

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 30, 2007

S. Venaas
UNINETT
February 26, 2007

ssmping Protocol
draft-venaas-mboned-ssmping-00

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/lid-abstracts.txt>.

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

This Internet-Draft will expire on August 30, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

ssmping is a tool that is used to check whether one can receive SSM, as well as obtaining some additional information. ssmping requires both a client and a server supporting the ssmping protocol to work. We here specify this protocol.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",

Internet-Draft

ssmping

February 2007

"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [1].

Table of Contents

1.	Introduction	3
2.	Protocol description	3
3.	Options	4
3.1.	Option format	4
3.2.	Defined Options	5
4.	Packet Format	7
5.	Acknowledgements	8
6.	IANA Considerations	8
7.	Security Considerations	8
8.	References	9
8.1.	Normative References	9
8.2.	Informative References	9
	Author's Address	9
	Intellectual Property and Copyright Statements	10

Internet-Draft

ssmping

February 2007

[1.](#) Introduction

ssmping is a tool that is used to check whether one can receive SSM, and it can also give other information like the time to establish the tree, number of router hops the packets have traveled, packet delay and loss. The ssmping functionality resembles ICMP echo request/reply using UDP and a client and a server that supports the ssmping protocol. It is used by a client to verify that it can receive multicast from the server, as well as some additional information. The protocol as specified here is based on an actual implementation of a tool [\[3\]](#) that has been found useful by many organisations.

[2.](#) Protocol description

Before going into the protocol details we will describe how it is used and what information it may provide. The typical usage is as follows. A server runs continuously in order to serve request from clients. At some point a client application may try to verify multicast reception from such a server. The client will need to know a unicast address of a server. The client joins an SSM channel (S,G) where S is a unicast address of the server, and G is a standardised multicast group for use by ssmping. After joining the channel, the client sends ssmping requests as UDP to a standardised ssmping port and the unicast address of the server. The requests are sent periodically, e.g. once per second, to the server. The requests contain a serial number, and typically a timestamp. The requests are typically, but not necessarily always, simply echoed back by the server. To each request, the server sends two replies. One as unicast back to the port and address the request was sourced from, and also as multicast back to the port the request came from. It is currently left open which port the request is sourced from, whether this port should be standardised or not. The TTL or Hop Limit of the replies are set to 64. The client should leave the SSM channel when it has finished its measurements.

By use of this protocol, a client can obtain information on several aspects of the multicast quality. First of all, by receiving unicast replies, it can verify that the server is receiving the unicast requests, is operational and responding. Hence provided that the client receives unicast replies, a failure in receiving multicast is indeed caused by a multicast problem. If it does receive multicast, it knows not only that it can receive, but it may get some information on how long it takes to establish the multicast tree (at least if it is in the range of seconds), whether there are packet drops, and the length and variation of round trip times (RTT). For unicast the RTT is the time from unicast request is sent to when the reply is received. For multicast we also talk about RTT, but then we

mean from the unicast request is sent to when the multicast reply is received. Since the server sets TTL or Hop Limit to 64, it can also know the number of router hops it is away from the source. By comparing with the unicast replies, it can see whether there are differences in RTT and number of hops etc for unicast and multicast. Provided that the server sends the unicast and multicast replies nearly simultaneously, it may also be able to measure difference in one way delay for unicast and multicast on the path from server to client, and also if there are differences in delay variation. Servers may optionally specify a timestamp. This may be useful if the unicast and multicast replies can not be sent nearly simultaneously, or if the client and server have synchronised clocks.


The ssmping requests and replies have a common format, one octet specifying the message type, followed by a number of options in TLV (Type, Length and Value) format. This makes the protocol easily extendible. Generally the client includes a number of options in the request, and a server may simply echo the content back (only changing the message type), without inspecting the options. However, there are a number of options that a server implementation may support, where the client may ask for a certain information or behaviour from the server. In some cases the server will need to add options in the response. The response will then first contain the exact options from the request, and then right after those, options appended by the server.

[3. Options](#)

There are a number of different options. Most of the options are only used by clients and simply echoed back by the server, where the server doesn't care about their contents. There are however some client options that the server may care about. There are also server options that may be requested by the client. Generally a simple client will only include a few options, and get exactly the same options and values echoed back. Strictly speaking the protocol could work without any options. Without sending any options a client would still be able to tell whether multicast is working or not, however with the use of some of the basic options a client can obtain a lot more information.

[3.1.](#) Option format

All options are TLVs formatted as specified below.



Type (2 octets) specifies the option. The different options are defined below.

Length (2 octets) specifies the length of the value. Depending on the option type it can be from 0 to 65535.

Value. The value must always be of the specified length. See the respective option definitions for possible values. If the length is 0, the value field is not included.

[3.2.](#) Defined Options

Client Identifier, type 1. Length MUST be non-zero. Only used by clients. A client SHOULD include this. The client may use any value it likes to be able to detect whether a reply is a reply to this query or not. A server should treat this as opaque data, and simply leave it unchanged in the reply. The value might be a process ID, perhaps process ID combined with an IP address because it may receive multicasted responses to queries from other clients. It is left to the client implementor how to make use of this.

Sequence number, type 2. Length MUST be 4. Only used by clients. A client SHOULD include this. This contains a 32 bit sequence number. The values would typically start at 1 and increase by one for each request in a sequence.

Timestamp, type 3. Length MUST be 8 bytes. A client SHOULD include this. A server MAY support this. If supported it SHOULD be included in the reply if requested by the client. The timestamp specifies the time when the message (query or reply) is sent. The first 4 bytes specify the number of seconds since the Epoch (beginning of the year 1970). The next 4 bytes specify the number of microseconds since the last second since the Epoch.

Multicast group, type 4. Length MUST be greater than 1. It is optional for clients and servers to support this. It allows a client to specify which group the server should send to. This is currently

used by a tool called "asmping" to test ASM connectivity. The server may have restrictions on which groups can be used. The format of the option value is as below.

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Addr Family | Multicast group address...                          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
                                     ....

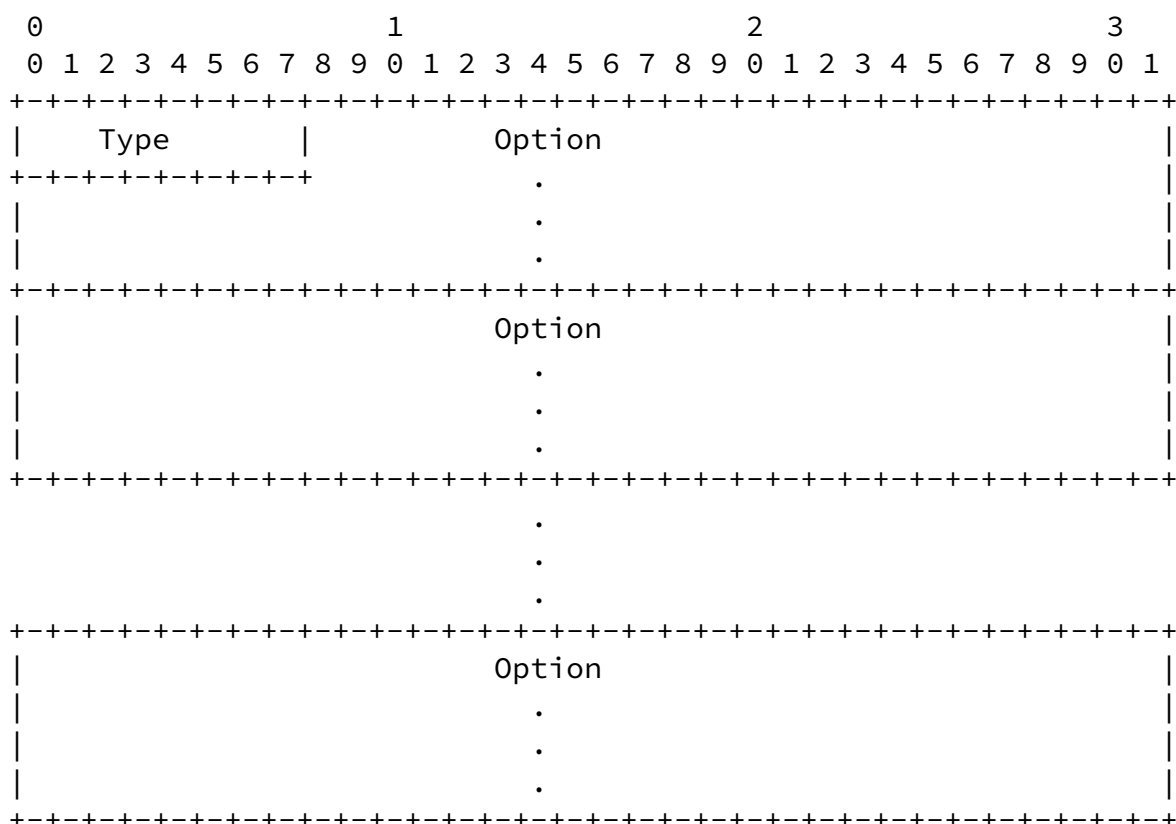
```

The address family is a value 0-127 as assigned by IANA for Internet Address Families [2]. This is followed by the group address. For IPv4 the option value length will be 5, for IPv6 17.

a reply that is larger than what the server normally would send. The addition of this option consumes a minimum of 3 octets, so it should only be added if the requested size is at least 3 octets more than the size of the normal (non-padded) response.

4. Packet Format

The format of the ssm ping messages is a one octet message type, followed by a variable number of options.



There are two message types defined. Type 81 (the character Q in ASCII) specifies a query. Type 65 (the character A in ASCII) specifies a response (answer).

The options follow right after the type octet and are not aligned in any way (no spacing or padding). I.e., options might start at any

5. Acknowledgements

The ssm ping idea was proposed by Pavan Namburi, Kamil Sarac and Kevin C. Almeroth in the paper SSM-Ping: A Ping Utility for Source Specific Multicast, and also the Internet Draft [draft-sarac-mping-00.txt](#). Mickael Hoerdt has contributed with several ideas. Alexander Gall, Nick Lamb and Dave Thaler have contributed in different ways to my implementation of the ssm ping tools [3]. Hugo Santos has made an independent implementation of an ssm ping server. Many people in communities like TERENA, Internet2 and the M6Bone have used early implementations of ssm ping and provided feedback that have influenced the current protocol. Thanks to Olav Kvittem, Kamil Sarac and Trond Skjesol for reviewing and providing feedback on this draft.

6. IANA Considerations

As currently specified, ssm ping would need a well known port number which the servers listen to. It might be desirable to use SRV records instead or in addition to this. For IPv6 SSM ssm ping should ideally have a reserved group ID. For the optional ASM functionality it would be useful to have a reserved IPv6 group ID, this may be the same as the one used for SSM. It may also be useful to have a dedicated group for the optional IPv4 ASM functionality. This section needs further work.

7. Security Considerations

There are some security issues to consider. One is that a host may send a request with an IP source address of another host, and make a random ssm ping server on the Internet send packets to this other host. This is fairly harmless. The worst case is if the host receiving the unicast replies also happen to be performing an ssm ping test towards that particular server. In this unlikely event there would be an amplification effect where the host receives twice as many replies as there are requests sent. An ssm ping server should perform rate limiting, to guard against this being used as an DoS attack. A client should also use the client identifier option to be able to distinguish replies to its own requests from replies that might be to other requests. How the protocol should be designed to cope with rate limiting at the server requires further study. One possibility might be that the server can choose to send generic replies, e.g. a packet every second without the usual client options but including sequence number and server time stamp, and where

clients do not send requests as long as they receive generic replies.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] "IANA, Address Family Numbers",
<<http://www.iana.org/assignments/address-family-numbers>>.

8.2. Informative References

- [3] "ssmping implementation",
<<http://www.venaas.no/multicast/ssmping/>>.

Author's Address

Stig Venaas
UNINETT
Trondheim NO-7465
Norway

Email: venaas@uninett.no

Internet-Draft

ssmping

February 2007

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at

ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

Venaas

Expires August 30, 2007

[Page 10]