

Intended Status: Best Current Practice  
BEHAVE WG  
Internet Draft  
Expires: December 30, 2007

P. Srisuresh  
Kazeon Systems  
B. Ford  
M.I.T.  
S. Sivakumar  
Cisco Systems  
S. Guha  
Cornell U.  
June 30, 2007

NAT Behavioral Requirements for ICMP protocol  
<[draft-ietf-behave-nat-icmp-04.txt](#)>

#### 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>.

#### Abstract

This document specifies the behavioral properties required of the Network Address Translator (NAT) devices in conjunction with the ICMP protocol. The objective of this memo is to make NAT devices more predictable and compatible with diverse application protocols that traverse the devices. Companion documents provide behavioral recommendations specific to TCP, UDP and other protocols.

## Table of Contents

1.	Introduction and Scope .....	
2.	Terminology .....	
3.	ICMP Query Handling .....	
3.1.	ICMP Query Mapping .....	
3.2.	ICMP Query Session Timeouts .....	
4.	ICMP Error Forwarding .....	
4.1.	ICMP Error Payload Validation .....	
4.2.	ICMP Error Packet Translation .....	
4.2.1.	ICMP Error Packet Received from External Realm ....	
4.2.2.	ICMP Error Packet Received from Private Realm .....	
4.3.	NAT Sessions Pertaining to ICMP Error Payload .....	
5.	Hairpinning Support for ICMP packets .....	
6.	Rejection of Outbound Flows Disallowed by NAT .....	
7.	Conformance to <a href="#">RFC 1812</a> .....	
7.1.	IP packet fragmentation .....	
7.1.1.	Generating "Packet too Big" ICMP error Message ....	
7.1.2.	Forwarding "Packet too big" ICMP Error Message ....	
7.2.	Generating "Time Exceeded" Error Message .....	
7.3.	<a href="#">RFC 1812</a> Conformance Requirements summary .....	
8.	Summary of Requirements .....	
9.	Security Considerations .....	
10.	IANA Considerations .....	
11.	Acknowledgements .....	

## 1. Introduction and Scope

As pointed out in [RFC 3424](#) [[UNSAF](#)], NAT implementations vary widely in terms of how they handle different traffic. The purpose of this document is to define a specific set of requirements for NAT behavior with regard to ICMP messages. The objective is to reduce the unpredictability and brittleness the NAT devices (NATs) introduce. This document is an adjunct to [[BEH-UDP](#)], [[BEH-TCP](#)], and other protocol-specific BEHAVE document(s) in the future which define requirements for NATs when handling protocol-specific traffic.

The requirements of this specification apply to Traditional NATs as described in [[NAT-TRAD](#)]. Traditional NAT has two variations, namely,

Basic NAT and Network Address Port Translator (NAPT). Of these, NAPT is by far the most commonly deployed NAT device. NAPT allows multiple private hosts to share a single public IP address simultaneously.

This document only covers the ICMP aspects of NAT traversal,

specifically the traversal of ICMP Query Messages and ICMP Error messages. Traditional NAT inherently mandates a certain level of firewall-like functionality. However, firewall functionality in general or any other middlebox functionality is out of the scope of this document.

In some cases, ICMP Message traversal behavior on a NAT device may be overridden by local administrative policies. Some administrators may choose to entirely prohibit forwarding of ICMP error messages across a NAT device. Some others may choose to prohibit ICMP query based applications across a NAT device. These are local policies and not within the scope of this document. For this reason, some of the ICMP behavior requirements listed in the document are preceded with a constraint of local policy permitting.

This document focuses strictly on the behavior of the NAT device, and not on the behavior of applications that traverse NATs. A separate document [[BEH-APP](#)] provides recommendations for application designers on how to make applications work robustly over NATs that follow the behavioral requirements specified here and the adjunct BEHAVE documents.

Per [[RFC1812](#)], ICMP is a control protocol that is considered to be an integral part of IP, although it is architecturally layered upon IP - it uses IP to carry its data end-to-end. As such, many of the ICMP behavioral requirements discussed in this document apply to all IP protocols.

In case a requirement in this document conflicts with protocol-specific BEHAVE requirement(s), protocol specific BEHAVE documents will take precedence. The authors are not aware of any conflicts between this and any other IETF document at the time of this writing.

[Section 2](#) describes the terminology used throughout the document.

Sections [3](#) is focused on requirements concerning ICMP query based applications traversing a NAT device. Sections [4](#) and [5](#) describe requirements concerning ICMP error messages traversing a NAT device. Sections [6](#) and [7](#) describe requirements concerning ICMP error messages generated by a NAT device. [Section 8](#) summarizes all the requirements in one place. [Section 9](#) has a discussion on security considerations.

## [2](#). Terminology

Definitions for majority of the NAT terms used throughout the document may be found in [[NAT-TERM](#)] and [[BEH-UDP](#)]. The term

Srisuresh, et. al.

[Page 3]

---

Internet-Draft      NAT Behavioral Requirements for ICMP

June 2007

"NAT Session" is adapted from [[NAT-MIB](#)] and denotes the entity within a NAT device that represents the translation glue for a session traversing the NAT device.

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](#)].

[Section 3.2.2 of \[RFC1122\]](#) broadly groups ICMP messages into two classes, namely "ICMP Query" messages and "ICMP Error" messages. In addition, there are ICMP messages, created after [[RFC1122](#)] that do not fall under either category. We refer them as "Post-1122 ICMP Messages". All three ICMP message classes are described as follows.

ICMP Query Messages - ICMP query messages are characterized by an Identifier field in the ICMP header. The Identifier field used by the ICMP Query messages is also referred as "Query Identifier" or "Query Id", for short throughout the document. A Query Id is used by query senders and responders as the equivalent of a TCP/UDP port to identify an ICMP Query session.

ICMP Error Messages - ICMP error messages provide signaling for IP. All ICMP error messages are characterized by the fact that they embed the original datagram that triggered the ICMP error message. The original datagram embedded within the ICMP Error payload is also referred as "Embedded packet", throughout the document. Unlike ICMP Query messages, ICMP error messages do not have a Query Id in the ICMP header.

Post-1122 ICMP Messages - ICMP messages that do not fall under either of the above two classes are referred as "Post-1122 ICMP Messages" throughout the document. For example, discovery ICMP messages ([\[RFC1256\]](#)) are "request/response" type ICMP messages. However, they are not characterized as ICMP Query Messages as they do not have an "Identifier" field within the messages. Likewise, there are other ICMP messages defined in [\[RFC4065\]](#) that do not fall in either of ICMP Query or ICMP Error message categories, but will be referred as Post-1122 ICMP error messages. A NAT MAY drop or appropriately handle a post-1122 ICMP messages.

### [3.](#) ICMP Query Handling

This section lists the behavioral requirements for a NAT device when processing ICMP Query packets. The following sub sections discuss requirements specific to ICMP Query handling in detail.

#### [3.1.](#) ICMP Query Mapping

Srisuresh, et. al.

[Page 4]

---

Internet-Draft      NAT Behavioral Requirements for ICMP

June 2007

A NAT device MUST permit ICMP query based applications to be initiated from private hosts to the external hosts. ICMP Query mapping by NAT devices is necessary for current ICMP query based applications to work. Specifically, a NAT device MUST transparently forward any ICMP Query packets initiated from the nodes behind NAT devices and the responses to these Query packets in the opposite direction. As specified in [\[NAT-TRAD\]](#), this requires translating the IP header. A NAT device further translates the ICMP Query Id and the associated checksum in the ICMP header prior to forwarding.

The mapping of ICMP Query identifier within the NAT device SHOULD be external endpoint independent. Say, an internal host A sent an ICMP query out to an external host B using Query Id X. And, say, the NAT assigned this an external mapping of Query id X' on the NAT's public address. If host A reused the Query Id X to send ICMP queries to the same or different external host, the NAT device SHOULD reuse the same Query Id mapping (i.e., map private host's Query id X to Query id X' on NAT's public IP address) instead of assigning a different mapping. This is similar to the "endpoint independent mapping" requirement specified in the TCP and UDP

BEHAVE requirements [[BEH-TCP](#), [BEH-UDP](#)].

Below is justification for making the endpoint independent mapping for ICMP query IDs a SHOULD [[RFC2119](#)] requirement. ICMP Ping ([RFC1470](#)) and ICMP traceroute ([MS-TRCRT](#)) are two most commonly known legacy applications built on top of ICMP query messages. Neither of these applications require the ICMP Query Id to be retained across different sessions with external hosts. But, that may not be case with future applications. In the future, when an end host application reuses the same Query identifier in sessions with different target hosts, the end host application might require that the endpoint identity (i.e., the tuple of IP address and Query Identifier) appears the same across all its target hosts. Such a requirement will be valid to make in an IP network without NAT devices. When NAT devices enforce endpoint mapping that is external host independent, the above assumption will be valid to make even in a world with NAT devices. Given the dichotomy between legacy applications not requiring endpoint independent mapping and future applications that might require it, the requirement level is kept at SHOULD [[RFC2119](#)].

REQ-1: When local policy permits, a NAT device MUST permit ICMP queries and their associated responses, when the query is initiated from a private host.

a) NAT mapping of ICMP Query identifiers SHOULD be external host independent.

### [3.2](#). ICMP Query Session Timeouts

When an application initiates an ICMP query that transits a NAT device, the NAT associates a timer to the ICMP query NAT session. This is so the ICMP query NAT session is freed up if the NAT session remains idle for longer than the timeout set by the timer. Query response times can vary. ICMP query based application are primarily request/response driven. The ICMP Query session timeout requirement is necessary for current ICMP query applications to work.

Ideally, the timeout should be set to Maximum Round Trip Time (Maximum RTT). For the purposes of constraining the maximum RTT, the Maximum Segment Lifetime (MSL), defined in [[RFC793](#)] could be considered a guideline to set packet lifetime. Per [[RFC793](#)], MSL is

the maximum amount of time a TCP segment can exist in a network before being delivered to the intended recipient. This is the maximum duration an IP packet can be assumed to take to reach the intended destination node before declaring that the packet will no longer be delivered. For an application initiating ICMP query message and waiting for a response for the query, the Maximum RTT could in practice be constrained to be sum total of MSL for the Query message and MSL for the response message. In other words, Maximum RTT could be constrained to no more than 2x MSL. The recommended value for MSL in [\[RFC793\]](#) is 120 seconds, even though several implementations set this to 60 seconds or 30 seconds. When MSL is 120 seconds, the Maximum RTT (2x MSL) would be 240 seconds.

In practice, ICMP Ping ([\[RFC1470\]](#)) and ICMP traceroute ([\[MS-TRCRT\]](#)), the two most commonly known legacy applications built on top of ICMP query messages take less than 10 seconds to complete a round trip, when the target node is operational on the network.

Setting the ICMP NAT session timeout to a very large duration (say, 240 seconds) could potentially tie up precious NAT resources such as query mappings and NAT Sessions for the whole duration. On the other hand, setting the timeout very low can result in premature freeing of NAT resources and applications failing to complete gracefully. The ICMP Query session timeout needs to be a balance between the two extremes. 60 seconds timeout is a balance between the two extremes. ICMP query session timer MUST not expire in less than 60 seconds. We RECOMMEND however that the administrator(s) be allowed to configure the timer.

REQ-2: An ICMP Query session timer MUST NOT expire in less than 60 seconds.

a) It is RECOMMENDED that the ICMP Query session timer be made configurable.

#### [4.](#) ICMP Error Forwarding

Many applications make use of ICMP error messages from end hosts and intermediate devices to shorten application timeouts. Some applications will not operate correctly without the receipt of ICMP error messages. The following sub-sections discuss the requirements

a NAT device MUST conform to in order to ensure reliable forwarding.

#### 4.1. ICMP Error Payload Validation

[Appendix C](#) of [\[ICMP-ATK\]](#) points out that newer revision end host TCP stacks do not accept ICMP error messages with a mismatched IP or TCP checksum in the embedded packet, if the embedded datagram contains full IP packet and the TCP checksum can be calculated. Whenever validation is possible, NAT devices should ensure that ICMP Error payload is not corrupted. Only after the payload is validated, should the NAT proceed to forward the ICMP error packet. This requirement is meant primarily for future applications. Current applications may not be checking for mismatched checksum.

If the IP checksum of the embedded packet does not validate, the NAT device SHOULD simply drop the error packet. [\[ICMP\]](#) stipulates that an ICMP error message should embed IP header and a minimum of 64 bits of the IP payload. [Section 4.3.2.3 of \[RFC1812\]](#) further recommends that an ICMP error originator SHOULD include as much of the original packet as possible in the payload without the length of the ICMP datagram exceeding 576 bytes. If the embedded packet is a complete IP packet, including the entire transport segment, and the transport protocol of the embedded packet requires the recipient to validate the checksum, the NAT device SHOULD validate the transport checksum. If the transport protocol is UDP and the checksum is set to zero, the UDP protocol does not require the recipient to validate the UDP checksum. In the case the ICMP Error payload includes ICMP extensions ([\[ICMP-EXT\]](#)), the NAT device SHOULD exclude the optional zero-padding and the ICMP extensions when evaluating transport checksum for the embedded packet. If the transport checksum fails, the NAT device SHOULD drop the error packet. Readers are urged to refer [\[ICMP-EXT\]](#) for identifying the presence of ICMP extensions in an ICMP message.

When the IP packet embedded within the ICMP error message includes IP options, the NAT device MUST NOT assume that the transport header of the embedded packet is at a fixed offset (as would be the case when there are no IP options associated with the packet) from the start of the embedded packet. Specifically, the NAT device SHOULD index past all IP options when locating the start of transport header for the embedded packet.

REQ-3: When an ICMP error packet is received, the NAT device SHOULD do the following.

- a) If the IP checksum of the embedded packet fails to validate, drop the error packet; and
- b) If the embedded packet includes IP options, traverse past the IP options to locate the start of transport header for the embedded packet; and
- c) If the ICMP Error payload contains ICMP extensions([\[ICMP-EXT\]](#)), exclude the optional zero-padding and the ICMP extensions when evaluating transport checksum for the embedded packet; and
- d) If the embedded packet contains the entire transport segment, and the transport protocol of the embedded packet requires the recipient to validate the transport checksum, and the checksum fails to validate, drop the error packet.

## 4.2. ICMP Error Packet Translation

[Section 4.3 of RFC 3022](#) describes the fields of an ICMP error message that a NAT device translates. In this section, we describe the requirements a NAT device MUST conform to while performing the translations. Requirements identified in this section are necessary for the current applications to work correctly.

Consider the following scenario in figure 1. Say, NAT-xy is a NAT device connecting hosts in private and external networks. Router-x and Host-x are in the external network. Router-y and Host-y are in the private network. The subnets in the external network are routable from the private as well as the external domains. By contrast, the subnets in the private network are only routable within the private domain. When Host-y initiated a session to Host-x, let us say that the NAT device mapped the endpoint on Host-y into Host-y' in the external network. The following subsections describe the processing of ICMP error messages on the NAT device(NAT-xy), when the NAT device receives an ICMP error message in response to a packet pertaining to this session.

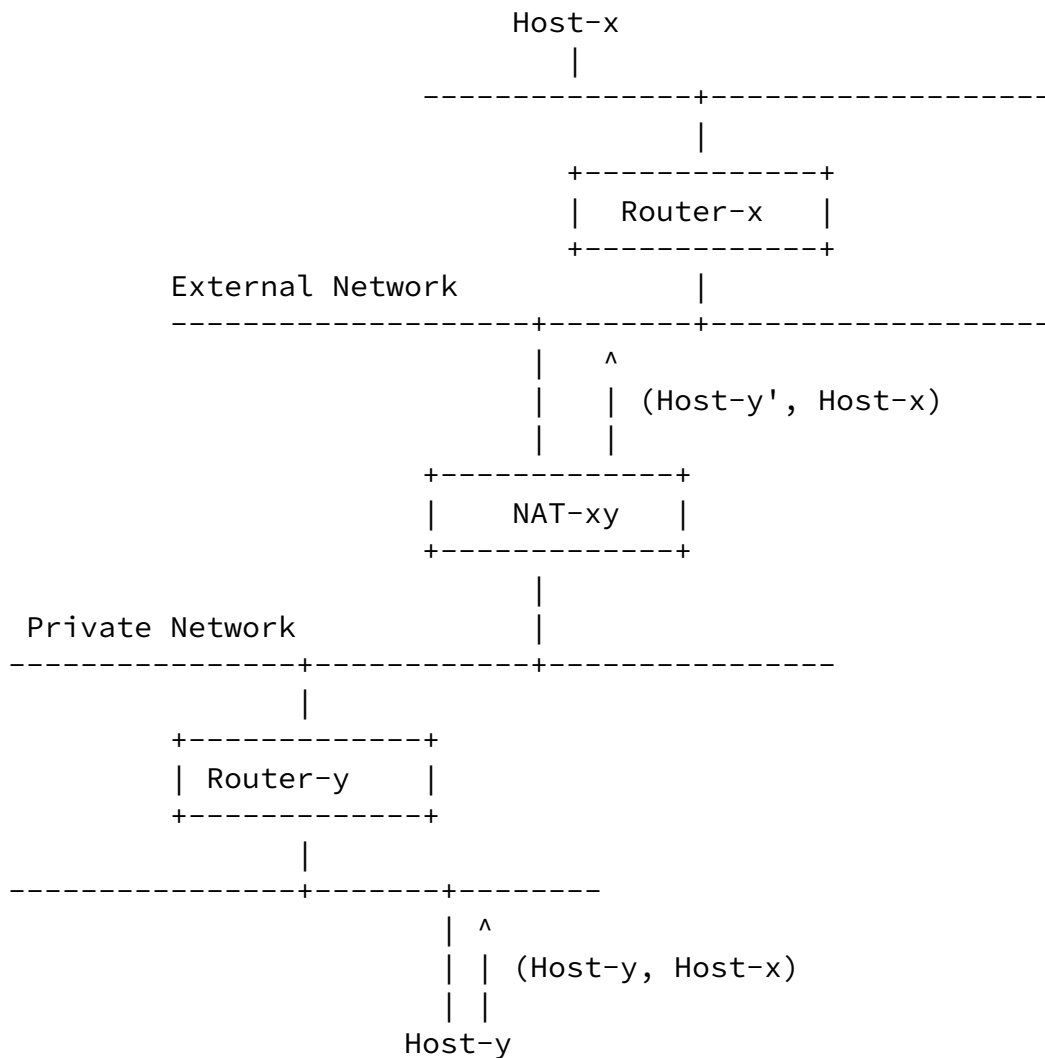


Figure 1. A Session from a private host traversing a NAT device.

#### [4.2.1.](#) ICMP Error Packet Received from External Realm

Say, a packet from Host-y to Host-x triggered an ICMP error message from one of Router-x or Host-x (both of which are in the external domain). Such an ICMP error packet will have one of Router-x or Host-x as the source IP address and Host-y' as the destination IP address as described in figure 2 below.

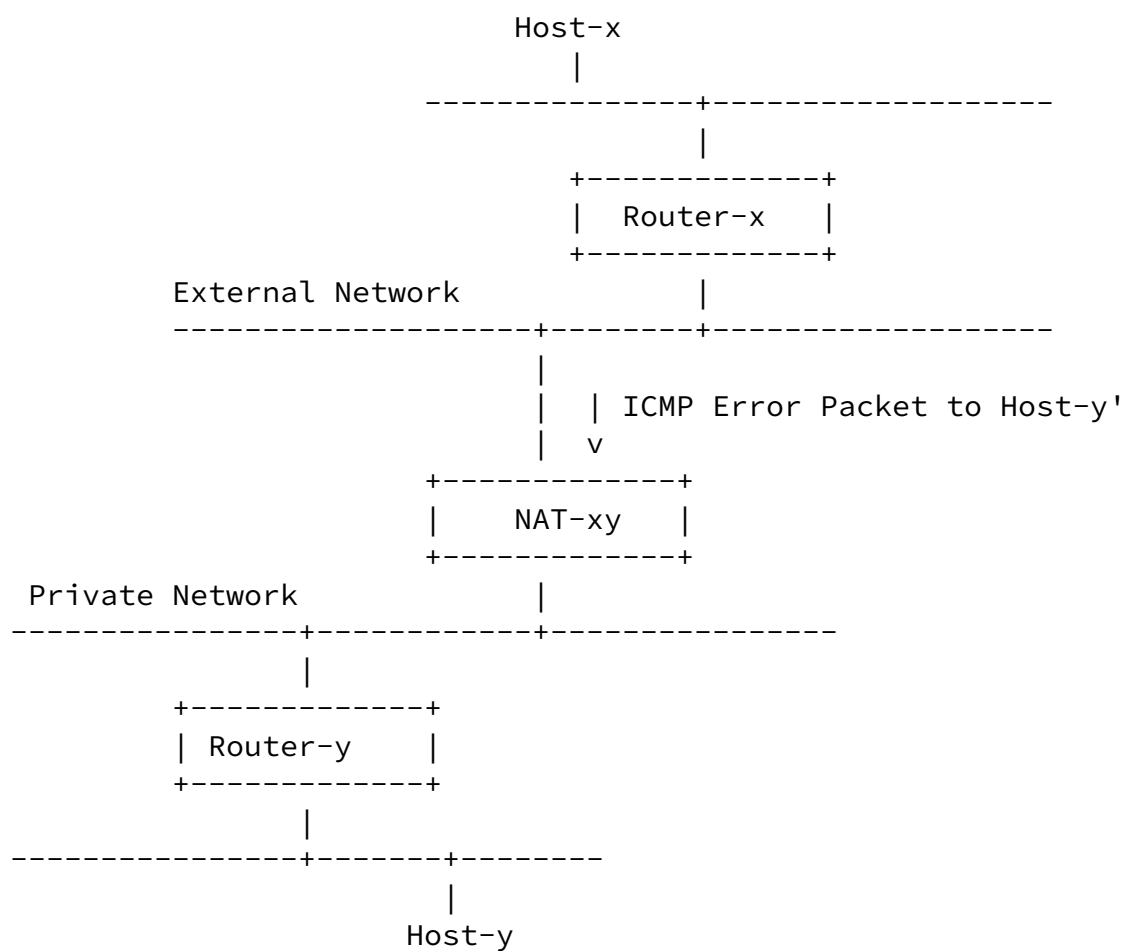


Figure 2. ICMP error Packet Received from External Network

When the NAT device receives the ICMP error packet, the NAT device MUST use the packet embedded within the ICMP error message (i.e., the IP packet from Host-y' to Host-x) to look up the NAT Session the embedded packet belongs to. If the NAT device does not have an active mapping for the embedded packet, the NAT SHOULD silently drop the ICMP error packet. Otherwise, the NAT device MUST use the matching NAT Session to translate the embedded packet. That is,

translate the source IP address of the embedded packet (e.g., Host-y' -> Host-y) and transport headers.

In addition, if the ICMP Error payload contains ICMP extensions ([[ICMP-EXT](#)]), the NATs are encouraged to support ICMP extension objects. At the time of this writing, the authors are not aware of any standard ICMP extension objects containing realm specific information.

The NAT device MUST also use the matching NAT Session to translate the destination IP address in the outer IP header. In the outer header, the source IP address will remain unchanged because the

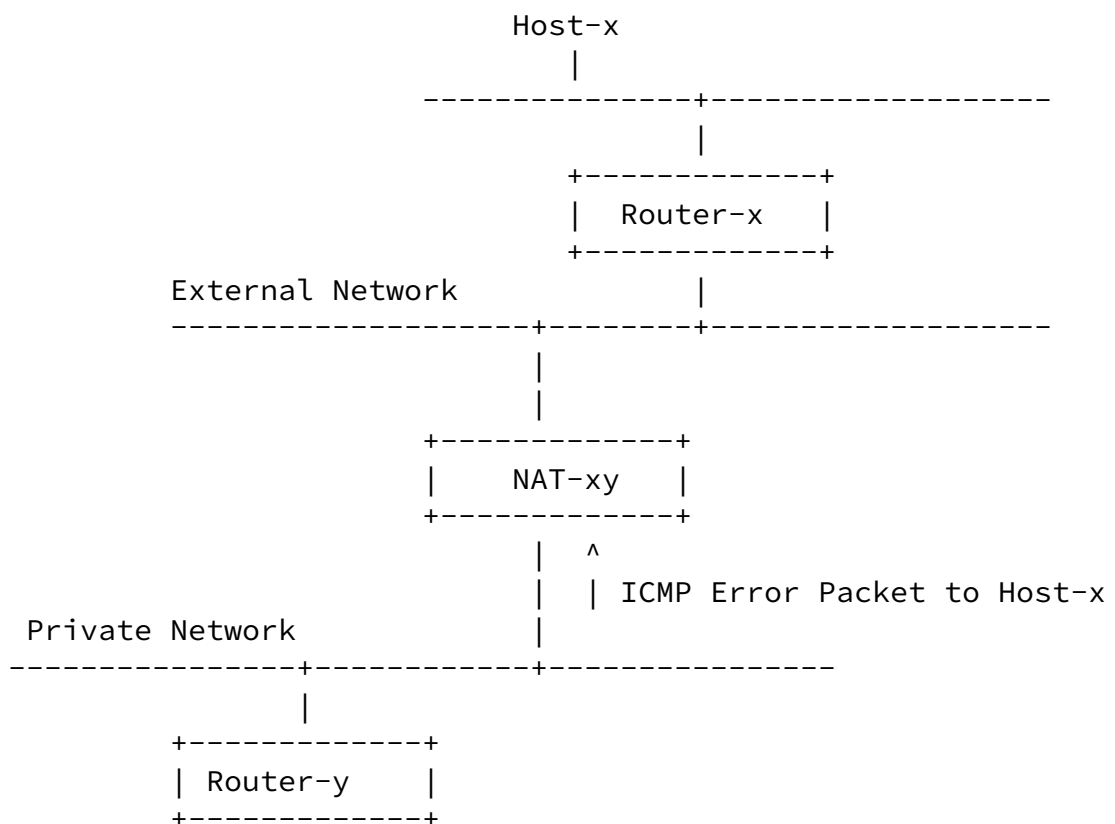
originator of the ICMP error message (Host-x or Router-x) is in external domain and routable from the private domain.

REQ-4: If a NAT device receives an ICMP error packet from external realm, and the NAT does not have an active mapping for the embedded payload, the NAT SHOULD silently drop the ICMP error packet. If the NAT has active mapping for the embedded payload and local policy permits, then the NAT MUST do the following prior to forwarding the packet.

- a) Revert the IP and transport headers of the embedded IP packet to their original form, using the matching mapping; and
- b) Leave the ICMP error type and code unchanged; and
- c) Modify the destination IP address of the outer IP header to be same as the source IP address of the embedded packet after translation.

#### [4.2.2](#). ICMP Error Packet Received from Private Realm

Now, say, a packet from Host-x to Host-y triggered an ICMP error message from one of Router-y or Host-y (both of which are in the private domain). Such an ICMP error packet will have one of Router-y or Host-y as the source IP address and Host-x as the destination IP address as specified in figure 3 below.



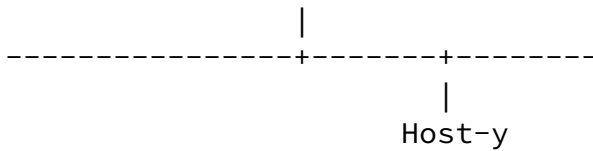


Figure 3. ICMP Error Packet Received from Private Network

When the NAT device receives the ICMP error packet, the NAT device MUST use the packet embedded within the ICMP error message (i.e., the IP packet from Host-x to Host-y) to look up the NAT Session the embedded packet belongs to. If the NAT device does not have an active mapping for the embedded packet, the NAT SHOULD silently drop the ICMP error packet. Otherwise, the NAT device MUST use the matching NAT Session to translate the embedded packet.

In addition, if the ICMP Error payload contains ICMP extensions ([[ICMP-EXT](#)]), the NATs are encouraged to support ICMP extension objects. At the time of this writing, the authors are not aware of any standard ICMP extension objects containing realm specific information.

In the outer header, the destination IP address will remain unchanged, as the IP addresses for Host-x is already in the external domain. If the ICMP error message is generated by Host-y, the NAT device must simply use the NAT Session to translate the source IP

address Host-y to Host-y'. If the ICMP error message is originated by the intermediate node Router-y, translation of the source IP address varies depending on whether Basic NAT or NAPT function ([[NAT-TRAD](#)]) is enforced by the NAT device. A NAT device enforcing Basic NAT function has a pool of public IP addresses and enforces address mapping (which is different from the endpoint mapping enforced by NAPT) when a private node initiates an outgoing session via the NAT device. So, if the NAT device has active mapping for the IP address of the intermediate node Router-y, the NAT device MUST translate the source IP address of the ICMP error packet with the public IP address in the mapping. In all other cases, the NAT device MUST simply use its own IP address in the external domain to translate the source IP address.

REQ-5: If a NAT device receives an ICMP error packet from private

realm, and the NAT does not have an active mapping for the embedded payload, the NAT SHOULD silently drop the ICMP error packet. If the NAT has active mapping for the embedded payload and local policy permits, then the NAT MUST do the following prior to forwarding the packet.

- a) Revert the IP and transport headers of the embedded IP packet to their original form, using the matching mapping; and
- b) Leave the ICMP error type and code unchanged; and
- c) If the NAT enforces Basic NAT function ([\[NAT-TRAD\]](#)), and the NAT has active mapping for the IP address that sent the ICMP error, translate the source IP address of the ICMP error packet with the public IP address in the mapping. In all other cases, translate the source IP address of the ICMP error packet with its own public IP address.

#### [4.3.](#) NAT Sessions Pertaining to ICMP Error Payload

While processing an ICMP error packet pertaining to an ICMP Query or Query response message, a NAT device MUST NOT refresh or delete the NAT Session that pertains to the embedded payload within the ICMP error packet. This is in spite of the fact that the NAT device uses the NAT Session to translate the embedded payload. This ensures that the NAT Session will not be modified if someone is able to spoof ICMP error messages for the session. [\[ICMP-ATK\]](#) lists a number of potential ICMP attacks that may be attempted by malicious users on the network. This requirement is necessary for current applications to work correctly.

REQ-6: While processing an ICMP error packet pertaining to an ICMP Query or Query response message, a NAT device MUST NOT refresh or delete the NAT Session that pertains to the embedded payload within the ICMP error packet.

## [5.](#) Hairpinning Support for ICMP packets

[\[BEH-UDP\]](#) and [\[BEH-TCP\]](#) mandate support for hairpinning for UDP and TCP sessions respectively on NAT devices. A NAT device needs to support hairpinning for ICMP Query sessions as well. Specifically, ICMP query hairpinning MUST be supported on Basic NATs. Say, for example, individual private hosts register their NAT assigned

external IP address with a rendezvous server. Other hosts that wish to initiate ICMP Query sessions to the registered hosts might do so using the public address registered with the Rendezvous server. For this reason, Basic NAT devices MUST support the traversal of hairpinned ICMP query sessions. This requirement is necessary for current applications to work correctly.

Packets belonging to any of the hairpinned sessions could in turn trigger ICMP error messages directed to the source of hairpinned IP packets. Such hairpinned ICMP error messages will traverse the NAT devices enroute. All NAT devices (i.e., Basic NAT as well as NAPT devices) MUST support the traversal of hairpinned ICMP error messages. Specifically, the NAT device must translate not only the embedded hairpinned packet, but also the outer IP header that is hairpinned. This requirement is necessary for current applications to work correctly.

A hairpinned ICMP error message is received from a node in private network. As such, the ICMP error processing requirement specified in Req-5 is applicable in its entirety in processing the ICMP error message. In addition, the NAT device MUST translate the destination IP address of the outer IP header to be same as the source IP address of the embedded IP packet after the translation.

REQ-7: NAT devices enforcing Basic NAT ([\[NAT-TRAD\]](#)) MUST support the traversal of hairpinned ICMP query sessions. All NAT devices (i.e., Basic NAT as well as NAPT devices) MUST support the traversal of hairpinned ICMP error messages.

a) When forwarding a hairpinned ICMP error message, the NAT device MUST translate the destination IP address of the outer IP header to be same as the source IP address of the embedded IP packet after the translation.

## [6.](#) Rejection of Outbound Flows Disallowed by NAT

A NAT device typically permits all outbound sessions. However, a NAT device may disallow some outbound sessions due to resource constraints or administration considerations. For example, a NAT device may not permit the first packet of a new outbound session,

if the NAT device is out of resources (out of addresses or TCP/UDP

ports or NAT Session resources) to set up a state for the session, or, the specific session is administratively restricted by the NAT device.

When the first packet of an outbound flow is prohibited by a NAT device due to resource constraints or administration considerations, the NAT device SHOULD send ICMP destination unreachable message. [Section 5.2.7.1 of \[RFC1812\]](#) recommends routers to use ICMP code 13 (Communication administratively prohibited) when they administratively filter packets. ICMP code 13 is a soft error and is on par with other soft error codes generated in response to transient events such as 'network unreachable' (ICMP type=3, code=0). A NAT device SHOULD use ICMP code 13 when generating an ICMP error message. This requirement is meant primarily for future use. Current applications do not require this for them to work correctly.

Some NAT designers opt to never reject an outbound flow. When a NAT runs short of resources, they prefer to steal a resource from an existing NAT Session rather than reject the outbound flow. Such a design choice may appear conformant to REQ-8 below. However, the design choice is in violation of the spirit of both REQ-8 and REQ-2. Such a design choice is strongly discouraged.

REQ-8: When a NAT device is unable to establish a NAT Session for a new transport-layer (TCP, UDP, ICMP, etc.) flow due to resource constraints or administrative restrictions, the NAT device SHOULD send an ICMP destination unreachable message, with a code of 13 (Communication administratively prohibited) to the sender, and drop the original packet.

## 7. Conformance to [RFC 1812](#)

A NAT device is inherently an intermediate router that forwards IP packets between private and external realms. As such, the NAT device MUST conform to all the requirements of a router, as specified in [\[RFC1812\]](#). [Section 5.2.7.1 of \[RFC1812\]](#) states that a router MUST also be able to generate ICMP Destination Unreachable messages and SHOULD choose a response code that most closely matches the reason the message is being generated.

Note, however, NAT devices also function as hosts on the Internet and are bound by the conformance requirements in [\[RFC1122\]](#). Protocol-specific BEHAVE documents ([\[BEH-UDP\]](#), [\[BEH-TCP\]](#)) identify instances where a NAT device should deviate from [RFC 1122](#). As such, the host behavior requirements of NAT devices specified in the

protocol-specific BEHAVE drafts take precedence over [RFC 1122](#).

The focus of this section is on conformance to router requirements. The following sub sections identify specific instances where a NAT device would be expected to conform to [RFC 1812](#).

### [7.1](#). IP packet fragmentation

Many networking applications (which include TCP as well as UDP based applications) depend on ICMP error messages from the network to perform end-to-end path MTU discovery [[PMTU](#)]. Once path MTU is discovered, an application that chooses to avoid fragmentation may do so by originating IP packets that fit within the Path MTU enroute and setting the DF (Don't Fragment) bit in the IP header, so the intermediate nodes enroute do not fragment the IP packets. The following sub-sections discuss the need for NAT devices to honor the DF bit in the IP header and be able to generate "Packet too big" ICMP error message when they cannot forward the IP packet without fragmentation. Also discussed is the need to seamlessly forward ICMP error messages generated by other intermediate devices.

#### [7.1.1](#). Generating "Packet too Big" ICMP error Message

When a router is unable to forward a datagram because it exceeds the MTU of the next-hop network and its Don't Fragment (DF) bit is set, the router is required by [[RFC1812](#)] to return an ICMP Destination Unreachable message to the source of the datagram, with the Code indicating "fragmentation needed and DF set". Further, [[PMTU](#)] states that the router MUST include the MTU of that next-hop network in the low-order 16 bits of the ICMP header field that is labeled "unused" in the ICMP specification[ICMP].

A NAT device MUST honor the DF bit in the IP header of the packets that transit the device. If the DF bit is set and the MTU on the forwarding interface of the NAT device is such that the IP datagram cannot be forwarded without fragmentation, the NAT device MUST issue a "packet too big" ICMP message (ICMP type 3, Code 4) with a suggested MTU back to the sender and drop the original IP packet. The sender will usually resend after taking the appropriate corrective action.

If the DF bit is not set and the MTU on the forwarding interface of the NAT device mandates fragmentation, the NAT device MUST fragment the packet and forward the fragments [[RFC1812](#)].

#### [7.1.2](#). Forwarding "Packet too big" ICMP Error Message

This is flip side of the argument for the above section. By virtue

of the address translation NAT performs, NAT may end up being the recipient of "Packet too big" message.

When NAT device is the recipient of "Packet too big" ICMP message from the network, the NAT device MUST forward the ICMP message back to the intended recipient, pursuant to the previously stated requirements REQ-3, REQ-4, REQ-5 and REQ-6.

### [7.2.](#) Generating "Time Exceeded" Error Message

[Section 5.2.7.3 of RFC 1812](#) says that a router MUST generate "Time Exceeded" ICMP error message when it discards a packet due to an expired TTL field. A router MAY have a per-interface option to disable origination of these messages on that interface, but that option MUST default to allowing the messages to be originated.

### [7.3.](#) [RFC 1812](#) Conformance Requirements summary

The requirements outlined in sections [7.1](#) and [7.2](#) are necessary for the current applications to work correctly. The following summarizes the requirements specified in sections [7.1](#) and [7.2](#),

REQ-9: A NAT device MUST conform to [RFC 1812](#) in IP packet handling. Below are specific instances where a NAT device MUST conform to [RFC 1812](#).

- a) A NAT MUST honor the DF bit in the IP header. If the DF bit is set on a transit IP packet and the NAT device cannot forward the packet without fragmentation, the NAT device MUST send a "Packet too big" ICMP message (ICMP type 3, Code 4) with a suggested MTU back to the sender and drop the original IP packet. If the DF-bit is clear and MTU mandates fragmentation, the NAT device MUST fragment the packet and forward the fragments.
- b) A NAT device MUST, by default, generate "Time Exceeded" ICMP error message when it discards a packet due to an expired TTL field, unless explicitly configured otherwise.

## [8.](#) Summary of Requirements

This section summarizes the requirements discussed in the preceding sections.

REQ-1: When local policy permits, a NAT device MUST permit ICMP queries and their associated responses, when the query is initiated from a private host.

a) NAT mapping of ICMP Query identifiers SHOULD be external host independent.

REQ-2: An ICMP Query session timer MUST NOT expire in less than 60 seconds.

a) It is RECOMMENDED that the ICMP Query session timer be made configurable.

REQ-3: When an ICMP error packet is received, the NAT device SHOULD do the following.

a) If the IP checksum of the embedded packet fails to validate, drop the error packet; and

b) If the embedded packet includes IP options, traverse past the IP options to locate the start of transport header for the embedded packet; and

c) If the ICMP Error payload contains ICMP extensions([\[ICMP-EXT\]](#)), exclude the optional zero-padding and the ICMP extensions when evaluating transport checksum for the embedded packet; and

d) If the embedded packet contains the entire transport segment, and the transport protocol of the embedded packet requires the recipient to validate the transport checksum, and the checksum fails to validate, drop the error packet.

REQ-4: If a NAT device receives an ICMP error packet from external realm, and the NAT does not have an active mapping for the embedded payload, the NAT SHOULD silently drop the ICMP error packet. If the NAT has active mapping for the embedded payload and local policy permits, then the NAT MUST do the following prior to forwarding the packet.

a) Revert the IP and transport headers of the embedded IP packet to their original form, using the matching mapping; and

b) Leave the ICMP error type and code unchanged; and

c) Modify the destination IP address of the outer IP header to be same as the source IP address of the embedded packet after translation.

REQ-5: If a NAT device receives an ICMP error packet from private realm, and the NAT does not have an active mapping for the embedded payload, the NAT SHOULD silently drop the ICMP error packet. If the NAT has active mapping for the embedded payload and local policy permits, then the NAT MUST do the following prior to forwarding the packet.

- a) Revert the IP and transport headers of the embedded IP packet to their original form, using the matching mapping; and
- b) Leave the ICMP error type and code unchanged; and
- c) If the NAT enforces Basic NAT function ([\[NAT-TRAD\]](#)), and the NAT has active mapping for the IP address that sent the ICMP error, translate the source IP address of the ICMP error packet with the public IP address in the mapping. In all other cases, translate the source IP address of the ICMP error packet with its own public IP address.

REQ-6: While processing an ICMP error packet pertaining to an ICMP Query or Query response message, a NAT device MUST NOT refresh or delete the NAT Session that pertains to the embedded payload within the ICMP error packet.

REQ-7: NAT devices enforcing Basic NAT ([\[NAT-TRAD\]](#)) MUST support the traversal of hairpinned ICMP query sessions. All NAT devices (i.e., Basic NAT as well as NAPT devices) MUST support the traversal of hairpinned ICMP error messages.

- a) When forwarding a hairpinned ICMP error message, the NAT device MUST translate the destination IP address of the outer IP header to be same as the source IP address of the embedded IP packet after the translation.

REQ-8: When a NAT device is unable to establish a NAT Session for a new transport-layer (TCP, UDP, ICMP, etc.) flow due to resource constraints or administrative restrictions, the NAT device SHOULD send an ICMP destination unreachable message, with a code of 13 (Communication administratively prohibited) to the sender, and drop the original packet.

REQ-9: A NAT device MUST conform to [RFC 1812](#) in IP packet handling. Below are specific instances where a NAT device MUST conform to [RFC 1812](#).

- a) A NAT MUST honor the DF bit in the IP header. If the DF bit is set on a transit IP packet and the NAT device cannot forward the packet without fragmentation, the NAT device MUST send a "Packet too big" ICMP message (ICMP type 3, Code 4) with a suggested MTU back to the sender and drop the original IP packet. If the DF-bit is clear and MTU mandates fragmentation, the NAT device MUST fragment the packet and forward the fragments.
- b) A NAT device MUST, by default, generate "Time Exceeded" ICMP error message when it discards a packet due to an expired TTL field, unless explicitly configured otherwise.

## 9. Security Considerations

This document does not introduce any new security concerns related to ICMP error message handling in the NAT devices. However, the document does propose counter measures to mitigate security concerns that already exist with ICMP error messages.

[ICMP-ATK] lists a number of ICMP attacks that can be directed against end host TCP stacks and suggests remedies to counter the attacks. [\[TCP-SOFT\]](#) describes improvements to the handling of ICMP error messages in many of the existing TCP/IP stacks, including

Linux. [Section 4](#) of this document describes a number of measures by which NAT devices should validate and update the embedded payload in ICMP error messages prior to forwarding. These measures ensure that NATs forward the ICMP error messages reliably, as stipulated in [\[ICMP-ATK\]](#).

For example, a rogue entity could bombard the NAT device with a large number of ICMP errors. If the NAT device did not validate the legitimacy of the ICMP error packets, the ICMP errors would be forwarded directly to the end nodes. End hosts not capable of defending themselves against such bogus ICMP error attacks could be adversely impacted by such attacks. Req-3 recommends validating embedded payload prior to forwarding. Checksum validation by itself does not protect end hosts from attacks. However, checksum validation mitigates end hosts from malformed ICMP error attacks. Req-4 and Req-5 further mandate that when a NAT device does not find a mapping selection for the embedded payload, the NAT should drop the ICMP error packets, without forwarding.

A rogue source could also try and send bogus ICMP error messages for the active NAT sessions, with an intent to destroy the sessions. Req-6 averts such an attack by ensuring that an ICMP error message does not effect the state of a session on the NAT device.

Req-8 recommends a NAT device sending ICMP error message when the NAT device is unable to create a NAT session due to lack of resources. Some administrators may choose not to have the NAT device send ICMP error message, as doing so could confirm to a malicious attacker that the attack has succeeded. For this reason, sending of the specific ICMP error message stated in REQ-8 should be left to the discretion of the NAT device administrator.

Unfortunately, ICMP messages are sometimes blocked at network boundaries due to local security policy. Thus, some of the requirements in this document allow local policy to override the recommendations of this document. Blocking such ICMP messages is known to break some protocol features (most notably Path MTU Discovery) and some applications (e.g., ping, traceroute), and such blocking is NOT RECOMMENDED.

## [10.](#) IANA Considerations

There are no IANA considerations.

## [11.](#) Acknowledgements

Srisuresh, et. al.

[Page 20]

The authors wish to thank Fernando Gont, Dan Wing, Carlos Pignataro, Philip Matthews and members of the BEHAVE work group for doing a thorough review of the document and providing valuable input and offering generous amount of their time in shaping the ICMP requirements. Their valuable feedback makes this document a better read. The authors highly appreciate that. Dan Wing and Fernando Gont have been a steady source of inspiration throughout, for each rev of the document. Fernando Gont spent many hours preparing slides and presenting the document in the IETF on behalf of the authors. For this, the authors are grateful to Fernando. The authors also thank and sincerely appreciate Dan Tappan and Carlos Pignataro, the

authors of the [[ICMP-EXT](#)] document for their feedback on operational experience with ICMP extensions across NAT devices.

## Normative References

- [BEH-UDP] F. Audet and C. Jennings, "NAT Behavioral Requirements for Unicast UDP", [RFC 4787](#), January 2007.
- [ICMP] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), September 1981.
- [ICMP-EXT] Bonica, R., Gan, D., Nikander, P., Tappan, D., and Pignataro, C., "Extended ICMP to Support Multi-part Messages", [RFC 4884](#), April 2007.
- [NAT-MIB] Rohit, R., Srisuresh, P., Raghunarayan, R., Pai, N., and Wang, C., "Definitions of Managed Objects for Network Address Translators (NAT)", [RFC 4008](#), March 2005.
- [RFC793] "Transmission Control Protocol DARPA Internet Program Protocol Specification", [RFC 793](#), September 1981.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", [RFC 1812](#), June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

## Informative References

- [BEH-APP] Ford, B., Srisuresh, P., and Kegel, D., "Application Design Guidelines for Traversal through Network Address Translators", [draft-ford-behave-app-05.txt](#) (Work In Progress), March 2007.

Srisuresh, et. al.

[Page 21]

---

Internet-Draft NAT Behavioral Requirements for ICMP

June 2007

- [BEH-TCP] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and Srisuresh, P., "NAT Behavioral Requirements for TCP", [draft-ietf-behave-tcp-07.txt](#) (Work In Progress), April 2007.

- [ICMP-ATK] Gont, F., "ICMP attacks against TCP", [draft-ietf-tcpm-icmp-attacks-02.txt](#) (Work In Progress), May 2007.
- [MS-TRCRT] Microsoft Support, "How to use the Tracert command-line utility to troubleshoot TCP/IP problems in Windows", <http://support.microsoft.com/kb/162326>, October, 2006.
- [NAT-TERM] Srisuresh, P. and Holdrege, M., "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [NAT-TRAD] Srisuresh, P., and Egevang, K., "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [PMTU] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#), November 1990.
- [RFC1122] Braden, R., "Requirements for Internet Hosts -- Communication Layers", [RFC 1122](#), October 1989.
- [RFC1256] Deering, S., "ICMP Router Discovery Messages", [RFC1256](#), September 1991.
- [RFC1470] Stine, R., "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices", [RFC 1470](#), June 1993.
- [RFC4065] Kempf, J., "Instructions for Seamoby and Experimental Mobility Protocol IANA Allocations", [RFC 4065](#), July 2005.
- [TCP-SOFT] Gont, F., "TCP's Reaction to Soft Errors", [draft-ietf-tcpm-tcp-soft-errors-06.txt](#) (Work In Progress), June 2007.
- [UNSAF] Daigle, L., and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", [RFC 3424](#), November 2002.

Author's Addresses:

Pyda Srisuresh  
Kazeon Systems, Inc.  
1161 San Antonio Rd.  
Mountain View, CA 94043  
U.S.A.  
Phone: (408)836-4773  
E-mail: srisuresh@yahoo.com

Bryan Ford  
Computer Science and Artificial Intelligence Laboratory  
Massachusetts Institute of Technology  
77 Massachusetts Ave.  
Cambridge, MA 02139  
U.S.A.  
Phone: (617) 253-5261  
E-mail: baford@mit.edu  
Web: <http://www.brynosaurus.com/>

Senthil Sivakumar  
Cisco Systems, Inc.  
7100-8 Kit Creek Road  
PO Box 14987  
Research Triangle Park, NC 27709-4987  
U.S.A.  
Phone: +1 919 392 5158  
Email: ssenthil@cisco.com

Saikat Guha  
Cornell University  
331 Upson Hall  
Ithaca, NY 14853  
U.S.A.  
Email: saikat@cs.cornell.edu

## Intellectual Property Statement

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](mailto:ietf-ipr@ietf.org).

#### 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.

#### Acknowledgment

Funding for the RFC Editor function is currently provided by the IETF Trust.

