

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 18, 2007

D. Wing  
J. Rosenberg  
Cisco Systems  
October 15, 2006

Controlling NAT Bindings using STUN  
draft-wing-behave-nat-control-stun-usage-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 April 18, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Simple Traversal Underneath NAT (STUN) is a mechanism for traversing NATs. STUN requests are transmitted through a NAT to external STUN servers. While this works very well, its two primary drawbacks are the inability to modify the properties of a NAT binding and the need to query a public STUN server for every NAT binding. These drawbacks require frequent messages which present a load on servers (like SIP servers and STUN servers) and are bad for low speed access networks, such as cellular. This document proposes that the STUN server be

Internet-Draft

NAT Control STUN Usage

October 2006

embedded in the NAT itself, and describes how these STUN servers can be readily discovered and utilized to reduce queries to public STUN servers and to reduce NAT keepalive traffic.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Overview of Operation . . . . .</a>	<a href="#">3</a>
<a href="#">2.1.</a>	<a href="#">Nested NATs . . . . .</a>	<a href="#">4</a>
<a href="#">2.2.</a>	<a href="#">Interacting with Legacy NATs . . . . .</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">NAT Control Usage . . . . .</a>	<a href="#">7</a>
<a href="#">3.1.</a>	<a href="#">Applicability . . . . .</a>	<a href="#">7</a>
<a href="#">3.2.</a>	<a href="#">Client Discovery of Server . . . . .</a>	<a href="#">7</a>
<a href="#">3.3.</a>	<a href="#">Server Determination of Usage . . . . .</a>	<a href="#">7</a>
<a href="#">3.4.</a>	<a href="#">New Requests or Indications . . . . .</a>	<a href="#">7</a>
<a href="#">3.5.</a>	<a href="#">New Attributes . . . . .</a>	<a href="#">7</a>
<a href="#">3.5.1.</a>	<a href="#">XOR-INTERNAL-ADDRESS . . . . .</a>	<a href="#">8</a>
<a href="#">3.5.2.</a>	<a href="#">REFRESH-INTERVAL . . . . .</a>	<a href="#">8</a>
<a href="#">3.6.</a>	<a href="#">Client Procedures . . . . .</a>	<a href="#">8</a>
<a href="#">3.7.</a>	<a href="#">Server Procedures . . . . .</a>	<a href="#">9</a>
<a href="#">4.</a>	<a href="#">Benefits . . . . .</a>	<a href="#">10</a>
<a href="#">4.1.</a>	<a href="#">Incremental Deployment . . . . .</a>	<a href="#">10</a>
<a href="#">4.2.</a>	<a href="#">Optimize SIP-Outbound . . . . .</a>	<a href="#">11</a>
<a href="#">4.3.</a>	<a href="#">Optimize ICE . . . . .</a>	<a href="#">11</a>
<a href="#">4.3.1.</a>	<a href="#">Candidate Gathering . . . . .</a>	<a href="#">11</a>
<a href="#">4.3.2.</a>	<a href="#">Keepalive . . . . .</a>	<a href="#">11</a>
<a href="#">4.3.3.</a>	<a href="#">Learning STUN Servers without Configuration . . . . .</a>	<a href="#">11</a>
<a href="#">5.</a>	<a href="#">Limitations . . . . .</a>	<a href="#">12</a>
<a href="#">5.1.</a>	<a href="#">Nested NATs . . . . .</a>	<a href="#">12</a>
<a href="#">5.2.</a>	<a href="#">Address Dependent Mapping . . . . .</a>	<a href="#">12</a>
<a href="#">5.3.</a>	<a href="#">Address Dependent Filtering . . . . .</a>	<a href="#">13</a>
<a href="#">6.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">13</a>
<a href="#">6.1.</a>	<a href="#">Authorization and Resource Exhaustion . . . . .</a>	<a href="#">14</a>
<a href="#">6.2.</a>	<a href="#">Comparison to Other NAT Control Techniques . . . . .</a>	<a href="#">14</a>
<a href="#">6.3.</a>	<a href="#">Rogue STUN Server . . . . .</a>	<a href="#">14</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">References . . . . .</a>	<a href="#">15</a>
<a href="#">8.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">15</a>
<a href="#">8.2.</a>	<a href="#">Informational References . . . . .</a>	<a href="#">15</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">16</a>
	<a href="#">Intellectual Property and Copyright Statements . . . . .</a>	<a href="#">17</a>

## 1. Introduction

Two common usages of STUN [[I-D.ietf-behave-rfc3489bis](#)] are Binding Discovery and NAT Keepalive. The Binding Discovery usage allows a STUN client to learn its public IP address (from the perspective of the STUN server it contacted) and the NAT keepalive usage allows a STUN client to keep an active NAT binding alive. Unlike some other techniques (UPnP [[UPnP](#)], MIDCOM [[RFC3303](#)], Bonjour [[Bonjour](#)]), STUN does not interact directly with the NAT. Because STUN doesn't interact directly with the NAT, STUN cannot request additional services from the NAT such as longer lifetimes (which would reduce keepalive messages).

This paper describes a mechanism for the STUN client to communicate directly with the NAT and request additional services. To achieve this optimization, the STUN client communicates to outer-most NAT's embedded STUN server. Thereafter, the STUN client need only ask that NAT's embedded STUN server for public IP addresses and UDP ports -- as it will return the same information as the public STUN server. Additionally, the STUN client can ask the NAT's embedded STUN server to extend the NAT binding for the flow, and the STUN client can learn the IP address of the next-outer-most NAT. By repeating this procedure with the next-outer-most NAT, all of the NATs along that path can have their bindings extended. By learning all of the STUN servers on the path between the public Internet and itself, an endpoint can optimize the path of peer-to-peer communications.

## 2. Overview of Operation

When a STUN client sends a STUN Request to a STUN server, it receives a STUN Response which indicates the IP address and UDP port seen by the STUN server. If the IP address and UDP port differs from the IP address and UDP port of the socket used to send the request, the STUN client knows there is at least one NAT between itself and the STUN server, and knows the 'public' IP address of that NAT. For example,

in the following diagram, the STUN client learns the public IP address of its NAT is 161.44.1.1:

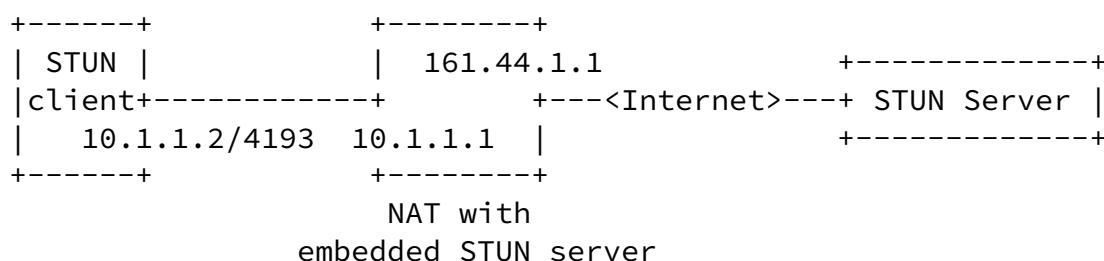


Figure 1: One NAT

After learning the public IP address of its outer-most NAT, the STUN client attempts to communicate with the STUN server embedded in that outer-most NAT. The STUN client does this by sending a STUN Binding Request to the public IP address of the NAT itself -- 161.44.1.1, in the figure above. If the NAT is running a STUN server and supports the usage described in this document, the NAT will return a STUN Binding Response message including the XOR-INTERNAL-ADDRESS attribute, which will indicate the IP address and UDP port seen on the \*internal\* side of the NAT. In the example above, the IP address and UDP port indicated in XOR-INTERNAL-ADDRESS are the same as that used by the STUN client (10.1.1.2/4193), which indicates there are no other NATs between the STUN client and that outer-most NAT.

The preceding technique allows the STUN client to communicate directly with all of the on-path NATs, even if there are multiple NATs in the path. This is described in section [Section 2.1](#).

The STUN client can request each NAT to increase the binding lifetime, as described in [Section 3.5](#). The STUN client receives

positive confirmation that the binding lifetime has been extended, allowing the STUN client to significantly reduce its NAT keepalive traffic. Additionally, as long as the NAT complies with [\[I-D.ietf-behave-nat-udp\]](#), the STUN client's keepalive traffic need only be sent to the outer-most NAT's IP address.

Additionally, the STUN client can request additional service from the NAT for that flow via the new attributes defined in .

## 2.1. Nested NATs

Nested NATs are controlled individually. The nested NATs are discovered, from outer-most NAT to the inner-most NAT, using the XOR-INTERNAL-ADDRESS attribute.

The following diagram shows how a STUN client iterates over the NATs to communicate with all of the NATs in the path. First, the STUN client would learn the outer-most NAT's IP address by performing the steps above. In the case below, however, the IP address and UDP port

indicated by the XOR-INTERNAL-ADDRESS will not be the STUN client's own IP address and UDP port -- rather, it's the IP address and UDP port on the *\*outer\** side of the NAT-B -- 10.1.1.2.

Because of this, the STUN client repeats the procedure and sends another STUN Binding Request to that newly-learned address (the *\*outer\** side of NAT-B). NAT-B will respond with a STUN Binding Response containing the XOR-INTERNAL-ADDRESS attribute, which will match the STUN client's IP address and UDP port. The STUN client knows there are no other NATs between itself and NAT-B, and finishes.

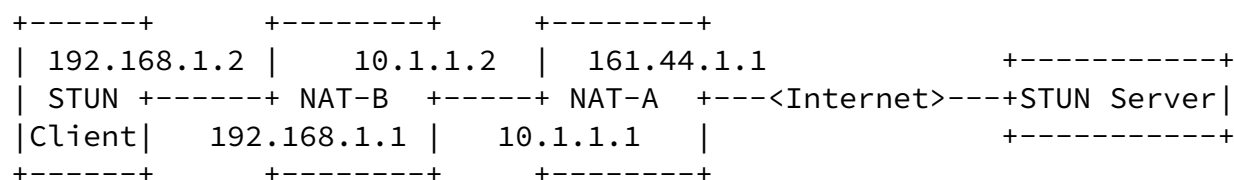
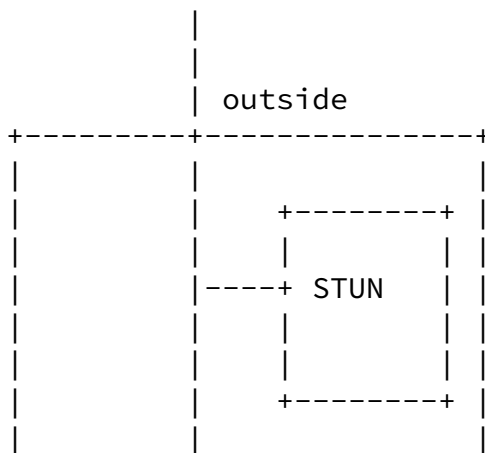


Figure 2: Two NATs

Internally, the NAT can be diagrammed to function like this, where the NAT operation occurs before the STUN server.



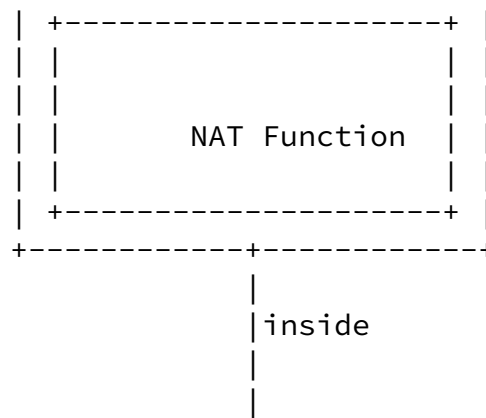


Figure 3: Block Diagram of Internal NAT Operation

## [2.2.](#) Interacting with Legacy NATs

There will be cases where the STUN client attempts to communicate with an on-path NAT which does not support the usage described in this document (that is, does not return the IP address of the internal NAT binding using the XOR-INTERNAL-ADDRESS attribute) or does not run a STUN server on its public interface. In both cases the optimizations described in this document won't be available to the STUN client and the STUN client will have to guess or assume the NAT binding timeout of the on-path NATs between itself and the first NAT that didn't run a STUN server.

Note: This is no worse than the condition today, and allows incremental upgrades of applications and NATs that implement the technique described in this document.

## [3.](#) NAT Control Usage

This section describes a new STUN usage, following the recommendation for defining a new usage in [[I-D.ietf-behave-rfc3489bis](#)].

### [3.1.](#) Applicability

This STUN usage MAY be used by a STUN client that discovers there is

a NAT between itself and its STUN server. Such discovery would most likely occur with a STUN Binding Request / Binding Response exchange to a STUN server on the Internet, and by noticing the IP address and UDP port indicated by the XOR-MAPPED-ADDRESS attribute of the STUN Binding Response differs from the local socket's IP address and UDP port. Such a difference indicates a NAT is present between the STUN client and its STUN server.

### [3.2.](#) Client Discovery of Server

As this usage involves communicating with on-path NATs directly, the client needs to find those NATs. The outer-most NAT is found by the normal XOR-MAPPED-ADDRESS attribute in a STUN Response. To iterate through the inner NATs, each NAT needs to support the usage described in this document, and the STUN client discovers each of those NATs by iterating through the XOR-INTERNAL-ADDRESS attribute returned by those NATs. This is described in diagrams and examples in [Section 2](#).

### [3.3.](#) Server Determination of Usage

If a STUN Binding Request is received from a NAT's private interface and sent to the IP address of its public interface, the STUN server can assume the NAT Control Usage.

### [3.4.](#) New Requests or Indications

This usage does not define any new message types.

### [3.5.](#) New Attributes

The following figure indicates which attributes are present in which messages for this usage. An M indicates that inclusion of the attribute in the message is mandatory, O means its optional, C means it's conditional based on some other aspect of the message, and - means that the attribute is not applicable to that message type.

Attribute	Error			
	Request	Response	Response	Indication
XOR-INTERNAL-ADDRESS	-	M	-	-

### 3.5.1. XOR-INTERNAL-ADDRESS

This attribute MUST be present in a Binding Response and may be used in other responses as well. This attribute is necessary to allow a STUN client to 'walk backwards' and communicate directly with all of the STUN-aware NATs along the path.

The format of the XOR-INTERNAL-ADDRESS attribute is:

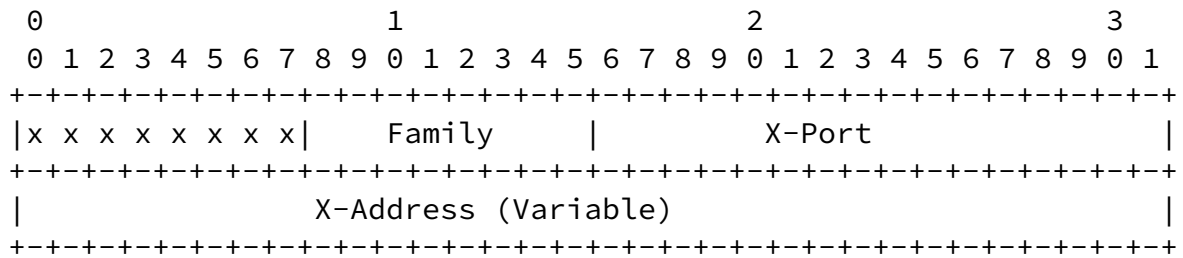


Figure 5: XOR-INTERNAL-ADDRESS Attribute

The meaning of Family, X-Port, and X-Address are exactly as in [I-D.ietf-behave-rfc3489bis]. The length of X-Address depends on the address family (IPv4 or IPv6).

### 3.5.2. REFRESH-INTERVAL

The REFRESH-INTERVAL attribute is defined in [I-D.ietf-behave-rfc3489bis] where it can only appear in a response. In the NAT Control usage defined in this document, the REFRESH-INTERVAL may also appear in a request.

In a Binding Request, the REFRESH-INTERVAL indicates the desired mapping timeout. In a Binding Response, the REFRESH-INTERVAL indicates the NAT's mapping timeout.

## 3.6. Client Procedures

The STUN client sends a STUN Binding Request to its STUN server and receives a STUN Binding Response, as normal. The STUN Binding Response contains the XOR-MAPPED-ADDRESS attribute which indicates the IP address and UDP port of the STUN Binding Request, as seen by the STUN server. If this IP address differs from the STUN client's IP address, the STUN client knows there is at least one NAT between itself and the STUN server, and it continues with the procedure; otherwise, it stops.

The STUN client now knows the public IP address of its outer-most NAT

-- it was returned in the XOR-MAPPED-ADDRESS attribute. The STUN client performs the Short-Term Password usage with the public IP address of its outer-most NAT, which is done over a TLS connection to the STUN TCP port (3478), following the procedures in the Short-Term Usage section of [[I-D.ietf-behave-rfc3489bis](#)]. As a result of the Short-Term Password usage, the STUN client now has a STUN USERNAME and PASSWORD which authenticate subsequent messages between the STUN client and server.

If subsequent messages from the STUN server fail authentication, the STUN client MUST re-obtain its IP address from a public STUN server, not from its outer-most NAT ([Section 6.3](#)).

To modify an existing NAT mapping's attributes, or to request a new NAT mapping, the STUN client can now send a STUN Binding Request to the IP address of address in its outer-most NAT's STUN UDP port (3478). The STUN server will respond with a STUN Binding Response containing an XOR-MAPPED-ADDRESS attribute (which points at the NAT's public IP address and port -- just as if the STUN Binding Request had been sent to a STUN server on the public Internet) and an XOR-INTERNAL-ADDRESS attribute (which points to the source IP address and UDP port the packet STUN Binding Request packet had prior to being NATted).

If the XOR-INTERNAL-ADDRESS attribute indicates an IP address and UDP port different from the STUN client's own IP address and port, the STUN client knows there is at least one NAT between itself and the STUN server it last contacted. If the STUN client wants to use multiple STUN servers, or wants to control the properties of the NAT bindings in each of those NATs, the STUN client can iteratively perform the Short-Term Password Usage followed by the Binding Discovery Usage with each NAT learned via the XOR-INTERNAL-ADDRESS attribute from the previous NAT.

In each case where the STUN client is sending STUN Binding Requests to the NATs, the STUN client can also include other STUN attributes to request certain properties for the flow. Requesting certain properties may require the STUN client to obtain short-term credentials. Defined in this document is a requested lifetime for the NAT binding in order to reduce keepalive traffic (REFRESH-INTERVAL).

### [3.7](#). Server Procedures

The server should listen for STUN Shared Secret Requests and STUN Binding Requests on the STUN UDP and TCP ports (UDP/3478, TCP/3478)

on its public IP address, from hosts connected to its private interface(s). The NAT SHOULD only respond to such messages which

arrive from its 'internal' interface. STUN Binding Requests sent to the NAT's public IP address which arrived from its public interface SHOULD be handled as if the NAT isn't listening on that port (e.g., return an ICMP error).

After receiving a STUN Shared Secret Request, the NAT follows the procedures described in the Short-Term Usage section of [\[I-D.ietf-behave-rfc3489bis\]](#). The embedded STUN server MUST store that username and password so long as any NAT bindings, created or adjusted with that same STUN username, have active mappings on the NAT.

After receiving a STUN Binding Request containing the REFRESH-INTERVAL attribute, the server SHOULD authenticate the request using the USERNAME attribute and the previously-shared STUN password (this is to defend against resource starvation attacks, see [Section 6.1](#)). If authenticated, the new binding's lifetime can be maximized against the NAT's configured sanity check and the lifetime indicated in the REFRESH-INTERVAL attribute of the STUN Binding Response.

In addition to its other attributes, the Binding Response always contains the XOR-MAPPED-ADDRESS and XOR-INTERNAL-ADDRESS attributes. The XOR-MAPPED-ADDRESS contains the public IP address and UDP port for this binding. The XOR-INTERNAL-ADDRESS contains the IP address and UDP port of the STUN Binding Request prior to the NAT translation. The XOR-INTERNAL-ADDRESS is used by the STUN client to walk backwards through nested NATs.

For example, looking at Figure 1, the XOR-INTERNAL-ADDRESS is the IP address and UDP port *\_prior to\_* the NAT operation. If there is only one NAT, as shown in Figure 1, XOR-INTERNAL-ADDRESS would contain the STUN client's own IP address and UDP port. If there are multiple NATs, XOR-INTERNAL-ADDRESS would indicate the IP address of the next NAT (that is, the next NAT closer to the STUN client). Iterating over this procedure allows the STUN client to find all of the NATs along the path.

#### [4.](#) Benefits

#### [4.1.](#) Incremental Deployment

NAT Control can be incrementally deployed. If the outer-most NAT does not support it, the STUN client behaves as normal. Where the outer-most STUN NAT does support it, the STUN client can gain some significant optimizations as described in the following sections.

Likewise, there is no change to applications if NATs are deployed

which support NAT Control.

#### [4.2.](#) Optimize SIP-Outbound

In sip-outbound [[I-D.ietf-sip-outbound](#)], the SIP proxy is also the STUN server. This document enables two optimizations of SIP-Outbound's keepalive mechanism:

1. STUN keepalive messages need only be sent to the outer-most NAT, rather than across the access link to the SIP proxy, which vastly reduces the traffic to the SIP proxy, and;
2. all of the on-path NATs can explicitly indicate their timeouts, reducing the frequency of keepalive messages.

#### [4.3.](#) Optimize ICE

The NAT Control usage provides several opportunities to optimize ICE [[I-D.ietf-mmusic-ice](#)].

##### [4.3.1.](#) Candidate Gathering

During its candidate gathering phase, an ICE endpoint normally contacts a STUN server on the Internet. If an ICE endpoint discovers that its outer-most NAT runs a STUN server, the ICE endpoint can use the outer-most NAT's STUN server rather than using the STUN server on the Internet. This saves access bandwidth and reduces the reliance on the STUN server on the Internet -- the STUN server on the Internet need only be contacted once.

##### [4.3.2.](#) Keepalive

ICE uses STUN as its primary media stream keepalive mechanism. This

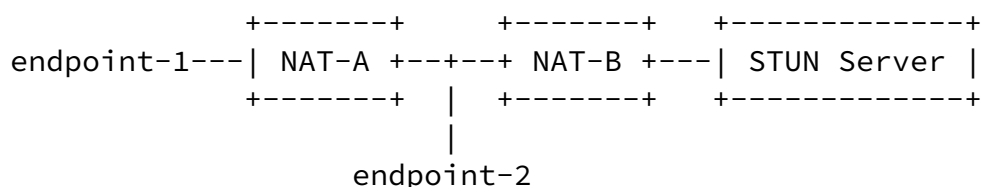
document enables two optimizations of ICE's keepalive techniques:

1. STUN keepalive messages need only be sent to the outer-most NAT, rather than across the access link to the remote peer, and;
2. all of the on-path NATs can explicitly indicate their timeouts, reducing the frequency of keepalive messages.

#### [4.3.3. Learning STUN Servers without Configuration](#)

ICE allows endpoints to have multiple STUN servers, but it is difficult to configure all of the STUN servers in the ICE endpoint -- it requires some awareness of network topology. By using the 'walk backward' technique described in this document, all the on-path NATs and their embedded STUN servers can be learned without additional configuration. By knowing the STUN servers at each address domain, ICE endpoints can optimize the network path between two peers.

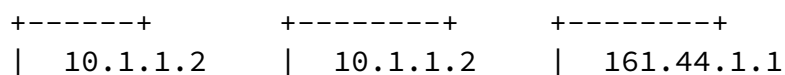
For example, if endpoint-1 is only configured with the IP address of the STUN server on the left, endpoint-1 can learn about NAT-B and NAT-A. Utilizing the STUN server in NAT-A, endpoint-1 and endpoint-2 can optimize their media path so they make the optimal path from endpoint-1 to NAT-A to endpoint-2:

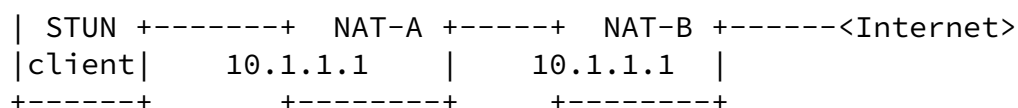


## [5. Limitations](#)

### [5.1. Nested NATs](#)

If nested NATs have overlapping IP address space, there will be undetected NATs on the path. Looking at the following figure, NAT-A and NAT-B are both using 10.1.1.x as their 'private' network. When this occurs, the STUN client will be unable to detect the presence of NAT-A if NAT-A assigns the same UDP port.

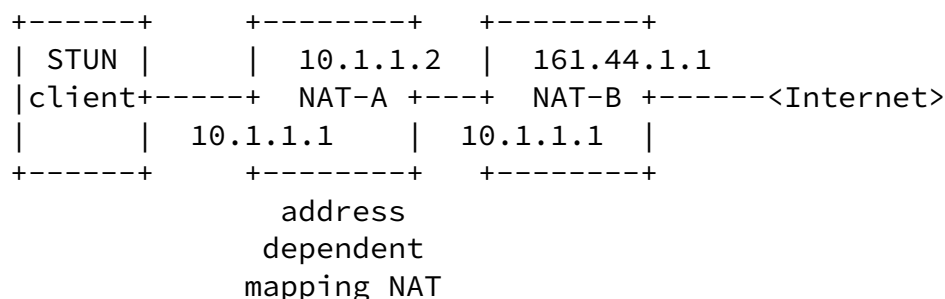




## 5.2. Address Dependent Mapping

In order to utilize the mechanisms described in this document, a STUN Request is sent from the same IP address and source port as the original STUN Binding Discovery message, but is sent to a different IP address (the IP address of an on-path NAT). If there is an on-path NAT, between the STUN client and the STUN server, with address dependent mapping behavior (as described in section 4.1 of [[I-D.ietf-behave-nat-udp](#)]), that NAT will prevent a STUN client from taking advantage of the technique described in this document.

An example of such a topology is shown in the following figure:



In this figure, NAT-A is a NAT that has address dependent mapping. Thus, when the STUN client sends a STUN Binding Request to 161.44.1.1 on UDP/3478, NAT-A will choose a new public UDP port for that communication. NAT-B will function normally, returning a different port in its XOR-MAPPED-ADDRESS, which indicates to the STUN client

that a symmetric NAT exists between the STUN client and the STUN server it just queried (NAT-B, in this example).

Open issue: We could resolve this problem by introducing a new STUN attribute which indicates the UDP port the STUN client wants to control. However, this changes the security properties of NAT Control, so this seems undesirable.

Open issue: When the STUN client detects this situation, should we recommend it abandon the NAT Control usage, and revert to operation as if it doesn't support the NAT Control usage?

### [5.3.](#) Address Dependent Filtering

If there is an NAT along the path that has address dependent filtering (as described in [section 5](#) of NAT-UDP), and the STUN client sends a STUN packet directly to any of the on-path NATs public addresses, the address-dependent filtering NAT will filter packets from the remote peer. Thus, after communicating with all of the on-path NATs the STUN client MUST send a UDP packet to the remote peer, if the remote peer is known.

Discussion: How many filter entries are in address dependent filtering NATs? If only one, this does become a real limitation if NATs are nested; if they're not nested, the outer-most NAT can avoid overwriting its own address in its address dependent filter.

## [6.](#) Security Considerations

This security considerations section will be expanded in a subsequent

version of this document. So far, the authors have identified the following considerations:

### [6.1.](#) Authorization and Resource Exhaustion

Only hosts that are 'inside' a NAT, which a NAT is already providing services for, can query or adjust the timeout of a NAT mapping.

A malicious STUN client could ask for absurdly long NAT bindings (days) for many UDP sessions, which would exhaust the resources in the NAT. To ensure the STUN client is not spoofing its IP address

when launching such an attack, the STUN server can challenge requests to extend the timeout by sending a NONCE to the STUN client. The STUN server can authorize an extension to the refresh timeout if a new request is sent with that same NONCE value.

Without considering this document and without considering STUN or other UNSAF NAT traversal techniques, a malicious TCP client can open many TCP connections, and keep them open, causing resource exhaustion in the NAT. A NAT which provide protection against such a TCP attack can provide a similar level of protection, via the NONCE challenge/response, as they can for TCP sessions.

## [6.2.](#) Comparison to Other NAT Control Techniques

Like UPnP, Bonjour, and host-initiated MIDCOM, the STUN usage described in this document alllows a host to learn its public IP address and UDP port mapping, and to request a specific lifetime for that mapping.

However, unlike those technologies, the NAT Control usage described in this dcoument only allows each UDP port on the host to create and adjust the mapping timeout of its own NAT mappings. Specifically, an application on a host can only adjust the duration of a NAT bindings for itself, and not for another application on that same host, and not for other hosts. This provides security advantages over other NAT control mechanisms where malicious software on a host can surreptitiously create NAT mappings to another application or to another host.

## [6.3.](#) Rogue STUN Server

As described in [Section 4](#), a STUN client can learn its outer-most NAT runs an embedded STUN server. Without the STUN client's knowledge, the outer-most NAT may acquire a new IP address (after moving to a new mobile network or DHCP lease expiration). When it acquires a new IP address, the STUN client will send a STUN Binding Request to the NAT's prior public IP address. If an attacker acquires that public

IP address and installs a rogue STUN server, the attacker has effectively acheived the attack "Compromise a Legitimate STUN Server" ([section 12.2.1 of \[RFC3489\]](#)). The attacker will send STUN Binding Responses indicating his IP address, which will be indistinguishable,

to the STUN client, from the behavior of the legitimate STUN server. To defend against this attack, the STUN client and STUN server need to obtain a short-term password as described in [\[I-D.ietf-behave-rfc3489bis\]](#).

## [7.](#) IANA Considerations

This document will add new IANA registrations for new STUN attributes.

[[This section will be completed in a later version of this document.]]

## [8.](#) References

### [8.1.](#) Normative References

- [I-D.ietf-behave-rfc3489bis]  
Rosenberg, J., "Simple Traversal Underneath Network Address Translators (NAT) (STUN)", [draft-ietf-behave-rfc3489bis-04](#) (work in progress), July 2006.
- [I-D.ietf-behave-nat-udp]  
Audet, F. and C. Jennings, "NAT Behavioral Requirements for Unicast UDP", [draft-ietf-behave-nat-udp-08](#) (work in progress), October 2006.
- [RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", [RFC 3489](#), March 2003.

### [8.2.](#) Informational References

- [UPnP] UPnP Forum, "Universal Plug and Play", 2000, <<http://www.upnp.org>>.
- [Bonjour] Apple Computer, "Bonjour", 2005, <<http://www.apple.com/macosx/features/bonjour/>>.
- [RFC3303] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A., and

A. Rayhan, "Middlebox communication architecture and framework", [RFC 3303](#), August 2002.

[I-D.ietf-mmusic-ice]

Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [draft-ietf-mmusic-ice-11](#) (work in progress), October 2006.

[I-D.ietf-sip-outbound]

Jennings, C. and R. Mahy, "Managing Client Initiated Connections in the Session Initiation Protocol (SIP)", [draft-ietf-sip-outbound-04](#) (work in progress), June 2006.

#### Authors' Addresses

Dan  
Cisco Systems  
170 West Tasman Drive  
San Jose, CA 95134  
USA

Email: [dwing@cisco.com](mailto:dwing@cisco.com)

Jonathan  
Cisco Systems  
600 Lanidex Plaza  
Parsippany, NJ 07054  
USA

Email: [jdrosen@cisco.com](mailto:jdrosen@cisco.com)

Internet-Draft

NAT Control STUN Usage

October 2006

## Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

## Acknowledgment

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