-timer and r-timer. When q-timer/r-timer expires, a GM updates its local SAD according to GSA_q/GSA_r. After that, GMs clears the "read_only" status of its local SAD to permit the modification to the SAD from other processes.

6. Handover of Q

Although ERs are usually stable, a new ER may be added into the network and an old ER may fail to work. In these cases, a querier election is caused and then a new Q may be elected in the link. The new Q will take over the work of old Q automatically and become GCKS soon in the link. All the EUs and NQs will discover the new Q since they will receive the general query sent by the new Q in SIGMP/SMLD. They initiate new GSAM sessions with the new Q. If they are authenticated successfully, the new Q will distribute new GSAM_TEK_SAs to them. SIGMP / SMLD messages will be protected by the new GSAM_TEK_SAs.

7. IANA Considerations

GSAM runs over UDP. A UDP port should be assigned to GSAM.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

- [I-D.atwood-pim-sigmp]
william.atwood@concordia.ca, w. and B. Li, "Secure Internet Group Management Protocol", draft-atwood-pim-sigmp-01 (work in progress), July 2014.
- [I-D.yeung-g-ikev2]
Rowles, S., Yeung, A., Tran, P., and Y. Nir, "Group Key Management using IKEv2", draft-yeung-g-ikev2-07 (work in progress), November 2013.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.
- [RFC5374] Weis, B., Gross, G., and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", RFC 5374, November 2008.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, October 2011.

Authors' Addresses

William Atwood
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Phone: +1(514)848-2424 ext3046
Email: william.atwood@concordia.ca
URI: <http://users.encs.concordia.ca/~bill>

Bing Li
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Email: leebingice@gmail.com

PIM
Internet-Draft
Intended status: Standards Track
Expires: September 4, 2014

W. Atwood
Concordia University/CSE
S. Venaas
Cisco
March 03, 2014

IANA Allocation of Experimental Code Points for PIM Join Attribute and
PIM Encoded-Source Address
draft-atwood-pim-reserve-exp-00

Abstract

This memo asks the IANA to allocate experimental code points to the PIM Join Attribute Types register and the Encoded-Source Address Encoding Type Field register.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 4, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	2
1.1. Terminology	2
2. Background	2
3. Security Considerations	3
4. IANA Considerations	3
5. Acknowledgements	3
6. References	3
6.1. Normative References	3
6.2. Informative References	4
Authors' Addresses	4

1. Introduction

To make it possible to experiment with protocol extensions safely, [RFC3692] recommends that "protocol documents should consider reserving a small set of protocol numbers for experimentation."

Two IANA registries related to Protocol Independent Multicast (PIM) do not reserve any numbers for experimentation.

This document requests the IANA to reserve two numbers for the Registry "PIM Join Attribute Types" and four numbers for the "Encoded-Source Address Encoding Type Field".

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Background

In the Protocol Independent Multicast (PIM) Parameters [URL] Protocol Registry, two sub-registries have no allocation for either Private or Experimental use. They are "Encoded-Source Address Encoding Type Field", for which the current definitions are given in [RFC4601] and [RFC5384], and "PIM Join Attribute Types", for which the current definitions are given in [RFC5384].

For the Encoded-Source Address Encoding Type Field, two values are assigned, and the remainder (2-255) are unassigned.

For the PIM Join Attribute Types, four values are assigned, and the remainder (4-63) are unassigned.

The remaining sub-registries all have values assigned for Private Use, for Experimental Use, or for extension of the type space.

The registrations proposed in this document are of type Experimental [RFC5226], because the expected usage of these reservations would not likely be confined to a single site.

3. Security Considerations

This document only assigns values in two IANA registries. The security implications of the use of these values must be considered by those who make use of them.

4. IANA Considerations

The requests in this document are for two registries that are part of the "Protocol Independent Multicast (PIM) Parameters" Registry.

This document requests the IANA to reserve four values in the "Encoded-Source Address Encoding Type Field" registry:

Type	Name	Reference
----	----	-----
251-255	Reserved (Experimental)	[RFC4601][RFC5384][this doc]

This document requests the IANA to reserve two values in the "PIM Join Attribute Types" registry:

Type	Name	Reference
----	----	-----
64-65	Reserved (Experimental)	[RFC5384][this doc]

5. Acknowledgements

Adrian Farrel and Brian Haberman observed that there were no reservations for Experimental Use for the PIM Join Attribute Type.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, November 2008.

6.2. Informative References

- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, January 2004.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

Authors' Addresses

William Atwood
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Phone: +1(514)848-2424 ext3046
Email: william.atwood@concordia.ca
URI: <http://users.encs.concordia.ca/~bill>

Stig Venaas
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA

Email: stig@cisco.com

PIM
Internet-Draft
Intended status: Standards Track
Expires: January 5, 2015

W. Atwood
B. Li
Concordia University/CSE
July 4, 2014

Secure Internet Group Management Protocol
draft-atwood-pim-sigmp-01

Abstract

This document specifies a Secure Internet Group Management Protocol (SIGMP), which is an extension to IGMP to enforce receiver access control for secured multicast groups. In SIGMP, only the hosts operated by authorized end users are permitted to report their interest in secured groups. IPsec is used to filter the messages that report or query the interest in secured groups. SIGMP provides two working modes that are fully compatible with IGMP v2 and IGMP v3 respectively.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 5, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	3
1.2. Assumptions	3
2. Overview of SIGMP	4
3. Packet Format	5
4. Router Operations	5
4.1. Router Operations Compatible with IGMP v2	5
4.1.1. Router Operations for a Received Report	5
4.2. Router Operations Compatible with IGMP v3	7
4.2.1. Router Operations on a Received Report	7
5. Host Operations	8
5.1. Host Operations Compatible with IGMP v2	8
5.1.1. Conditions for Unsolicited Report	8
5.1.2. Host Operations for a Received Query	8
5.2. Host Operations Compatible with IGMP v3	9
5.2.1. Host Operations for a Received General Query	9
6. IANA Considerations	9
7. References	9
7.1. Normative References	9
7.2. Informative References	9
Authors' Addresses	10

1. Introduction

The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to any neighboring multicast routers. There are two popular versions: IGMP v2, as specified in [RFC2236] and IGMP v3, as specified in [RFC3376]. However, both versions establish a fully "open" multicast network, where any host can join any multicast group as a recipient without receiver access control.

This document specifies a Secure Internet Group Management Protocol (SIGMP) working in a "hybrid" multicast network. In a hybrid network, multicast groups are classified into two categories: open groups and secured groups. Open groups refer to multicast groups that any host can join unconditionally as a receiver. Secured groups refer to multicast groups with receiver access control, e.g., only hosts operated by authenticated and authorized end users are permitted to join as receivers. SIGMP retains most mechanisms of IGMP and enforces receiver access control to secured groups in a multicast network. On the one hand, any host could report its

interest in open groups freely as in IGMP. On the other hand, only hosts operated by the authenticated and authorized end users are permitted to report their interest in secured groups.

Instead of a new specific mechanism, SIGMP uses IPsec [RFC4301] to implement receiver access control to secured groups at the IP layer. Some Security Associations (SAs) are created to secure the SIGMP packets that are used to report or query secured groups. The packets coming from the unauthorized hosts will be discarded by the IPsec subsystem if they are used to report or query interest in secured groups.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

It is assumed that the reader is familiar with the defining documents for IGMP [RFC2236] and [RFC3376]. Unless otherwise noted, terms defined in these documents are used with the same meaning in this one.

In addition, the following terms are used in this document.

open group: A multicast group without receiver access control. Any host can unconditionally join any open group as a receiver, e.g. the data in a open group can be received by any host.

secured group: A multicast group with receiver access control. Only hosts operated by authenticated and authorized end users are permitted to join a secured group as a receiver, e.g. the data in a secured group can only be received by hosts operated by authenticated and authorized end users.

1.2. Assumptions

In order to focus on the actions of group membership (e.g., joining and leaving groups), the following topics are assumed to be discussed elsewhere:

1. how to distinguish between secured groups and open groups;
2. how to authenticate and authorize the operators of the devices (hosts and routers);

3. how to distribute the necessary Security Associations to participant devices (hosts and routers).

The existence of the group property (secured or open) defines the hybrid nature of the environment in which SIGMP works. A variety of existing protocols (e.g., LDAP) can be used to enquire as to the status of a particular multicast group.

The hosts that show interest in secured groups MUST be operated by authenticated and authorized end users. One approach to the task of authentication and authorization of end users is based on the use of PANA [RFC5191] and EAP [RFC3748], and is described in [I-D.atwood-mboned-mrac-req], [I-D.atwood-mboned-mrac-arch] and draft-atwood-mboned-pana (not yet published).

A coordination protocol may be needed to manage and distribute the Security Associations (SAs) for secured groups among the routers and the hosts that correspond to authenticated and authorized end users. One set of possible procedures for SA creation and maintenance is specified in draft-atwood-pim-gsam (not yet published).

2. Overview of SIGMP

SIGMP is an extension to IGMP and performs receiver access control for groups in a multicast network. It retains most mechanisms of IGMP and has two working modes: 1) mode compatible with IGMP v2 and 2) mode compatible with IGMP v3. It works in either mode and is transparent for hosts that support only IGMP, i.e., that do not support SIGMP. In addition, SIGMP uses IPsec to secure part of its packets. For an open group, it delivers the data to any host unconditionally as IGMP does. However, for a secured group, SIGMP only delivers the data to the hosts that have established SAs in the IPsec subsystem in order to perform access control.

In a network segment, hosts show their interest in secured groups using IPsec protected packets although their interest for open groups is still reported using unprotected packets. Similarly, routers query the membership interest for a secured group using IPsec protected packets, although the general query and the query for the membership of open groups are performed using unprotected packets.

In general, the packets in SIGMP are classified into four categories, which are Query for Open Group (OGQ), Query for Secured Group (SGQ), Report for Open Group (OGR) and Report for Secured Group (SGR). OGQ and SGQ are sent by the the Querier and are used to learn the membership of open groups (or all groups for general query) and secured groups respectively. In detail, OGQ includes general query, specific-group query for open group and group-and-source-specific

query for the source of open group. SGQ includes specific-group query for secured group and group-and-source-specific query for the source of secured group. OGR and SGR are sent by hosts and used to report the membership of open groups and secured groups respectively. In detail, OGR includes report to specific-group query for open group, report to group-and-source-specific query for the source of open group, unsolicited report for open group and part of reports to general query. SGR includes report to specific-group query for secured group, report to group-and-source-specific query for the source of secured group, unsolicited report for secured group and part of reports to general query. SGQ and SGR are protected by IPsec at IP layer while OGQ and OGR are delivered without IPsec protection.

The destination address of packets in IP layer is specified as follows. In SGQ and SGR, the destination address is a secured group address. In OGQ, it is 224.0.0.1 if the packet is general query and otherwise it is an open group address. In OGR, it is 224.0.0.22 if the packet is the report to general query compatible with IGMP v3 and otherwise it is an open group address. The two addresses of 224.0.0.1 and 224.0.0.22 are the open group addresses. NOTE: When SIGMP works in the mode compatible with IGMP v3, the response to a general query contains zero or one OGR and zero or more SGR. It is described in detail in Section 5.2.1.

3. Packet Format

The packet format of SIGMP is identical to the packet format for IGMP. In detail, the format is the same as IGMP v2 when SIGMP works in the mode compatible with IGMP v2. The format is the same as IGMP v3 when SIGMP works in the mode compatible with IGMP v3.

4. Router Operations

Router operations in SIGMP are based on router operations in IGMP. However, some additional operations must be appended since access control to secured groups is extended into SIGMP. This section describes the additional operations for the two working modes.

4.1. Router Operations Compatible with IGMP v2

The additional router operations focus on the operations for a received report.

4.1.1. Router Operations for a Received Report

On receiving a report, a router checks the group address in the received report. If the group address indicates an open group, the report is considered as an OGR. A router will process an OGR as it

does that in IGMP v2 directly. Otherwise, the received report is an SGR that SHOULD just have been authenticated (and decrypted) by the IPsec subsystem (e.g., AH [RFC4302]). For SGR, a router must perform two verifications: address consistency and SA existence.

In the address consistency verification, a router compares two addresses: the group address in the SIGMP report and the destination address in the IP header. The verification fails if the two addresses are not the same. In the failure case, the sender of the IGMP Report has attempted to hide a request for a specific group (probably a secured group) in an IGMP Report for a different group (probably an open group). This will cause the IPsec subsystem to deliver the IGMP Report without requiring it to be protected. Therefore a router must discard the report if this address consistency verification fails.

In the SA existence verification, a router checks whether SAs have been established for the secured group whose address is contained in the received report. The verification fails if there are no valid SAs for the group in the router's IPsec subsystem. Since the IPsec subsystem is used to enforce the access control, no access to a secured group is permitted until its SAs have been established. Therefore a router must discard the report if this verification fails.

If the two verifications succeed on SGR, a router will proceed to update the group memberships and refresh the timers as it does in IGMP v2. In summary, the router operations for a received report are shown in Table 1.

#	Group Address	Address Consistency	SA Existence	Operations for Report
1	Open	-	-	Process as IGMP v2
2	Secured	No	-	Discard
3	Secured	Yes	No	Discard
4	Secured	Yes	Yes	Process as IGMP v2

Table 1: Router Operations for a Received Report for the Mode Compatible with IGMP v2

4.2. Router Operations Compatible with IGMP v3

The additional router operations still focus on the operations for a received report. However, there is a little difference between the operations in the mode compatible with IGMP v3 and the operations in the mode compatible with IGMP v2, since the formats of received reports in the two modes are different.

4.2.1. Router Operations on a Received Report

On receiving a report, a router checks the number of group records in the report. If the number is more than one, it indicates that the report is an OGR, but not an SGR, since only one group record is included in an SGR. In this case, every group record in the report must be verified further as follows. A router checks the multicast address in the group record. If the multicast address is an open group address, a router will process the group record as it does in IGMP v3. Otherwise, a secured group address is in the group record and a router must discard the group record. The OGR including more than one group records is not protected by IPsec systems and is not permitted to contain any information related to any secured group.

In contrast, if the number of the group records is just one, a router still checks the multicast address in the single group record. If the multicast address indicates an open group address, the received report is considered as an OGR and a router will process the group record as it does that in IGMP v3 directly. Otherwise, the received report SHOULD be an SGR that SHOULD just be authenticated (and decrypted) by the IPsec subsystem. For the single group record in the SGR, a router must perform two verifications, address consistency and SA existence, similar to Section 4.1.

In the address consistency verification, a router compares two addresses: the multicast address in the group record of the SIGMP report and the destination address in the IP header. A router must discard the report if the two addresses are not the same.

In SA existence verification, a router checks whether SAs have been established for the secured group whose address is contained in the group record of the received report. A router must discard the report if there are no SAs established in the router's IPsec subsystem.

If the two verifications succeed on an SGR, a router will proceed to update the group memberships and refresh the timers as it does in IGMP v3. In summary, router operations for a received report are shown in Table 2.

#	#Group record in report	Multicast Address in Group Record	Address Consistency	SA Existence	Operations for Group Record
1	>1	Open	-	-	Process as IGMP v2
2	>1	Secured	-	-	Discard
3	=1	Open	-	-	Process as IGMP v2
4	=1	Secured	No	-	Discard
5	=1	Secured	Yes	No	Discard
6	=1	Secured	Yes	Yes	Process as IGMP v2

Table 2: Router Operations for a Received Report for Mode Compatible with IGMP v3

5. Host Operations

Host operations in SIGMP are based on host operations in IGMP. However, some additional operations must be appended since access control to secured group is extended into SIGMP. This section describes the additional operations for the two working modes.

5.1. Host Operations Compatible with IGMP v2

The additional host operations focus on the conditions for unsolicited report and the operations for a received query.

5.1.1. Conditions for Unsolicited Report

Before creating an unsolicited report, a host must check the reported group. If the report group is open, a host will do as in IGMP v2. If secured, a host must continue to check whether SAs have been established for the secured group. If no SA is defined for this group address, a host MUST return an error indication to the issuer of the request that provoked the unsolicited report. [[Is this the right behavior?]]

5.1.2. Host Operations for a Received Query

On receiving the query, a host does the additional operation as a router does in Section 4.2.1.

5.2. Host Operations Compatible with IGMP v3

The additional host operations focus on three aspects: 1) the conditions for unsolicited report, 2) the operations for a received non-general query and 3) the operations for a received general query. The first two are identical to those described in Section 5.1.1 and Section 5.1.2. In this subsection, only the last case is explained.

5.2.1. Host Operations for a Received General Query

When it determines to respond to a general query, a host creates zero or one OGR and zero or more SGR in SIGMP instead of one report in IGMP v3. The OGR reports the current state of all the open groups that the host is interested in. Each SGR reports the current state of one secured group that the host is interested in.

At the IP layer, the destination address of OGR is 224.0.0.22. In contrast, at the IP layer the destination addresses of SGRs are the secured group addresses. Since IPsec has established SAs for secured groups, SGRs will be protected and the OGR will not.

6. IANA Considerations

The protocol number of SIGMP is the same as IGMP.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

7.2. Informative References

- [I-D.atwood-mboned-mrac-arch]
william.atwood@concordia.ca, w., Li, B., and S. Islam,
"Architecture for IP Multicast Receiver Access Control",
draft-atwood-mboned-mrac-arch-00 (work in progress),
October 2013.
- [I-D.atwood-mboned-mrac-req]
william.atwood@concordia.ca, w., Islam, S., and B. Li,
"Requirements for IP Multicast Receiver Access Control",
draft-atwood-mboned-mrac-req-00 (work in progress),
October 2013.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H.
Levkowetz, "Extensible Authentication Protocol (EAP)", RFC
3748, June 2004.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December
2005.
- [RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A.
Yegin, "Protocol for Carrying Authentication for Network
Access (PANA)", RFC 5191, May 2008.

Authors' Addresses

William Atwood
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Phone: +1(514)848-2424 ext3046
Email: william.atwood@concordia.ca
URI: <http://users.encs.concordia.ca/~bill>

Bing Li
Concordia University/CSE
1455 de Maisonneuve Blvd, West
Montreal, QC H3G 1M8
Canada

Email: leebingice@gmail.com

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: January 4, 2015

IJ. Wijnands
S. Venaas
Cisco Systems, Inc.
M. Brig
Aegis BMD Program Office
July 3, 2014

PIM flooding mechanism and source discovery
draft-ietf-pim-source-discovery-bsr-01

Abstract

PIM Sparse-Mode uses a Rendezvous Point (RP) and shared trees to forward multicast packets to Last Hop Routers (LHR). After the first packet is received by the LHR, the source of the multicast stream is learned and the Shortest Path Tree (SPT) can be joined. This draft proposes a solution to support PIM Sparse Mode (SM) without the need for PIM registers, RPs or shared trees. Multicast source information is flooded throughout the multicast domain using a new generic PIM flooding mechanism. This mechanism is defined in this document, and is modeled after the PIM Bootstrap Router protocol. By removing the need for RPs and shared trees, the PIM-SM procedures are simplified, improving router operations, management and making the protocol more robust.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 4, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Conventions used in this document	3
1.2. Terminology	3
2. A generic PIM flooding mechanism	3
2.1. PFP message format	4
3. Distributing Source to Group Mappings	5
3.1. Group Source Holdtime TLV	5
4. Originating SG messages	6
5. Processing SG messages	7
6. The first packets and bursty sources	7
7. Resiliency to network partitioning	8
8. Security Considerations	9
9. IANA considerations	9
10. Acknowledgments	9
11. References	9
11.1. Normative References	9
11.2. Informative References	9
Authors' Addresses	10

1. Introduction

PIM Sparse-Mode uses a Rendezvous Point (RP) and shared trees to forward multicast packets to Last Hop Routers (LHR). After the first packet is received by the LHR, the source of the multicast stream is learned and the Shortest Path Tree (SPT) can be joined. This draft proposes a solution to support PIM Sparse Mode (SM) without the need for PIM registers, RPs or shared trees. Multicast source information is flooded throughout the multicast domain using a new generic PIM flooding mechanism. This mechanism is defined in this document, and is modeled after the Bootstrap Router protocol [RFC5059]. By removing the need for RPs and shared trees, the PIM-SM procedures are simplified, improving router operations, management and making the protocol more robust.

1.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Terminology

RP: Rendezvous Point.

BSR: Bootstrap Router.

RPF: Reverse Path Forwarding.

SPT: Shortest Path Tree.

FHR: First Hop Router, directly connected to the source.

LHR: Last Hop Router, directly connected to the receiver.

SG Mapping: Multicast source to group mapping.

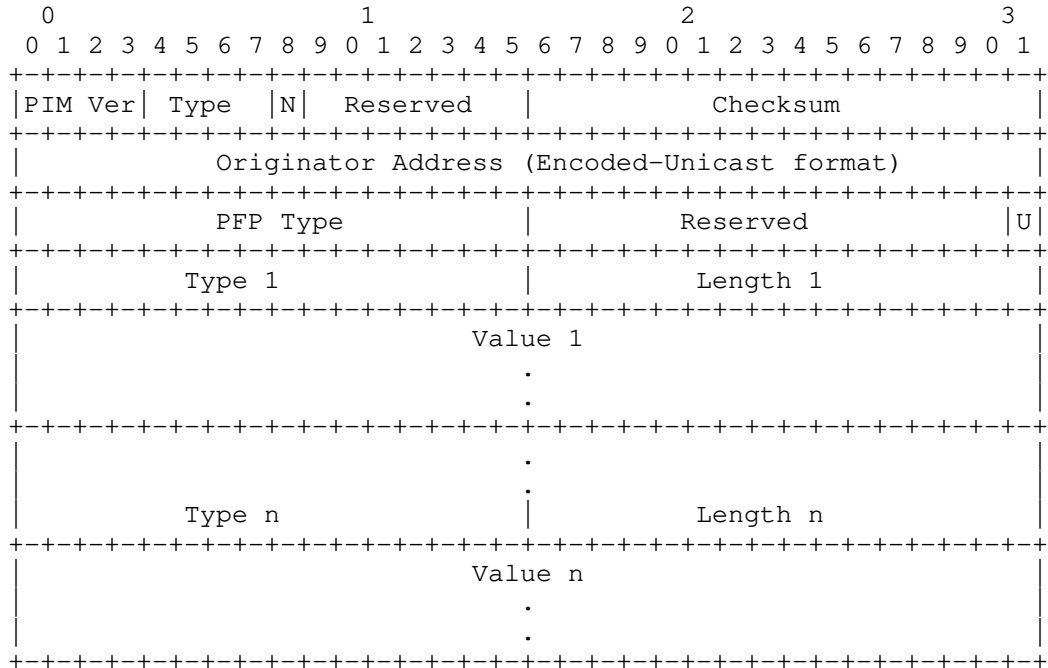
SG Message: A PIM message containing SG Mappings.

2. A generic PIM flooding mechanism

The Bootstrap Router protocol (BSR) [RFC5059] is a commonly used protocol for distributing dynamic Group to RP mappings in PIM. It is responsible for flooding information about such mappings throughout a PIM domain, so that all routers in the domain can have the same information. BSR as defined, is only able to distribute Group to RP mappings. We are defining a more generic mechanism that can flood any kind of information throughout a PIM domain. It is not necessarily a domain though, it depends on the administrative boundaries being configured. The forwarding rules are identical to BSR, except that there is no BSR election and that one can control whether routers should forward messages of unsupported types. For some types of information it is quite useful that it can be distributed without all routers having to support the particular type, while there may also be types where it is necessary for every single router to support it. The protocol includes an originator address which is used for RPF checking to restrict the flooding, just like BSR. Just like BSR it is also sent hop by hop. Note that there is no built in election mechanism as in BSR, so there can be multiple originators. It is still possible to add such an election mechanism on a type by type bases if this protocol is used in scenarios where this is desirable. We include a type field, which can allow

boundaries to be defined, and election to take place, independently per type. We call this protocol the PIM Flooding Protocol (PFP).

2.1. PFP message format



PIM Version: Reserved, Checksum Described in [RFC4601].

Type: PIM Message Type. Value (pending IANA) for a PFP message.

[N]o-Forward bit: When set, this bit means that the PFP message is not to be forwarded.

Originator Address: The address of the router that originated the message. This can be any address assigned to this router, but MUST be routable in the domain to allow successful forwarding (just like BSR address). The format for this address is given in the Encoded-Unicast address in [RFC4601].

PFP Type: There may be different sub protocols or different uses for this generic protocol. The PFP Type specifies which sub protocol it is used for.

[U]nknown-No-Forwarding bit: Some sub protocols may require that each router do some processing of the contents and not simply

forwarding. This bit controls how a router should treat an unknown PFP Type. When set, a router MUST NOT forward the message when the PFP Type is unknown. When clear, a router MUST forward the message when possible. If the PFP Type is known, then the specification of that type will specify how to handle the message, including whether it should be forwarded.

Type 1..n: A message contains one or more TLVs, in this case n TLVs. The Type specifies what kind of information is in the Value. Note that the Type space is shared between all PFP types. Not all types make sense for all PFP types though.


Length 1..n: The length of the the value field.

Value 1..n: The value associated with the type and of the specified length.

3. Distributing Source to Group Mappings

We want to provide information about active multicast sources throughout a PIM domain by making use of the generic flooding mechanism defined in the previous section. We request PFP Type 0 to be assigned for this purpose. We call a message with PFP Type 0 an SG Message. We also define a PFP TLV which we request to be type 0. How this TLV is used with PFP Type 0 is defined in the next section. Other PFP Types may specify the use of this TLV for other purposes. For PFP Type 0 the U-bit MUST NOT be set. This means that routers not supporting PFP Type 0 would still forward the message.

3.1. Group Source Holdtime TLV



Type: This TLV has type 0.

Length: The length of the value.

Group Address: The group we are announcing sources for. The format for this address is given in the Encoded-Group format in [RFC4601].

Src Count: How many unicast encoded sources address encodings follow.

Src Holdtime: The Holdtime (in seconds) for the corresponding source(s).

Src Address: The source address for the corresponding group. The format for these addresses is given in the Encoded-Unicast address in [RFC4601].

4. Originating SG messages

An SG Message, that is a PFP message of Type 0, may contain one or more Group Source Holdtime TLVs. This is used to flood information about active multicast sources. Each FHR that is directly connected to an active multicast source originates SG BSR messages. How a multicast router discovers the source of the multicast packet and when it considers itself the FHR follows the same procedures as the registering process described in [RFC4601]. After it is decided that a register needs to be sent, the SG is not registered via the PIM SM register procedures, but the SG mapping is included in an SG message.

Note, only the SG mapping is distributed in the message, not the entire packet as would have been done with a PIM register. The router originating the SG messages includes one of its own addresses in the originator field. Note that this address must be routeable due to RPF checking. The SG messages are periodically sent for as long as the multicast source is active, similar to how PIM registers are periodically sent. The default announcement period is 60 seconds, which means that as long as the source is active, it is included in an SG message originated every 60 seconds. The holdtime for the source is by default 210 seconds. Other values can be configured, but the holdtime must be larger than the announcement period. It is RECOMMENDED to be 3.5 times the announcement period. Note that as a special case a source MAY be announced with a holdtime of 0 to indicate that the source is no longer active.

5. Processing SG messages

A router that receives an SG message should parse the message and store the SG mappings with a holdtimer started with the advertised holdtime for that group. If there are directly connected receivers for that group this router should send PIM (S,G) joins for all the SG mappings advertised in the message. The SG mappings are kept alive for as long as the holdtimer for the source is running. Once the holdtimer expires a PIM router SHOULD send a PIM (S,G) prune to remove itself from the tree. Note that a holdtime of 0 has a special meaning. It is to be treated as if the source just expired, causing a prune to be sent and state to be removed. Source information MUST not be removed due to it being omitted in a message. For instance, if there are a large number of sources for a group, there may be multiple SG messages for the same group, each message containing a different list of sources.

6. The first packets and bursty sources

The PIM register procedure is designed to deliver Multicast packets to the RP in the absence of a native SPT tree from the RP to the source. The register packets received on the RP are decapsulated and forwarded down the shared tree to the LHRs. As soon as an SPT tree is built, multicast packets would flow natively over the SPT to the RP or LHR and the register process would stop. The PIM register process ensures packet delivery until an SPT tree is in place reaching the FHR. If the packets were not unicast encapsulated to the RP they would be dropped by the FHR until the SPT is setup. This functionality is important for applications where the initial packet(s) must be received for the application to work correctly. Another reason would be for bursty sources. If the application sends out a multicast packet every 4 minutes (or longer), the SPT is torn down (typically after 3:30 minutes of inactivity) before the next

packet is forwarded down the tree. This will cause no multicast packet to ever be forwarded. A well behaved application should really be able to deal with packet loss since IP is a best effort based packet delivery system. But in reality this is not always the case.

With the procedures proposed in this draft the packet(s) received by the FHR will be dropped until the LHR has learned about the source and the SPT tree is built. That means for bursty sources or applications sensitive for the delivery of the first packet this proposal would not be very applicable. This proposal is mostly useful for applications that don't have strong dependency on the initial packet(s) and have a fairly constant data rate, like video distribution for example. For applications with strong dependency on the initial packet(s) we recommend using PIM Bidir [RFC5015] or SSM [RFC4607]. The protocol operations are much simpler compared to PIM SM, it will cause less churn in the network and both guarantee best effort delivery for the initial packet(s).

Another solution to address the problems described above is documented in [I-D.ietf-magma-msnip]. This proposal allows for a host to tell the FHR its willingness to act as Source for a certain Group before sending the data packets. LHRs have time to join the SPT tree before the host starts sending which would avoid packet loss. The SG mappings announced by [I-D.ietf-magma-msnip] can be advertised directly in SG messages, allowing a very nice integration of both proposals. The life time of the SPT is not driven by the liveliness of Multicast data packets (which is the case with PIM SM), but by the announcements driven via [I-D.ietf-magma-msnip]. This will also prevent packet loss due to bursty sources.

7. Resiliency to network partitioning

In a PIM SM deployment where the network becomes partitioned, due to link or node failure, it is possible that the RP becomes unreachable to a certain part of the network. New sources that become active in that partition will not be able to register to the RP and receivers within that partition are not able to receive the traffic. Ideally you would want to have a candidate RP in each partition, but you never know in advance which routers will form a partitioned network. In order to be fully resilient, each router in the network may end up being a candidate RP. This would increase the operational complexity of the network.

The solution described in this document does not suffer from that problem. If a network becomes partitioned and new sources become active, the receivers in that partitioned will receive the SG Mappings and join the source tree. Each partition works

independently of the other partition(s) and will continue to have access to sources within that partition. As soon as the network heals, the SG Mappings are re-flooded into the other partition(s) and other receivers can join to the newly learned sources.

8. Security Considerations

The security considerations are mainly similar to what is documented in [RFC5059]. It may be a concern that rogue devices can inject packets that are flooded throughout a domain. PFP packets SHOULD only be accepted from a PIM neighbor. Deployments may use mechanisms for authenticating PIM neighbors.

9. IANA considerations

This document requires the assignment of a new PIM Protocol type for the PIM Flooding Protocol (PFP). IANA is also requested to create a registry for PFP Types with type 0 allocated to "Source-Group Message". IANA is also requested to create a registry for PFP TLVs, with type 0 allocated to the "Source Group Holdtime" TLV. The allocation procedures are yet to be determined.

10. Acknowledgments

The authors would like to thank Arjen Boers for contributing to the initial idea and Yiqun Cai for his comments on the draft.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC5059] Bhaskar, N., Gall, A., Lingard, J., and S. Venaas, "Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)", RFC 5059, January 2008.

11.2. Informative References

- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, October 2007.

[I-D.ietf-magma-msnip] Fenner, B., Haberman, B., Holbrook, H., Kouvelas, I., and S. Venaas, "Multicast Source Notification of Interest Protocol (MSNIP)", draft-ietf-magma-msnip-06 (work in progress), March 2011.

Authors' Addresses

IJsbrand Wijnands
Cisco Systems, Inc.
De kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Stig Venaas
Cisco Systems, Inc.
Tasman Drive
San Jose CA 95134
USA

Email: stig@cisco.com

Michael Brig
Aegis BMD Program Office
17211 Avenue D, Suite 160
Dahlgren VA 22448-5148
USA

Email: michael.brig@mda.mil