Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 8, 2009

M. Petit-Huguenin
8x8, Inc.

July 7, 2008

Path MTU Discovery Using Session Traversal Utilities for NAT (STUN)
draft-petithuguenin-behave-stun-pmtud-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/1id-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 January 8, 2009.

Abstract

This document describes a Session Traversal Utilities for NAT (STUN) usage for discovering the Path MTU between a client and a server.

Table of Contents

1. Introduction

The Packetization Layer Path MTU Discovery specification [RFC4821] describes a method to discover the PMTU but does not describe a practical protocol to discover the Path MTU when using UDP.

This document only describe how the probing mechanism is implemented with STUN. The algorithm to find the Path MTU is described in [RFC4821].

The probing mechanism is implemented by sending a Probe Request with a PADDING [I-D.ietf-behave-nat-behavior-discovery] attribute and the DF bit set over UDP. A router on the path to the server can reject this request with an ICMP message or drop it. The STUN retransmission algorithm is modified so the third and next retransmissions do not include the PADDING attribute, so they are not dropped by an intermediate router. The server responds by indicating if the request received contained the PADDING attribute or not. This permits to quickly find if a probe packet if too big for the Path MTU or not.

2. 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].

3. Probe Mechanism

A client MUST NOT send a Probe Request if it does not have knowledge that the server supports this specification. This is done by an external mechanism which is specific to each UDP protocol. Section 4 describes some of this mechanisms.

The probe mechanism is used to measure the Path MTU in one direction only, from the client to the server.

3.1. Sending a Probe Request

A client forms a Probe Request by following the rules in [I-D.ietf-behave-rfc3489bis] section 7.1. No authentication method is used. The client adds a PADDING [I-D.ietf-behave-nat-behavior-discovery] attribute with a size that, when added to the IP and UDP headers and the other STUN components, is equal to the Selected Probe Size, as defined in [RFC4821] section 7.3. If the IP address and port tuple used as destination for the

Probe Request is also used by another protocol then the client MUST add the FINGERPRINT attribute.

Then the client sends the Probe Request to the server over UDP. The UDP retransmission mechanism described in [I-D.ietf-behave-rfc3489bis] section 7.2.1 is modified so that if the Probe Request has to be retransmitted three times or more then it is stripped of its PADDING attribute before been sent. The UDP packets MUST be sent with the DF bit set.

3.2. Receiving a Probe Request

A server receiving a Probe Request MUST process it as specified in [I-D.ietf-behave-rfc3489bis]. The server MUST NOT challenge the client.

The server then creates a Probe Response. If the Probe Request contains a PADDING attribute, then a PADDING-RECEIVED attribute is added to the response, with a value equals to the size of the PADDING attribute received. If the IP address and port used to send the Probe Request is also used by another protocol, then the server MUST add the FINGERPRINT attribute. The server then sends the response to the client.

3.3. Receiving a Probe Response

A client receiving a Probe Response processes it as specified in [I-D.ietf-behave-rfc3489bis]. If the response contains a PADDING-RECEIVED attribute, then this is interpreted as a Probe Success as defined in [RFC4821] section 7.6.1. If an ICMP packet "Fragmentation needed" is received or if the response does not contain a PADDING-RECEIVED attribute, then this is interpreted as a Probe Failure as defined in [RFC4821] section 7.6.2. If the Probe transactions fails in timeout, then this is interpreted as a Probe Inconclusive as defined in [RFC4821] section 7.6.4.

4. Probe Support Discovery Mechanisms

4.1. Implicit Mechanism

An endpoint acting as a client for the STUN usage described in this specification MUST also act as a server for this STUN usage. This means that a server receiving a Probe Request can assumes that it can acts as a client to discover the Path MTU to the IP address and port from which it received the Probe Request.

4.2. Probe Support Discovery with TURN

A TURN client supporting this STUN usage will add a PMTUD-SUPPORTED attribute to the Allocate Request sent to the TURN server. The TURN server can immediately start to send Probe Requests to the TURN client on reception of an Allocation Request with a PMTUD-SUPPORTED attribute. The TURN client will then use the Implicit Mechanism described above to send probes.

4.3. Probe Support Discovery with ICE

An ICE [I-D.ietf-mmusic-ice] client supporting this STUN usage will add a PMTUD-SUPPORTED attribute to the Binding Request sent during a connectivity check. The ICE server can immediately start to send Probe Requests to the ICE client on reception of a Binding Request with a PMTUD-SUPPORTED attributed. The ICE client will then use the Implicit Mechanism described above to send probes.

5. New STUN Method

This specification defines one new STUN method:

```
Request/Response Transaction
   0x801 : Probe
```

6. New STUN Attributes

This specification defines the following new STUN attributes:

```
   0x4001 : PADDING-RECEIVED
   0xC002 : PMTUD-SUPPORTED
```

6.1. PADDING-RECEIVED

The PADDING-RECEIVED attribute contains the size of the PADDING attribute received. It is a 16-bit unsigned integer, followed by two reserved bytes which MUST be set to 0 on transmission and MUST be ignored on reception.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Padding size           |         Reserved              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

6.2. PMTUD-SUPPORTED

The PMTUD-SUPPORTED attribute is used in STUN usages and extensions to signal the support of this specification. This attribute has no content.

7. Security Considerations

TBD

8. IANA Considerations

TBD

9. Normative References

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

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.

[I-D.ietf-behave-rfc3489bis]
Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)", draft-ietf-behave-rfc3489bis-16 (work in progress), July 2008.

[I-D.ietf-behave-nat-behavior-discovery]
MacDonald, D. and B. Lowekamp, "NAT Behavior Discovery Using STUN", draft-ietf-behave-nat-behavior-discovery-03 (work in progress), February 2008.

[I-D.ietf-mmusic-ice]
Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", draft-ietf-mmusic-ice-19 (work in progress), October 2007.

Author's Address

Marc Petit-Huguenin
8x8, Inc.
3151 Jay Street
Santa Clara, CA 95054
US

Phone: +1 408 654 0875
Email: marc@8x8.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.