BEHAVE WG	J. Rosenberg
Internet-Draft	Cisco
Intended status: Standards Track	R. Mahy
Expires: January 14, 2009	Plantronics
	P. Matthews
	(Unaffiliated)
	July 13, 2008

TOC

Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)
draft-ietf-behave-turn-09

Status of this Memo

Abstract

If a host is located behind a NAT, then in certain situations it can be impossible for that host to communicate directly with other hosts (peers) located behind other NATs. In these situations, it is necessary for the host to use the services of an intermediate node that acts as a communication relay. This specification defines a protocol, called TURN (Traversal Using Relays around NAT), that allows the host to control the operation of the relay and to exchange packets with its peers using the relay.

The TURN protocol can be used in isolation, but is more properly used as part of the ICE (Interactive Connectivity Establishment) approach to NAT traversal.

Table of Contents

---

1. Introduction

TOC

Session Traversal Utilities for NAT (STUN) [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.) provides a suite of tools for facilitating the traversal of NAT. Specifically, it defines the Binding method, which is used by a client to determine its reflexive transport address towards the STUN server. The reflexive transport address can be used by the client for receiving packets from peers, but only when the client is behind "good" NATs. In particular, if a client is behind a NAT whose mapping behavior [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.) is address or address and port dependent (sometimes called "bad" NATs), the reflexive transport address will not be usable for communicating with a peer.
The only reliable way to obtain a UDP transport address that can be used for corresponding with a peer through such a NAT is to make use of a relay. The relay sits on the public side of the NAT, and allocates transport addresses to clients reaching it from behind the private side

of the NAT. These allocated transport addresses, called relayed transport address, are IP addresses and ports on the relay. When the relay receives a packet on one of these allocated addresses, the relay forwards it toward the client.
This specification defines an extension to STUN, called TURN, that allows a client to request a relayed transport address on a TURN server.
Though a relayed transport address is highly likely to work when corresponding with a peer, it comes at high cost to the provider of the relay service. As a consequence, relayed transport addresses should only be used as a last resort. Protocols using relayed transport addresses should make use of mechanisms to dynamically determine whether such an address is actually needed. One such mechanism, defined for multimedia session establishment protocols based on the offer/answer protocol in RFC 3264 (Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)," June 2002.) [RFC3264], is Interactive Connectivity Establishment (ICE) [I-D.ietf-mmusic-ice] (Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," October 2007.).
TURN was originally invented to support multimedia sessions signaled using SIP. Since SIP supports forking, TURN supports multiple peers per client; a feature not supported by other approaches (e.g., SOCKS [RFC1928] (Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5," March 1996.)). However, care has been taken in the later stages of its development to make sure that TURN is suitable for other types of applications.

---

2. Overview of Operation

TOC

This section gives an overview of the operation of TURN. It is non-normative.
In a typical configuration, a TURN client is connected to a private network (Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," February 1996.) [RFC1918] and through one or more NATs to the public Internet. On the public Internet is a TURN server. Elsewhere in the Internet are one or more peers that the TURN client wishes to communicate with. These peers may or may not be behind one or more NATs.

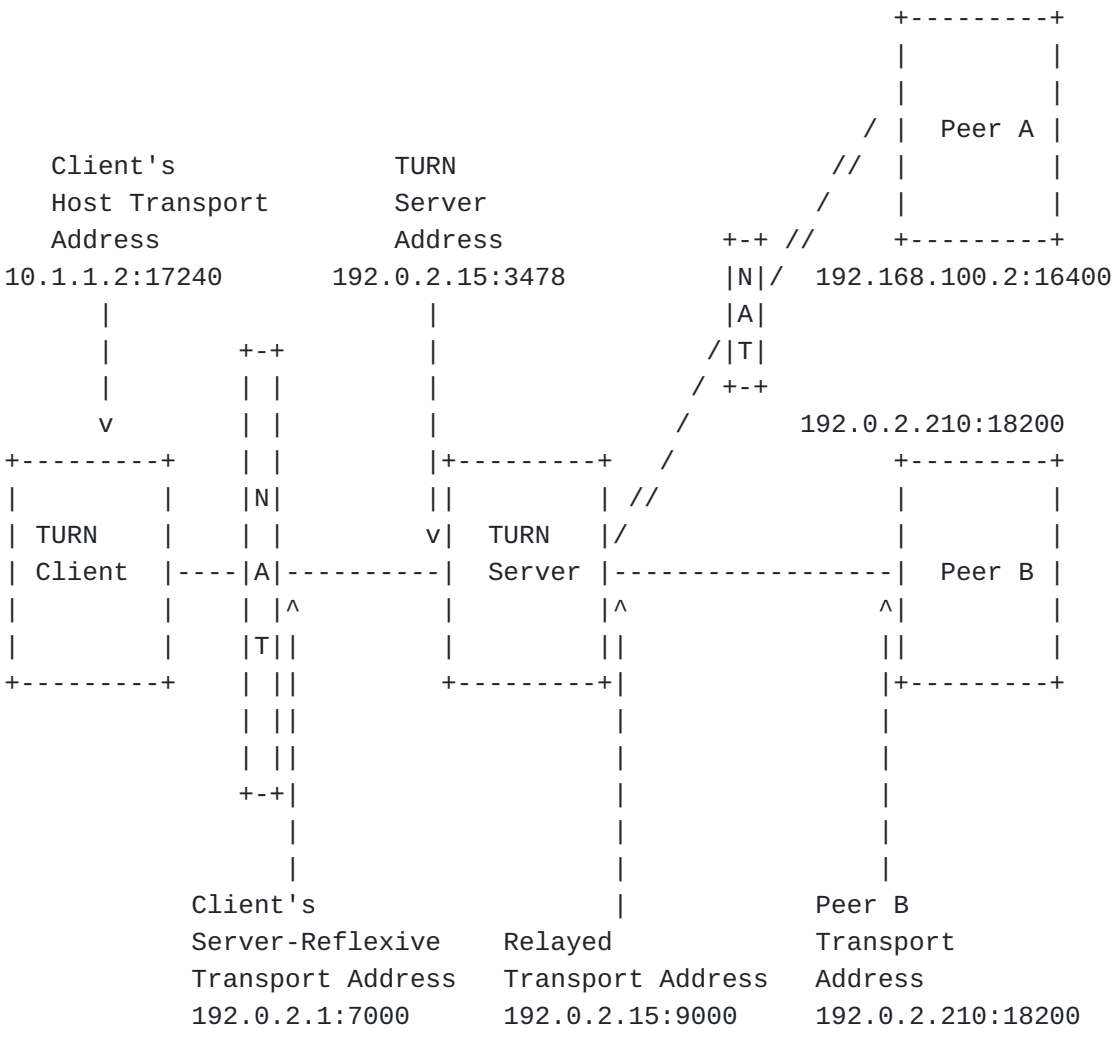

Figure 1

Figure 1 shows a typical deployment. In this figure, the TURN client and the TURN server are separated by a NAT, with the client on the private side and the server on the public side of the NAT. This NAT is assumed to be a “bad” NAT; for example, it might have a mapping property of address-and-port-dependent mapping (see [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.) for a description of what this means).

The client has allocated a local port on one of its addresses for use in communicating with the server. The combination of an IP address and a port is called a TRANSPORT ADDRESS and since this (IP address, port) combination is located on the client and not on the NAT, it is called the client’s HOST transport address.

The client sends TURN messages from its host transport address to a transport address on the TURN server which is known as the TURN SERVER

ADDRESS. The client learns the server's address through some unspecified means (e.g., configuration), and this address is typically used by many clients simultaneously. The TURN server address is used by the client to send both commands and data to the server; the commands are processed by the TURN server, while the data is relayed on to the peers.

Since the client is behind a NAT, the server sees these packets as coming from a transport address on the NAT itself. This address is known as the client's SERVER-REFLEXIVE transport address; packets sent by the server to the client's server-reflexive transport address will be forwarded by the NAT to the client's host transport address.

The client uses TURN commands to allocate a RELAYED TRANSPORT ADDRESS, which is an transport address located on the TURN server. The server ensures that there is a one-to-one relationship between the client's server-reflexive transport address and the relayed transport address; thus a packet received at the relayed transport address can be unambiguously relayed by the server to the client.

The client will typically communicate this relayed transport address to one or more peers through some mechanism not specified here (e.g., an ICE offer or answer [I-D.ietf-mmusic-ice] (Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," October 2007.)). Once this is done, the client can send data to the server to relay towards its peers. In the reverse direction, peers can send data to the relayed transport address of the client. The server will relay this data to the client as long as the client explicitly created a permission (see Section 2.5 (Permissions)) for the IP address of the peer.

2.1. Transports

TOC

TURN as defined in this specification only allows the use of UDP between the server and the peer. However, this specification allows the use of any one of UDP, TCP, or TLS over TCP to carry the TURN messages between the client and the server.

TURN client to TURN server	TURN server to peer
UDP	UDP
TCP	UDP
TLS over TCP	UDP

If TCP or TLS over TCP is used between the client and the server, then the server will convert between these transports and UDP transport when relaying data to/from the peer.

TURN supports TCP transport between the client and the server because some firewalls are configured to block UDP entirely. These firewalls block UDP but not TCP in part because TCP has properties that make the intention of the nodes being protected by the firewall more obvious to the firewall. For example, TCP has a three-way handshake that makes in clearer that the protected node really wishes to have that particular connection established, while for UDP the best the firewall can do is guess which flows are desired by using filtering rules. Also, TCP has explicit connection teardown, while for UDP the firewall has to use timers to guess when the flow is finished

TURN supports TLS over TCP transport between the client and the server because TLS provides additional security properties not provided by TURN's default digest authentication; properties which some clients may wish to take advantage of. In particular, TLS provides a way for the client to ascertain that it is talking to the server that it intended to, and also provides for confidentiality of TURN control messages. TURN does not require TLS because the overhead of using TLS is higher than that of digest authentication; for example, using TLS likely means that most application data will be doubly encrypted (once by TLS and once to ensure it is still encrypted in the UDP datagram).

There is a planned extension to TURN to add support for TCP between the server and the peers [I-D.ietf-behave-turn-tcp] (Perreault, S. and J. Rosenberg, “Traversal Using Relays around NAT (TURN) Extensions for TCP Allocations,” March 2010.). For this reason, allocations that use UDP between the server and the peers are known as UDP allocations, while allocations that use TCP between the server and the peers are known as TCP allocations. This specification describes only UDP allocations.

2.2. Allocations

TOC

To allocate a relayed transport address, the client uses an Allocate transaction. The client sends a Allocate request to the server, and the server replies with an Allocate response containing the allocated relayed transport address. The client can include attributes in the Allocate request that describe the type of allocation it desires (e.g., the lifetime of the allocation). And since relaying data may require lots of bandwidth, the server typically requires that the client authenticate itself using STUN’s long-term credential mechanism, to show that it is authorized to use the server.

Once a relayed transport address is allocated, a client must keep the allocation alive. To do this, the client periodically sends a Refresh request to the server with the allocated related transport address. TURN deliberately uses a different method (Refresh rather than Allocate) for refreshes to ensure that the client is informed if the allocation vanishes for some reason.

The frequency of the Refresh transaction is determined by the lifetime of the allocation. The client can request a lifetime in the Allocate request and may modify its request in a Refresh request, and the server always indicates the actual lifetime in the response. The client must issue a new Refresh transaction within 'lifetime' seconds of the previous Allocate or Refresh transaction. If a client no longer wishes to use an Allocation, it should do a Refresh transaction with a requested lifetime of 0.
Note that sending or receiving data from a peer DOES NOT refresh the allocation.
Both the server and the client keeps track of the client transport address and port, the server transport address and port, and the protocol used by the client to communicate with the server. These 5 values are collectively referred to as the 5-TUPLE. The server remembers the 5-tuple used in the Allocate request. Subsequent transactions between the client and the server use this same 5-tuple. In this way, the server knows which client owns the allocated relayed transport address. If the client wishes to allocate a second relayed transport address, it must use a different 5-tuple for this allocation (e.g., by using a different client host address or port).,

> NOTE: While the terminology used in this document refers to 5-tuples, the TURN server can store whatever identifier it likes that yields identical results. Specifically, many implementations use a file-descriptor in place of a 5-tuple to represent a TCP connection.

> NOTE: In some applications of TURN, a client may send and receive packets other than TURN packets on the address and port it is using to communicate with the TURN server. This can happen, for example, when using TURN with ICE [I-D.ietf-mmusic-ice] (Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," October 2007.). In these cases, the client can examine the 5-tuple for an arriving packet and use the 5-tuple to distinguish packets received from the TURN server from packets received from other nodes.

2.3. Exchanging Data with Peers

TOC

There are two ways for the client and peers to exchange data using the TURN server. The first way uses Send and Data indications, the second way uses channels. Common to both ways is the ability of the client to communicate with multiple peers using a single allocated relayed transport address; thus both ways include a means for the client to indicate to the server which peer to forward the data to, and for the server to indicate which peer sent the data.

When using the first way, the client sends a Send indication to the TURN server containing, in attributes inside the indication, the transport address of the peer and the data to be sent to that peer. When the TURN server receives the Send indication, it extracts the data from the Send indication and sends it in a UDP datagram to the peer, using the allocated relay address as the source address. In the reverse direction, UDP datagrams arriving at the relay address on the TURN server are converted into Data indications and sent to the client, with the transport address of the peer included in an attribute in the Data indication.

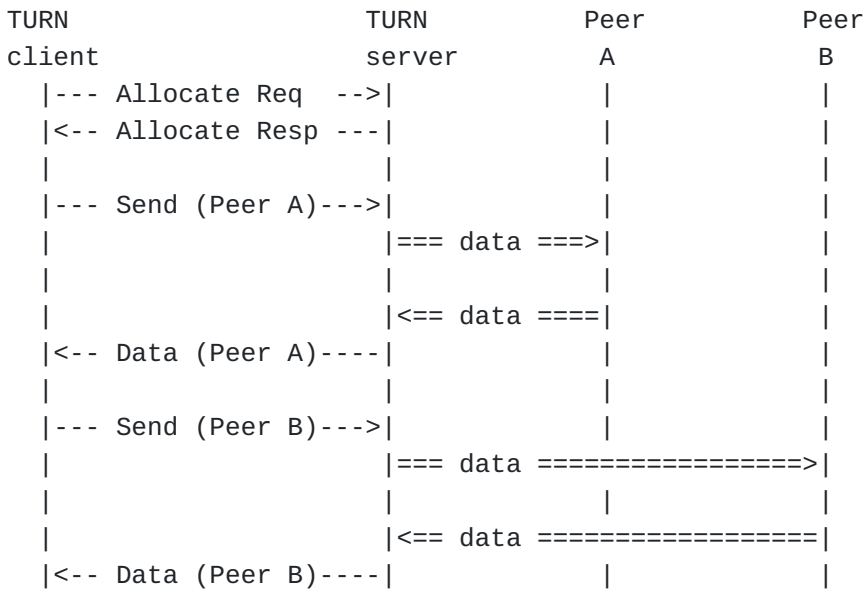

Figure 2

In the figure above, the client first allocates a relayed transport address. It then sends data to Peer A using a Send indication; at the server, the data is extracted and forwarded in a UDP datagram to Peer A, using the relayed transport address as the source transport address. When a UDP datagram from Peer A is received at the relayed transport address, the contents are placed into a Data indication and forwarded to the client. A similar exchange happens with Peer B.

2.4. Channels

TOC

For some applications (e.g. Voice over IP), the 36 bytes of overhead that a Send or Data indication adds to the application data can

substantially increase the bandwidth required between the client and the server. To remedy this, TURN offers a second way for the client and server to associate data with a specific peer.

This second way uses an alternate packet format known as the ChannelData message. The ChannelData message does not use the STUN header used by other TURN messages, but instead has a 4-byte header that includes a number known as a channel number. Each channel number in use is bound to a specific peer and thus serves as a shorthand for the peer's address.

To bind a channel to a peer, the client sends a ChannelBind request to the server, and includes an unbound channel number and the transport address of the peer. Once the channel is bound, the client can use a ChannelData message to send the server data destined for the peer. Similarly, the server can relay data from that peer towards the client using a ChannelData message.

Channel bindings last for 10 minutes unless refreshed. Channel bindings are refreshed by sending ChannelData messages from the client to the server, or by rebinding the channel to the peer.

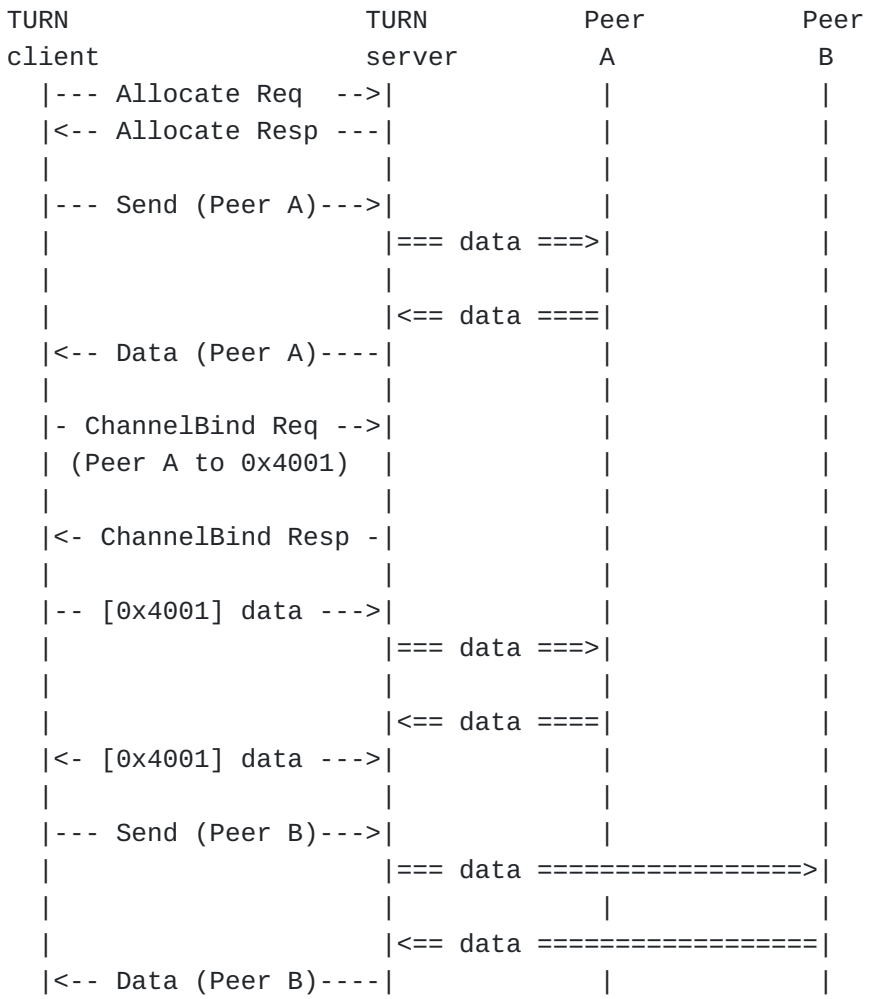

Figure 3

The figure above shows the channel mechanism in use. The client begins by allocating a relayed transport address, and then uses that address to exchange data with Peer A. After a bit, the client decides to bind a channel to Peer A. To do this, it sends a ChannelBind request to the server, specifying the transport address of Peer A and a channel number (0x4001). After that, the client can send application data encapsulated inside ChannelData messages to Peer A: this is shown as "[0x4001] data" where 0x4001 is the channel number.

Note that ChannelData messages can only be used for peers to which the client has bound a channel. In the example above, Peer A has been bound to a channel, but Peer B has not, so application data to and from Peer B uses Send and Data indications.

Channel bindings are always initiated by the client.

2.5. Permissions

TOC

To ease concerns amongst enterprise IT administrators that TURN could be used to bypass corporate firewall security, TURN includes the notion of permissions. TURN permissions mimic the address-restricted filtering mechanism of NATs that comply with [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.).
The client can install a permission by sending data to a peer (or by doing certain other things). Once a permission is installed, any peer with the same IP address (the ports numbers can differ) is permitted to send data to the client. After 5 minutes, the permission times out and the server drops any UDP datagrams arriving at the relayed transport from that IP address. Note that permissions are within the context of an allocation, so adding or expiring a permission in one allocation does not affect other allocations.
Data received from the peer DOES NOT refresh the permission.

2.6. Preserving vs. Non-Preserving Allocations

TOC

Some applications that use TURN are quite tolerant of the different possible ways a TURN server could set the Diff-Serv, ECN, TTL / Hop Limit, and Flow Label fields in the IP header of the outgoing packet. Other applications require that the TURN server set these fields in a specific way, and also require that the TURN server relay ICMP error packets. Applications in the second class typically wish to do Path MTU Discovery or end-to-end QOS.
Unfortunately, reading and manipulating fields in the IP header and relaying ICMP messages usually requires the server to have special permissions (e.g., access to RAW sockets or be loaded into the kernel), something that the person setting up the server may be unwilling or unable to grant. This is especially true when the server is part of a larger application, for example a peer-to-peer application. It is also significantly more difficult to implement this type of server than just relaying at the UDP layer.
To allow TURN to cater to both usage scenarios, TURN defines the concept of Preserving vs. Non-Preserving allocations. A Preserving allocation sets the fields in the outgoing IP header correctly, and also relays ICMP messages, while a Non-Preserving allocation may not relay correctly in every case. The relaying rules for a Preserving are designed to guarantee the following:

*Path MTU Discovery works end-to-end (i.e. client-to-peer), using either the old algorithm ([RFC1191] (Mogul, J. and S. Deering, “Path MTU discovery,” November 1990.) and [RFC1981] (McCann, J., Deering, S., and J. Mogul, “Path MTU Discovery for IP version 6,”

August 1996.)) or the new one ([RFC4821] (Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” March 2007.));

*ECN and Diff-Serv works end-to-end;

*Loops are prevented by copying and decrementing the TTL/Hop Count field.

If the client knows its application or usage scenario requires a Preserving allocation, then it can request one in its Allocate request. If the server is unable to grant this request, then it rejects the Allocate request.
Note that a Preserving allocation only makes sense when the transport protocol to the client is UDP; when the transport is TCP or TLS, the allocation is always Non-Preserving.

---

3. Terminology

TOC

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 (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) [RFC2119].
Readers are expected to be familar with [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.) and the terms defined there.
The following terms are used in this document:

TURN: A protocol spoken between a TURN client and a TURN server. It is an extension to the STUN protocol [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.). The protocol allows a client to allocate and use a relayed transport address.

TURN client: A STUN client that implements this specification.

TURN server: A STUN server that implements this specification. It relays data between a TURN client and its peer(s).

Peer: A host with which the TURN client wishes to communicate. The TURN server relays traffic between the TURN client and its peer(s). The peer does not interact with the TURN server using the protocol defined in this document; rather, the peer receives data sent by the TURN server and the peer sends data towards the TURN server.

Host Transport Address: A transport address allocated on a host.

Server-Reflexive Transport Address: A transport address on the "public side" of a NAT. This address is allocated by the NAT to correspond to a specific host transport address.

Relayed Transport Address: A transport address that exists on a TURN server. If a permission exists, packets that arrive at this address are relayed towards the TURN client.

Allocation: The relayed transport address granted to a client through an Allocate request, along with related state, such as permissions and expiration timers.

5-tuple: The combination (client IP address and port, server IP address and port, and transport protocol (UDP or TCP)) used to communicate between the client and the server . The 5-tuple uniquely identifies this communication stream. The 5-tuple also uniquely identifies the Allocation on the server.

Permission: The IP address and transport protocol (but not the port) of a peer that is permitted to send traffic to the TURN server and have that traffic relayed to the TURN client. The TURN server will only forward traffic to its client from peers that match an existing permission.

Preserving Allocation An allocation that sets the the fields in the IP header in a specific manner when relaying application data, and which also relays ICMP messages. An allocation that may not do this in some cases is called a Non-Preserving allocation.

---

4. General Behavior

TOC

This section contains general TURN processing rules that apply to all TURN messages.
TURN is an extension to STUN. All TURN messages, with the exception of the ChannelData message, are STUN-formatted messages. All the base processing rules described in [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.) apply to STUN-formatted messages. This means that all the message-forming and -processing descriptions in this document are implicitly prefixed with the rules of [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.).
In addition, the server SHOULD require that all TURN requests use the Long-Term Credential mechanism described in

[I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.), and the client MUST be prepared to authenticate requests if required. The server's administrator MUST choose a realm value that will uniquely identify the username and password combination that the client must use, even if the client uses multiple servers under different administrations. The server's administrator MAY choose to allocate a unique username to each client, or MAY choose to allocate the same username to more than one client (for example, to all clients from the same department or company).
The client and/or the server MAY include the FINGERPRINT attribute in any of the methods defined in this document. The client and server SHOULD include the SOFTWARE-TYPE attribute in all requests and responses, but SHOULD NOT include it in Send and Data indications. TURN does not use the backwards-compatibility mechanism described in [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.).
By default, TURN runs on the same port as STUN. However, either the SRV procedures or the ALTERNATE-SERVER procedures described in Section 6 (Creating an Allocation) may be used to run TURN on a different port.

---

5. Allocations

TOC

All TURN operations revolve around allocations, and all TURN messages are associated with an allocation. An allocation conceptually consists of the following state data:

- the relayed transport address
- The 5-tuple: client IP address, client port, server IP address, server port, transport protocol
- the username
- the transaction ID of the Allocate request
- the time-to-expiry
- A list of permissions
- A list of channel to peer bindings
- A flag indicating whether or not the allocation is Preserving

The relayed transport address is the transport address allocated by the server for communicating with peers, while the 5-tuple describes the

communication path between the client and the server. Both of these MUST be unique across all allocations, so either one can be used to uniquely identify the allocation.

When a TURN message arrives at the server from the client, the server uses the 5-tuple in the message to identify the associated allocation. For all TURN messages (including ChannelData) EXCEPT an Allocate request, if the 5-tuple does not identify an existing allocation, then the message MUST either be rejected with a 437 Allocation Mismatch error (if it is a request), or silently ignored (if it is an indication or a ChannelData message). A client receiving a 437 error response to a request other than Allocate MUST assume the allocation no longer exists.

The username and password of the allocation is the username and password of the authenticated Allocate request that creates the allocation. Subsequent requests on an allocation use the same username as that used to create the allocation, to prevent attackers from hijacking the client's allocation. Specifically, if the server requires the use of the Long-Term Credential mechanism, and if a non-Allocate request passes authentication under this mechanism, and if the 5-tuple identifies an existing allocation, but the request does not use the same username as used to create the allocation, then the request MUST be rejected with a 441 (Wrong Credentials) error.

The transaction ID of the allocation is the transaction ID used in the Allocate request. This is used to detect retransmissions of the Allocate request over UDP (see Section 6.2 (Receiving an Allocate Request) for details).

The time-to-expiry is the time in seconds left until the allocation expires. Each Allocate or Refresh transaction sets this timer, which then ticks down towards 0. By default, each Allocate or Refresh transaction resets this timer to 600 seconds (10 minutes), but the client can request a different value in the Allocate and Refresh request. Allocations can only be refreshed using the Refresh request; sending data to a peer does not refresh an allocation. When an allocation expires, the state data associated with the allocation can be freed. However the server MUST ensure that neither the relayed transport address nor the client reflexive transport address from the 5-tuple are re-used in other allocations until 2 minutes after the allocation expires; this ensures that any messages that are in transit when the allocation expires are gone before either of these transport addresses are re-used.

The list of permissions is described in Section 8 (Permissions) and the list of channels is described in Section 10 (Channels).

The differences between a Preserving and a Non-Preserving allocation are described in Section 11 (IP and ICMP).

TOC

6. Creating an Allocation

An allocation on the server is created using an Allocate transaction.

6.1. Sending an Allocate Request

TOC

The client forms an Allocate request as follows.
The client first needs to pick a host transport address that the server does not think is currently in use, or was recently in use. The client SHOULD pick a currently-unused transport address on the client's host (typically by allowing its OS to pick a currently-unused port for a new socket).
The client needs to pick a transport protocol to use between the client and the server. The transport protocol MUST be one of UDP, TCP, or TLS over TCP. Since this specification only allows UDP between the server and the peers, it is RECOMMENDED that the client pick UDP unless it has a reason to use a different transport. One reason to pick a different transport would be that the client believes, either through configuration or by experiment, that it is unable to contact any TURN server using UDP. See Section 2.1 (Transports) for more discussion.
The client must also pick a server transport address. Typically, this is done by the client learning (perhaps through configuration) one or more domain names for TURN servers. In this case, the client uses the DNS procedures described in [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.), but using an SRV service name of "turn" (or "turns" for TURN over TLS) instead of "stun" (or "stuns"). For example, to find servers in the example.com domain, the client performs a lookup for '_turn._udp.example.com', '_turn._tcp.example.com', and '_turns._tcp.example.com' if the client wants to communicate with the server using UDP, TCP, or TLS over TCP, respectively.
The client MUST include a REQUESTED-TRANSPORT attribute in the request. This attribute specifies the transport protocol between the server and the peers (note that this is NOT the transport protocol that appears in the 5-tuple). In this specification, the REQUESTED-TRANSPORT type is always UDP. This attribute is included to allow future extensions specify other protocols (e.g., [I-D.ietf-behave-turn-tcp] (Perreault, S. and J. Rosenberg, “Traversal Using Relays around NAT (TURN) Extensions for TCP Allocations,” March 2010.)).
If the client wishes the server to initialize the time-to-expire field of the allocation to some value other the default lifetime, then it MAY include a LIFETIME attribute specifying its desired value. This is just a request, and the server may elect to use a different value. Note that the server will ignore requests to initialize the field to less than the default value.

If the client required the allocation to satisfy certain properties, then the client includes the REQUESTED-PROPS attribute. This attribute is optional, and can be omitted if no special properties are required. Using the E and R bits in the REQUESTED-PROPS attribute, the client can request:

*(E=1, R=0) That the server allocate a relayed transport address with an even port number; OR

*(E=1, R=1) That the server reserve a pair of relayed transport addresses with adjacent port numbers N and N+1, where N is even and N+1 is odd, and then use port N for the current allocation. In this case, the server returns a RESERVATION-TOKEN attribute in the response which the client can then include in a subsequent Allocate request to create an allocation with port number N+1.

Note that the client cannot request a pair of adjacent ports unless it also requests that the lower numbered port be even. Thus the combination (E=0, R=1) is not allowed.
Similarly, by setting the P bit to 1 in the REQUESTED-PROPS attribute, the client can request that the server allocate a Preserving allocation.
For all the various REQUESTED-PROPS flags, if the server cannot satisfy the request, the Allocate request is rejected.
The client MAY also include a RESERVATION-TOKEN attribute in the request to ask the server to use a previously reserved port for the allocation. If the RESERVATION-TOKEN attribute is included, then the client MUST either omit the REQUESTED-PROPS attribute or set E=0 and R=0, since doing otherwise would make no sense.
Once constructed, the client sends the Allocate request on the 5-tuple.

6.2. Receiving an Allocate Request

TOC

When the server receives an Allocate request, it performs the following checks:

1. The server checks the credentials of the request, as per the Long-Term Credential mechanism of [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.).

2. The server checks if the 5-tuple is currently in use by an existing allocation, or was it in use by another allocation

within the last 2 minutes. If yes, then there are two sub-cases:

- If the transport protocol in the 5-tuple is UDP, and if the 5-tuple is currently in use by an existing allocation, and if the transaction id of the request matches the transaction id stored with the allocation, then the request is a retransmission of the original request. The server replies either with a stored copy of the original response, or with a response rebuilt from the stored state data. If the server chooses to rebuild the response, then (a) it need not parse the request further, but can immediately start building a success response, (b) the value of the LIFETIME attribute can be set to the current value of the time-to-expire timer, and (c) the server may need to include an extra field in the allocation to store the token returned in a RESERVATION-TOKEN attribute.

- Otherwise, the server rejects the request with a 437 (Allocation Mismatch) error.

NOTE: If the request includes credentials that are acceptable to server, but the 5-tuple is already in use, then it is important that the server reject the request with a 437 (Allocation Mismatch) error rather than a 401 (Unauthorized) error. This ensures that the client knows that the problem is with the 5-tuple, rather than (wrongly) believing that the problem lies with its credentials.

3. The server checks if the request contain a REQUESTED-TRANPORT attribute. If the REQUESTED-TRANSPORT attribute is not included or is malformed, the server rejects the request with a 400 (Bad Request) error. Otherwise, if the attribute is included but specifies a protocol other that UDP, the server rejects the request with a 422 (Unsupported Transport Protocol) error.

4. The server checks if the request contains a REQUESTED-PROPS attribute. If yes, then the server checks that it understands and can satisfy all the flags that are set to 1. If a flag is not understood, or if the server cannot satisfy the request, then the server rejects the request with a 508 (Insufficient Port Capacity) error. The server includes in its error response a REQUESTED-PROPS attribute with all the flags the server understands set to 1 and all others set to 0. Note that the combination (E=0, R=1) MUST be treated as unsupported.

5. The server checks if the request contains a RESERVATION-TOKEN attribute. If yes, and the request also contains a REQUESTED-PROPS attribute with the E and R flags set to any combination

other than E=0 and R=0, then the server rejectes the request with a 400 (Bad Request) error. Otherwise it checks to see if the token is valid (i.e., the token is in range and has not expired, and the corresponding relayed transport address is still available). If the token is not valid for some reason, the server rejects the request with a 508 (Insufficient Port Capacity) error.

6. At any point, the server MAY also choose to reject the request with a 486 (Allocation Quota Reached) error if it feels the client is trying to exceed some locally-defined allocation quota. The server is free to define this allocation quota any way it wishes, but SHOULD define it based on the username used to authenticate the request, and not on the client's transport address.

If the server rejects the request with one of the error codes 422 (Unsupported Transport Protocol), 486 (Allocation Quota Reached) or 508 (Insufficient Port Capacity), it MAY include an ALTERNATE-SERVER attribute in the error response redirecting the client to another server that it believes will accept the request. If the attribute is included, the address MUST be from the same address family as the server's transport address. Note that, if the attribute is included, the client will try this alternate server before trying the other servers given by the SRV procedures.

> NOTE: When UDP transport is used between the client and the server, the client will restransmit an Allocate request if it does not receive a response within a certain timeout period [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.). Because of this, the server may receive two (or more) Allocate requests with the same 5-tuple and same transaction id. Check #2 (above) handles the case where the first Allocate request is accepted and generates a success response, but it does not handle the case where the first request is rejected but the second request is accepted (because conditions on the server have changed in the brief intervening time period). If the client receives the first (failure) response, it will ignore the second (success) response and believe that an allocation was not created. An allocation created in this matter will eventually timeout, since the client will not refresh it. Furthermore, if the client later retries with the same 5-tuple but different transaction id, it will receive a 437 (Allocation Mismatch), which should cause it to retry with a different 5-tuple.
>
> Server implementors MAY elect to prevent this second case by remembering recent failure responses and returning the saved failure response when receiving a retransmitted Allocate request. This

> optional behavior may be appropriate when the server implements some sort of charging mechanism or a per-user quota. Alternatively, servers may use a smaller maximum lifetime value to mimize the lifetime of this "orphaned" allocation (see below).
>
> Server implementors debating whether to implement this optional feature should be aware that there are other scenarios in TURN that lead to such "orphaned" allocations.

If all the checks pass, the server creates the allocation. The 5-tuple is set to the 5-tuple from the Allocate request, while the list of permissions and the list of channels are initially empty.
When allocating a relayed transport address for the allocation, the server MUST allocate an IP address from the same family (e.g, IPv4 vs. IPv6) as that on which the request was received (i.e., the server's IP address in the 5-tuple for the allocation).

> NOTE: An extension to TURN to allow an address from a different address family is currently in progress [I-D.ietf-behave-turn-ipv6] (Camarillo, G., Novo, O., and S. Perreault, “Traversal Using Relays around NAT (TURN) Extension for IPv6,” March 2010.).

In addition, the server SHOULD only allocate ports from the range 49152 – 65535 (the Dynamic and/or Private Port range [Port-Numbers] (, “IANA Port Numbers Registry,” .)), unless the TURN server application knows, through some means not specified here, that other applications running on the same host as the TURN server application will not be impacted by allocating ports outside this range. This condition can often be satisfied by running the TURN server application on a dedicated machine and/or by arranging that any other applications on the machine allocate ports before the TURN server application starts. In any case, the TURN server SHOULD NOT allocate ports in the range 0 - 1023 (the Well-Known Port range) to discourage clients from using TURN to run standard services.

> NOTE: The IETF is currently investigating the topic of randomized port assignments to avoid certain types of attacks (see [I-D.ietf-tsvwg-port-randomization] (Larsen, M. and F. Gont, “Transport Protocol Port Randomization Recommendations,” April 2010.)). It is recommended that a TURN implementor keep abreast of this topic and, if appropriate, implement a randomized port assignment algorithm. This is especially applicable to servers that choose to pre-allocate a number of ports from the underlying OS and then later assign them to allocations; for example, a server may choose this technique to implement the E and R flags in the REQUESTED-PROPS attribute (see below).

If the request contains a REQUESTED-PROPS attribute with the E flag set, then the server looks for an even port number to use for the relayed transport address.
If the request contains a REQUESTED-PROPS attribute with both the E and R flags set, then the server looks for a pair of port numbers N and N+1 on the same IP address, where N is even. Port N is used in the current allocation, while the relayed transport address with port N+1 is assigned a token and reserved for a future allocation. The server MUST hold this reservation for at least 30 seconds, and MAY choose to hold longer (e.g. until the allocation with port N expires). The server then includes the token in a RESERVATION-TOKEN attribute in the success response.
If the request contains a RESERVATION-TOKEN, the server uses the previously-reserved transport address corresponding to the included token (if it is still available).

> NOTE: The port N+1 reservation is a global reservation and is not specific to a particular allocation, since the Allocate request containing the RESERVATION-TOKEN will use a different 5-tuple and will create a different allocation. The 5-tuple for the subsequent Allocate request can be any allowed 5-tuple; the subsequent Allocate request can use a 5-tuple with a different client IP address and port, a different transport protocol, and even different server IP address and port (provided, of course, that the server IP address and port is one that the server is listening for TURN requests on).

Otherwise (i.e., the E and R flags are not set, and RESERVATION-TOKEN is not included), the server allocates any port in the range described above.
The server determines the initial value of the time-to-expire field as follows. If the request contains a LIFETIME attribute, and the proposed lifetime value is greater than the default lifetime, and the proposed lifetime value is otherwise acceptable to the server, then the server uses that value. Otherwise, the server uses the default lifetime. It is RECOMMENDED that the server impose a maximum lifetime of no more than 3600 seconds (1 hour). Servers that implement allocation quotas or charge users for allocations in some way may wish to use a smaller maximum lifetime (perhaps as small as the default lifetime) to more quickly remove orphaned allocations (that is, allocations where the corresponding client has crashed or terminated or the client connection has been lost for some reason). Also note that the time-to-expire is recomputed with each successful Refresh request, and thus the value computed here applies only until the first refresh.
Once the allocation is created, the server replies with a success response. The success response contains:

*A RELAYED-ADDRESS attribute containing the relayed transport address;

*A LIFETIME attribute containing the current value of the time-to-expire timer;

*A RESERVATION-TOKEN attribute (if a second relayed transport address was reserved).

*An XOR-MAPPED-ADDRESS attribute containing the client's IP address and port (from the 5-tuple);

> NOTE: The XOR-MAPPED-ADDRESS attribute is included in the response as a convenience to the client. TURN itself does not make use of this value, but clients running ICE can often need this value and can thus avoid having to do an extra Binding transaction with some STUN server to learn it.

The response (either success or error) is sent back to the client on the 5-tuple.

6.3. Receiving an Allocate Response

TOC

If the client receives a success response, then it MUST check that the relayed transport address is in an address family that the client understands and is prepared to deal with. This specification only covers the case where the relayed transport address is of the same address family as the client's transport address. If the relayed transport address is not in an address family that the client is prepared to deal with, then the client MUST delete the allocation (Section 7 (Refreshing an Allocation)) and MUST NOT attempt to create another allocation on that server until it believes the mismatch has been fixed.

> The IETF is currently considering mechanisms for transitioning between IPv4 and IPv6 that could result in a client originating an Allocate request over IPv4, but the request would arrive at the server over IPv6, or vica-versa. Hence the importance of this check.

Otherwise, the client creates its own copy of the allocation data structure to track what is happening on the server. In particular, the client needs to remember the actual lifetime received back from the server, rather than the value sent to the server in the request. The client must also remember the 5-tuple used for the request and the username and password it used to authenticate the request to ensure that it reuses them for subsequent messages. The client also needs to track the channels and permissions it establishes on the server.
The client will probably wish to send the relayed transport address to peers (using some method not specified here) so the peers can communicate with it. The client may also wish to use the server-

reflexive address it receives in the XOR-MAPPED-ADDRESS attribute in its ICE processing.

If the client receives an error response, then the processing depends on the actual error code returned:

*(Request timed out): There is either a problem with the server, or a problem reaching the server with the chosen transport. The client MAY choose to try again using a different transport (e.g., TCP instead of UDP), or the client MAY try a different server.

*400 (Bad Request): The server believes the client's request is malformed for some reason. The client MAY notify the user or operator and SHOULD NOT retry the same request with this server until it believes the problem has been fixed. The client MAY try a different server.

*401 (Unauthorized): If the client has followed the procedures of the Long-Term Credential mechanism and still gets this error, then the server is not accepting the client's credentials. The client SHOULD notify the user or operator and SHOULD NOT send any further requests to this server until it believes the problem has been fixed. The client MAY try a different server.

*437 (Allocation Mismatch): This indicates that the client has picked a 5-tuple which the server sees as already in use or which was recently in use. One way this could happen is if an intervening NAT assigned a mapped transport address that was recently used by another allocation. The client SHOULD pick another client transport address and retry the Allocate request (using a different transaction id). The client SHOULD try three different client transport addresses before giving up on this server. Once the client gives up on the server, it SHOULD NOT try to create another allocation on the server for 2 minutes.

*441 (Wrong Credentials): The client should not receive this error in response to a Allocate request. The client MAY notify the user or operator and SHOULD NOT retry the same request with this server until it believes the problem has been fixed. The client MAY try a different server.

*442 (Unsupported Transport Address): The client should not receive this error in response to a request for a UDP allocation. The client MAY notify the user or operator and SHOULD NOT retry the same request with this server until it believes the problem has been fixed. The client MAY try a different server.

*486 (Allocation Quota Reached): The server is currently unable to create any more allocations with this username. The client SHOULD

wait at least 1 minute before trying to create any more allocations on the server. The client MAY try a different server.

*508 (Insufficient Port Capacity): The server has no more relayed transport addresses avaiable, or has none with the requested properties, or the one that was reserved is no longer available. If the client is using either the REQUESTED-PROPS or the RESERVATION-TOKEN attribute, then the client MAY choose to remove or modify this attribute and try again immediately. Otherwise, the client SHOULD wait at least 1 minute before trying to create any more allocations on this server. The client MAY try a different server.

If the error response contains an ALTERNATE-SERVER attribute, and the client elects to try a different server, the the client SHOULD try the alternate server specified in that attribute (while obeying the rules in [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)," July 2008.) for avoiding redirection loops) before trying any other servers found using the SRV procedures of [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)," July 2008.).

---

7. Refreshing an Allocation

TOC

A Refresh transaction can be used to either (a) refresh an existing allocation and update its time-to-expire, or (b) delete an existing allocation.
If a client wishes to continue using an allocation, then the client MUST refresh it before it expires. It is suggested that the client refresh the allocation roughly 1 minute before it expires. If a client no longer wishes to use an allocation, then it SHOULD explicitly delete the allocation. A client MAY also change the time-to-expire of an allocation at any time for other reasons.

7.1. Sending a Refresh Request

TOC

If the client wishes to immediately delete an existing allocation, it includes a LIFETIME attribute with a value of 0. All other forms of the request refresh the allocation.
The Refresh transaction updates the time-to-expire timer of an allocation. If the client wishes the server to set the time-to-expire timer to something other than the default lifetime, it includes a LIFETIME attribute with the requested value. The server then computes a

new time-to-expire value in the same way as it does for an Allocate transaction, with the exception that a requested lifetime of 0 causes the server to immediately delete the allocation.
The Refresh transaction is sent on the 5-tuple for the allocation.

7.2. Receiving a Refresh Request

TOC

When the server receives a Refresh request, it processes it as follows:

1. The server checks the credentials of the request, as per the Long-Term Credential mechanism of [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)," July 2008.).

2. The server computes a value called the "desired lifetime" as follows: If the request contains a LIFETIME attribute and the attribute value is 0, then the desired lifetime is 0. Otherwise, if the request contains a LIFETIME attribute and the attribute value is greater than the default lifetime, and if the attribute value is otherwise acceptable to the server, then the the desired lifetime is the attribute value. Otherwise the desired lifetime is the default value.

3. The processing then depends on whether or not the 5-tuple corresponds to an existing allocation:

   * If there is no existing allocation and the desired lifetime is 0, then the request suceeeds (as it is OK to delete a non-existent allocation).

   * If there is no existing allocation and the desired lifetime is non-zero, then the server rejects the request with a 437 Allocation Mismatch error.

   * If there is an existing allocation and the desired lifetime is 0, then the request succeeds and the allocation is deleted.

   * If there is an existing allocation and the desired lifetime is non-zero, then the request succeeds and the allocation's time-to-expiry is set to the desired lifetime

   If the request succeeds, then server sends a success response containing:

   * A LIFETIME attribute containing the current value of the time-to-expire timer.

If the Refresh request is carried over UDP, then it is possible that it can be retransmitted. The server need not do anything special to handle this case since it is OK to delete a non-existent allocation and it is also OK to refresh an existing allocation twice in rapid succession.

7.3. Receiving a Refresh Response

TOC

If the client receives a success response to its Refresh request, it updates its copy of the allocation data structure with the time-to-expire value contained in the response.
If the client receives an 437 (Allocation Mismatch) error response to its Refresh request, then it must consider the allocation as having expired, as described in Section 4 (General Behavior). All other errors indicate a software error on the part of either the client or the server.

---

8. Permissions

TOC

For each allocation, the server keeps a list of zero or more permissions. Each permission consists of an IP address which uniquely identifies the permission, and an associated time-to-expiry. The IP address describes a peer that is allowed to send data to the client, and the time-to-expiry is the number of seconds until the permission expires.
Various events, as described in subsequent sections, can cause a permission for a given IP address to be installed or refreshed. This causes one of two things to happen:

- If no permission for that IP address exists, then a permission is created with the given IP address and a time-to-expiry equal to the default permission lifetime.

- If a permission for that IP address already exists, then the lifetime for that permission is reset to the default permission lifetime.

The default permission lifetime MUST be 300 seconds (= 5 minutes). Each permission’s time-to-expire decreases down once per second until it reaches 0, at which point the permission expires and is deleted.
When a UDP datagram arrives at the relayed transport address for the allocation, the server checks the list of permissions for that allocation. If there is a permission with an IP address that is equal to the source IP address of the UDP datagram, then the UDP datagram can be relayed to the client. Otherwise, the UDP datagram is silently

discarded. Note that only IP addresses are compared; port numbers are irrelevant.
The permissions for one allocation are totally unrelated to the permissions for a different allocation. If an allocation expires, all its permissions expire with it.

> NOTE: Though TURN permissions expire after 5 minutes, many NATs deployed at the time of publication expire their UDP bindings considerably faster. Thus an application using TURN will probably wish to send some sort of keep-alive traffic at a much faster rate. Applications using ICE should follow the keep-alive guidelines of ICE [I-D.ietf-mmusic-ice] (Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," October 2007.), and applications not using ICE are advised to do something similar.

---

9. Send and Data Indications

TOC

TURN supports two ways to send and receive data from peers. This section describes the use of Send and Data indications, while Section 10 (Channels) describes the use of the Channel Mechanism.

9.1. Sending a Send Indication

TOC

A client can use a Send indication to pass data to the server for relaying to a peer. A client can also use a Send indication without a DATA attribute to install or refresh a permission for the specified IP address. A client may use a Send indication to send data to a peer even if a channel is bound to that peer.
When forming a Send indication, the client MUST include a PEER-ADDRESS attribute and MAY include a DATA attribute. If the DATA attribute is included, then the DATA attribute contains the actual application data to be sent to the peer, and the PEER-ADDRESS attribute contains the transport address of the peer to which the data is to be sent. If the DATA attribute is not present, then the PEER-ADDRESS attribute contains the IP address for which a permission is to be installed or refreshed; in this case the port specified in the attribute is ignored.
Note that no authentication attributes are included, since indications cannot be authenticated using the Long-Term Credential mechanism.
The Send indication MUST be sent using the same 5-tuple used for the original allocation.

9.2. Receiving a Send Indication

TOC

When the server receives a Send indication, it processes it as follows. If the received Send indication contains a DATA attribute, then it forms a UDP datagram as follows:

- the source transport address is the relayed transport address of the allocation, where the allocation is determined by the 5-tuple on which the Send indication arrived;
- the destination transport address is taken from the PEER-ADDRESS attribute;
- the data following the UDP header is the contents of the value field of the DATA attribute.

The resulting UDP datagram is then sent to the peer. If any errors are detected during this process (e.g., the Send indication does not contain a PEER-ADDRESS attribute), the received indication is silently discarded and no UDP datagram is sent.
Clients are not allowed to use Send indications to send ICMP messages to peers. Thus the server MUST silently ignore a Send indication containing the ICMP attribute.
When the server receives a valid Send indication, either with or without a DATA attribute, it also installs or refreshes a permission for the IP address contained in the PEER-ADDRESS attribute (see Section 8 (Permissions)).

9.3. Receiving a UDP Datagram

TOC

When the server receives a UDP datagram at a currently allocated relayed transport address, the server looks up the allocation associated with the relayed transport address. It then checks to see if relaying is permitted, as described in Section 8 (Permissions).
If relaying is permitted, then the server checks if there is a channel bound to the peer that sent the UDP datagram (see Section 10 (Channels)). If a channel is bound, then processing proceeds as described in Section 10.7 (Relaying Data from the Peer).
If relaying is permitted but no channel is bound to the peer, then the server forms and sends a Data indication. The Data indication MUST contain both a PEER-ADDRESS and a DATA attribute and MUST NOT contain an ICMP attribute. The DATA attribute is set to the value of the ‘data octets’ field from the datagram, and the PEER-ADDRESS attribute is set to the source transport address of the received UDP datagram. The Data indication is then sent on the 5-tuple associated with the allocation.

9.4. Receiving a Data Indication

TOC

When the client receives a Data indication, it checks that the Data indication contains both a PEER-ADDRESS and a DATA attribute, and discards the indication if it does not.
The client then checks for the presence of the ICMP attribute. If it is present, the Data indication contains an ICMP message as described in Section 11 (IP and ICMP).
If the Data indication does not contain an ICMP attribute, the client delivers the data octets inside the DATA attribute to the application, along with an indication that they were received from the peer whose transport address is given by the PEER-ADDRESS attribute.

10. Channels

TOC

Channels provide a way for the client and server to send application data using ChannelData messages, which have less overhead than Send and Data indications.
Channel bindings are always initiated by the client. The client can bind a channel to a peer at any time during the lifetime of the allocation. The client may bind a channel to a peer before exchanging data with it, or after exchanging data with it (using Send and Data indications) for some time, or may choose never to bind a channel it. The client can also bind channels to some peers while not binding channels to other peers.
Channel bindings are specific to an allocation, so that a binding in one allocation has no relationship to a binding in any other allocation. If an allocation expires, all its channel bindings expire with it.
A channel binding consists of:

- A channel number;
- A transport address (of the peer);
- A time-to-expiry timer.

Within the context of an allocation, a channel binding is uniquely identified either by the channel number or by the transport address. Thus the same channel cannot be bound to two different transport addresses, nor can the same transport address be bound to two different channels.
A channel binding lasts for 10 minutes unless refreshed. Refreshing the binding (by the server receiving either a ChannelBind request rebinding the channel to the same peer, or by the server receiving a ChannelData

message on that channel) resets the time-to-expire timer back to 10 minutes. When the channel binding expires, the channel becomes unbound and available for binding to a different transport address.
When binding a channel to a peer, the client SHOULD be prepared to receive ChannelData messages on the channel from the server as soon as it has sent the ChannelBind request. Over UDP, it is possible for the client to receive ChannelData messages from the server before it receives a ChannelBind success response.
In the other direction, the client MAY elect to send ChannelData messages before receiving the ChannelBind success response. Doing so, however, runs the risk of having the ChannelData messages dropped by the server if the ChannelBind request does not succeed for some reason (e.g., packet lost if the request is sent over UDP, or the server being unable to fulfill the request). A client that wishes to be safe should either queue the data, or use Send indications until the channel binding is confirmed.

10.1. Sending a ChannelBind Request

TOC

A channel binding is created using a ChannelBind transaction. A channel binding can also be refreshed using a ChannelBind transaction.
To initiate the ChannelBind transaction, the client forms a ChannelBind request. The channel to be bound is specified in a CHANNEL-NUMBER attribute, and the peer's transport address is specified in a PEER-ADDRESS attribute. Section 10.2 (Receiving a ChannelBind Request) describes the restrictions on these attributes.
Note that rebinding a channel to the same transport address that it is already bound to provides a way to refresh a channel binding without sending data to the peer.
Once formed, the ChannelBind request is sent using the 5-tuple for the allocation.

10.2. Receiving a ChannelBind Request

TOC

When the server receives a ChannelBind request, it checks the following:

*The request contains both a CHANNEL-NUMBER and a PEER-ADDRESS attribute;

*The channel number is in the range 0x4000 to 0xFFFE (inclusive);

*The channel number is not currently bound to a different transport address (same transport address is OK);

*The transport address is not currently bound to a different channel number.

If any of these tests fail, the server replies with an error response with error code 400 "Bad Request". Otherwise, the ChannelBind request is valid and the server replies with a ChannelBind success response. There are no required attributes in a ChannelBind response.
If ChannelBind request is valid, then the server creates or refreshes the channel binding using the channel number in the CHANNEL-ADDRESS attribute and the transport address in the PEER-ADDRESS attribute. The server also installs or refreshes a permission for the IP address in the PEER-ADDRESS attribute.

10.3. Receiving a ChannelBind Response

TOC

When the client receives a successful ChannelBind response, it updates its data structures to record that the channel binding is now active.

10.4. The ChannelData Message

TOC

The ChannelData message is used to carry application data between the client and the server. It has the following format:

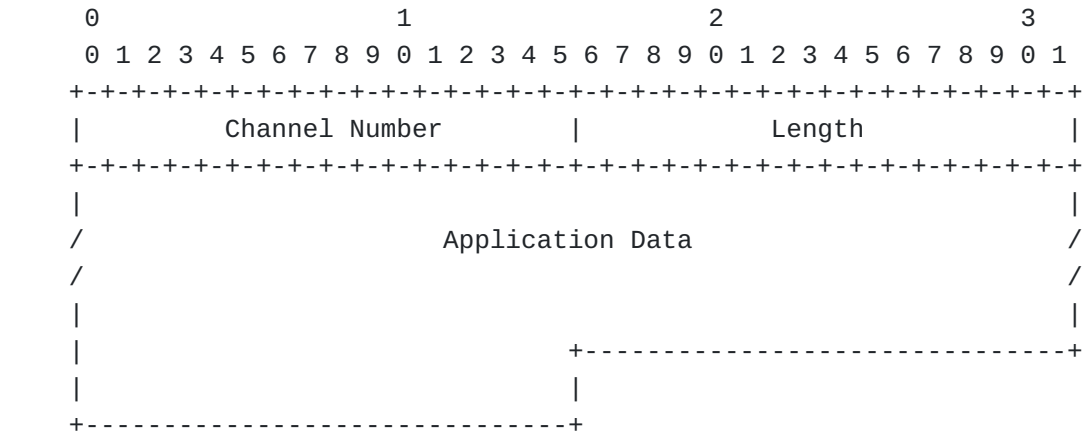

The Channel Number field specifies the number of the channel on which the data is traveling, and thus the address of the peer that is sending or is to receive the data. The channel number MUST be in the range 0x4000 – 0xFFFF, with channel number 0xFFFF being reserved for possible future extensions.
Channel numbers 0x0000 – 0x3FFF cannot be used because bits 0 and 1 are used to distinguish ChannelData messages from STUN-formatted messages (i.e., Allocate, Send, Data, ChannelBind, etc). STUN-formatted messages

always have bits 0 and 1 as “00”, while ChannelData messages use combinations “01”, “10”, and “11”.
The Length field specifies the length in bytes of the application data field (i.e., it does not include the size of the ChannelData header). Note that 0 is a valid length.
The Application Data field carries the data the client is trying to send to the peer, or that the peer is sending to the client.

10.5. Sending a ChannelData Message

TOC

Once a client has bound a channel to a peer, then when the client has data to send to that peer it may use either a ChannelData message or a Send indication; that is, the client is not obligated to use the channel when it exists and may freely intermix the two message types when sending data to the peer. The server, on the other hand, MUST use the ChannelData message if a channel has been bound to the peer.
The fields of the ChannelData message are filled in as described in Section 10.4 (The ChannelData Message).
Over stream transports, the ChannelData message MUST be padded to a multiple of four bytes in order to ensure the alignment of subsequent messages. The padding is not reflected in the length field of the ChannelData message, so the actual size of a ChannelData message (including padding) is (4 + Length) rounded up to the nearest multiple of 4. Over UDP, the padding is not required but MAY be included.
The ChannelData message is then sent on the 5-tuple associated with the allocation.

10.6. Receiving a ChannelData Message

TOC

The receiver of the ChannelData message uses bits 0 and 1 to distinguish it from STUN-formatted messages, as described in Section 10.4 (The ChannelData Message).
If the ChannelData message is received in a UDP datagram, and if the UDP datagram is too short to contain the claimed length of the ChannelData message (i.e., the UDP header length field value is less than the ChannelData header length field value + 4 + 8), then the message is silently discarded.
If the ChannelData message is received over TCP or over TLS over TCP, then the actual length of the ChannelData message is as described in Section 10.5 (Sending a ChannelData Message).
If the ChannelData message is received on a channel which is not bound to any peer, then the message is silently discarded.

If no errors are detected, the server relays the application data to the peer by forming a UDP datagram as follows:

- the source transport address is the relayed transport address of the allocation, where the allocation is determined by the 5-tuple on which the ChannelData message arrived;

- the destination transport address is the transport address to which the channel is bound;

- the data following the UDP header is the contents of the data field of the ChannelData message.

The resulting UDP datagram is then sent to the peer. Note that if the Length field in the ChannelData message is 0, then there will be no data in the UDP datagram, but the UDP datagram is still formed and sent.
If the ChannelData message is valid, then the server refreshes the channel binding, and also installs or refreshes a permission for the IP address part of the transport address to which the UDP datagram is sent (see Section 8 (Permissions)).

10.7. Relaying Data from the Peer

TOC

When the server receives a UDP datagram on the relayed transport address associated with an allocation, the server processes it as described in Section 9.3 (Receiving a UDP Datagram). If that section indicates that a ChannelData message should be sent (because there is a channel bound to the peer that sent to UDP datagram), then the server forms and sends a ChannelData message as described in Section 10.5 (Sending a ChannelData Message).

---

11. IP and ICMP

TOC

This section describes how the server sets various fields in the IP header when relaying between the client and the peer or vica-versa. It also describes how the server relays ICMP messages. The descriptions in this section apply: (a) when the server receives a Send indication or ChannelData message from the client and sends a UDP datagram to the peer, (b) when the server receives a UDP datagram on the relayed-transport address and sends a Data indication or ChannelData message to the client, or (c) when the server receives an ICMP message. This section does not apply when the server sends TURN control messages.

The descriptions below have two parts: a preferred behavior and an alternate behavior. A Preserving allocation MUST implement the preferred behavior. A non-preserving allocation with UDP transport to the client SHOULD implement the preferred behavior, but if that is not possible for a particular field, then it SHOULD implement the alternative behavior. A non-preserving allocation with TCP or TLS transport to client SHOULD implement the alternate behavior, except where this conflicts with standard TCP or TLS behavior.

11.1. IP

TOC

This section describes the preferred and alternate behavior for various fields in the IP header.
Time to Live (IPv4) or Hop Count (IPv6)

> Preferred Behavior: If the incoming value is 0, then send an ICMP Time Exceeded message back to the sender. Otherwise set the outgoing Time to Live/Hop Count to one less than the incoming value.
>
> Alternate Behavior: Set the outgoing value to the default for outgoing packets.

Diff-Serv Code Point

> Preferred Behavior: Set the outgoing value to the incoming value, unless the server includes a differentiated services classifier and marker [RFC2474] (Nichols, K., Blake, S., Baker, F., and D. Black, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,” December 1998.).
>
> Alternate Behavior: Set the outgoing value to a fixed value, which by default is Best Effort unless configured otherwise.
>
> In both cases, if the server is immediately adjacent to a differentiated services classifier and marker, then DSCP MAY be set to any arbitrary value in the direction towards the classifier.

ECN

> Preferred Behavior: Set the outgoing value to the incoming value, UNLESS the server is doing Active Queue Management, the incoming ECN field is 01 or 10, and the server wishes to indicate that congestion has been experienced, in which case set the outgoing value to 11.

> Alternate Behavior: Set the outgoing value to 00 (ECN not supported)

Flow Label

> Preferred Behavior: Set the outgoing flow label to 0.
>
> Alternate Behavior: Same as the Preferred behavior.

IPv4 Fragmentation

> Preferred Behavior:
>
> > If the outgoing packet size does not exceed the outgoing link's MTU, then send the outgoing packet unfragmented. Set the DF bit in the outgoing packet to the value of the DF bit in the incoming packet, and set the other fragmentation fields (Identification, MF, Fragment Offset) as appropriate for a packet originating from the server.
> >
> > Otherwise, if the outgoing link's MTU is exceeded and the incoming DF bit is 0, then fragment the packet before sending. Set the outgoing DF to 0, and set the other fragmentation fields as appropriate for fragments originated from the server.
> >
> > Otherwise [link MTU exceeded and incoming DF set], drop the outgoing packet and send an ICMP message of type 3 code 4 ("fragmentation needed and DF set") to the sender of the incoming packet.
>
> Alternate Behavior: As described in the Preferred Behavior, except always assume the incoming DF bit is 0.

IPv6 Fragmentation

> Preferred Behavior:
>
> > If the incoming packet did not include a Fragmentation header and the outgoing packet size does not exceed the outgoing link's MTU, then send the outgoing packet without a Fragmentation header.
> >
> > If the incoming packet included a Fragment header and if the outgoing packet size (with a Fragmentation header included) does not exceed the outgoing link's MTU, then send the outgoing packet

> with a Fragmentation header. Set the fields of the Fragmentation header as appropriate for a packet originating from the server.
>
> If the incoming packet did not include a Fragmentation header and the outgoing packet size exceeds the outgoing link's MTU , then drop the outgoing packet and send an ICMP message of type 2 code 0 ("Packet too big") to the sender of the incoming packet. If the packet is being sent to the peer, then reduce the MTU reported in the ICMP message by 48 bytes to allow room for the overhead of a Data indication.
>
> Otherwise, if the link's MTU is exceeded and the incoming packet contained a Fragmentation header, then fragment the outgoing packet into fragments of no more than 1280 bytes. Set the fields of the Fragmentation header as appropriate for a packet originating from the server.

Alternate Behavior: As described in the Preferred Behavior, except always assume incoming packet has a Fragmentation header.

IPv4 Options

> Preferred Behavior: The outgoing packet is sent without any IPv4 options.
>
> Alternate Behavior: Same as preferred.

IPv6 Extention Headers

> Preferred Behavior: The outgoing packet is sent without any IPv6 extension headers, with the exception of the Fragmentation header as described above
>
> Alternate Behavior: Same as preferred.

11.2. ICMP

TOC

This sub-section describes the preferred behavior of ICMP relaying. The corresponding alternate behavior is to not relay ICMP messages.
When an ICMP message arrives at the server, the copy of the original IP packet present inside the ICMP message is examined. The server first checks that the original IP packet header is immediately followed by a UDP protocol header, such that the original source transport address

was X and the original destination transport address was Y. The server also checks that the type and code values in the ICMP header are one of those relayed (see below). Other ICMP messages are either ignored, or used by the server internally in an unspecified manner.
The server then checks if one of the following two cases applies:
Case 1: X is a relayed-transport-address currently assigned to an active allocation on the server, and there exists a permission for the IP address of Y in the allocation.
In this case, the original IP packet was traveling from the server to a peer, so the the server relays the ICMP message back to the client. The server creates a Data indication where the PEER-ADDRESS attribute contains Y, and the ICMP attribute contains the type and code from the incoming ICMP message, and the DATA attribute contains application data from the original IP packet starting AFTER the UDP header. The server SHOULD include as much application data as possible consistent with not exceeding a total IP packet size of either 576 bytes (for IPv4) or 1280 bytes (for IPv6).

> Note that there is no point in including the original IP or UDP header in the DATA attribute because those headers were generated by the server, not the client.

Case 2: There is an active allocation where X is the server transport address, Y is the client transport address, and UDP is used as transport between the client and the server. Furthermore, the packet after the UDP header is either (a) a ChannelData header which contains an active channel number in the allocation, or (b) a Data indication whose PEER-ADDRESS attribute contains an IP address for which there exists a permission in the allocation.
In this case, the original IP packet was traveling from the server to the client, so the server creates and sends an ICMP message to the peer. The outgoing ICMP message contains the type and code fields from the incoming ICMP message and then contains an approximation to the original IP packet sent from the peer to the server (the one the server was trying to relay to the client inside the ChannelData or Data indication). This approximation contains a synthesized IP header, a synthesized UDP header, and some application data. The synthesis is done as follows:

- The destination transport address is the relayed-transport-address of the allocation;

- The source transport address is the peer's transport address determined from either (a) the channel number or (b) the PEER-ADDRESS attribute;

- The application data is taken from either (a) the ChannelData message or (b) the DATA attribute. The server SHOULD include as

much application data as possible consistent with not exceeding either 576 bytes (for IPv4) or 1280 bytes (for IPv6).

The remaining fields in the IP and UDP headers are simply set to sensible values, since for most of them there is no way to reconstruct the original values.
The server SHOULD relay all ICMP type/code combinations and MUST relay at least the following combinations. For IPv4:

> Type 3, code 4: Fragmentation needed and DF set

For IPv6:

> Type 2, code <any>: Packet too big

Note that the ICMP attribute appears only in Data indications; the client cannot use the ICMP attribute in a Send indication to send ICMP messages to the peer.

---

12. New STUN Methods

TOC

This section lists the codepoints for the new STUN methods defined in this specification. See elsewhere in this document for the semantics of these new methods.

```
Request/Response Transactions
  0x003  :  Allocate
  0x004  :  Refresh
  0x009  :  ChannelBind

Indications
  0x006  :  Send
  0x007  :  Data
```

---

13. New STUN Attributes

TOC

This STUN extension defines the following new attributes:

```
0x000C: CHANNEL-NUMBER
0x000D: LIFETIME
0x0010: Reserved (was BANDWIDTH)
0x0012: PEER-ADDRESS
0x0013: DATA
0x0016: RELAYED-ADDRESS
0x0018: REQUESTED-PROPS
0x0019: REQUESTED-TRANSPORT
0x0022: RESERVATION-TOKEN
0x0030: ICMP
```

13.1. CHANNEL-NUMBER

TOC

The CHANNEL-NUMBER attribute contains the number of the channel. It is a 16-bit unsigned integer, followed by a two-octet RFFU (Reserved For Future Use) field which MUST be set to 0 on transmission and MUST be ignored on reception.

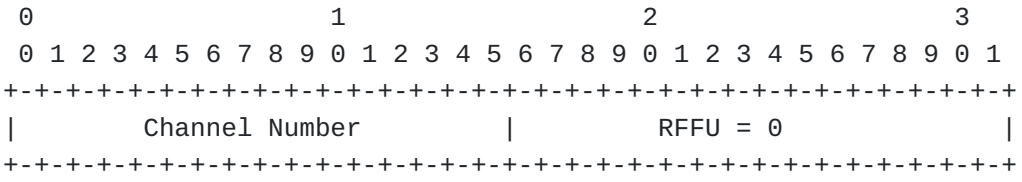

13.2. LIFETIME

TOC

The lifetime attribute represents the duration for which the server will maintain an allocation in the absence of a refresh. It is a 32-bit unsigned integral value representing the number of seconds remaining until expiration.

13.3. PEER-ADDRESS

TOC

The PEER-ADDRESS specifies the address and port of the peer as seen from the TURN server. It is encoded in the same way as XOR-MAPPED-ADDRESS.

13.4. DATA

TOC

The DATA attribute is present in all Data indications and most Send indications. The contents of DATA attribute is the application data (that is, the data that would immediately follow the UDP header if the data was been sent directly between the client and the peer).

13.5. RELAYED-ADDRESS

TOC

The RELAYED-ADDRESS is present in Allocate responses. It specifies the address and port that the server allocated to the client. It is encoded in the same way as XOR-MAPPED-ADDRESS.

13.6. REQUESTED-PROPS

TOC

This attribute allows the client to request that the allocation have certain properties, and by the server to indicate which properties are supported. The attribute is 32 bits long. Its format is:

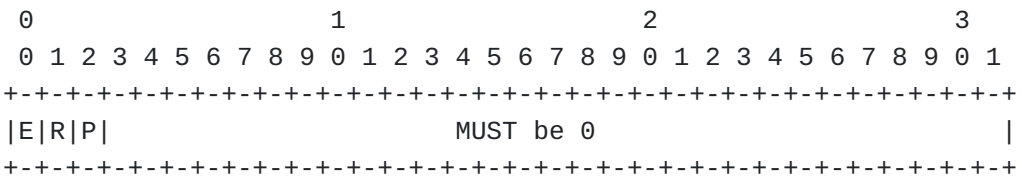

The first part of the attribute value contains a number of one-bit flags. These are:

E: If 1, the port number for the relayed-transport-address must be even. If 0, the port number can be even or odd.

R: If 1, the server must reserve the next highest port for a subsequent allocation. If 0, no such reservation is requested. If the client sets the R bit to 1, it MUST also set the E bit to 1 (however, the E bit may be 1 when the R bit is 0).

P: If 1, the allocation must be a Preserving allocation. If 0, the allocation can be either Preserving or Non-Preserving.

All these flags have the property that if the bit is 1, and the server cannot create an allocation that satisfies the request, then the Allocate request is rejected. To allow future TURN extensions to define new flags that also have this property, the client MUST set the rest of

the attribute to zero, and the server MUST fail the Allocate request if any bits which the server does not support are set to 1. By doing this, any new flags that are not recognized by the server will cause the Allocate request to fail.

13.7. REQUESTED-TRANSPORT

TOC

This attribute is used by the client to request a specific transport protocol for the allocated transport address. It has the following format:

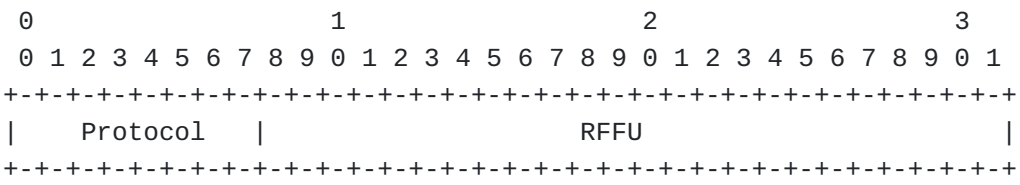

The Protocol field specifies the desired protocol. The codepoints used in this field are taken from those allowed in the Protocol field in the IPv4 header and the NextHeader field in the IPv6 header [Protocol-Numbers] (, “IANA Protocol Numbers Registry,” 2005.). This specification only allows the use of codepoint 17 (User Datagram Protocol).
The RFFU field MUST be set to zero on transmission and MUST be ignored on reception. It is reserved for future uses.

13.8. RESERVATION-TOKEN

TOC

The RESERVATION-TOKEN attribute contains a token that uniquely identifies a relayed transport address being held in reserve by the server. The server includes this attribute in a success response to tell the client about the token, and the client includes this attribute in a subsequent Allocate request to request the server use that relayed transport address for the allocation.
The attribute value is a 64-bit-long field containing the token value.

13.9. ICMP

TOC

This attribute is included by the server in a Data indication to indicate that the Data indication contains information from an ICMP

message that was received by the server. The attribute has the following format:

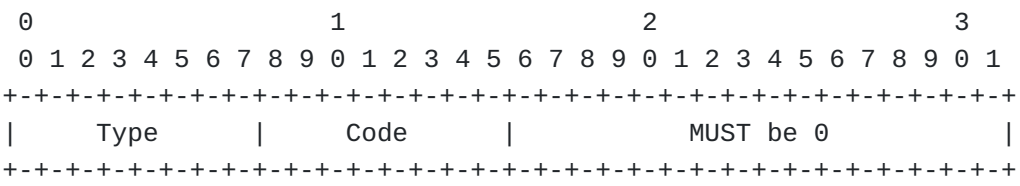

The Type and Code fields of the attribute are taken from the Type and Code fields in the ICMP message received by the server.

---

14. New STUN Error Response Codes

TOC

This document defines the following new error response codes:

437 (Allocation Mismatch): A request was received by the server that requires an allocation to be in place, but there is none, or a request was received which requires no allocation, but there is one.

441 (Wrong Credentials): The credentials in the (non-Allocate) request, though otherwise acceptable to the server, do not match those used to create the allocation.

442 (Unsupported Transport Protocol): The Allocate request asked the server to use a transport protocol between the server and the peer that the server does not support. NOTE: This does NOT refer to the transport protocol used in the 5-tuple.

486 (Allocation Quota Reached): No more allocations using this username can be created at the present time.

508 (Insufficient Port Capacity): The server has no more relayed transport addresses available right now, or has none with the requested properties, or the one that corresponds to the specified token is not available.

---

15. Security Considerations

TOC

TBD: Update this section to match changes to the TURN protocol.
TURN servers allocate resources to clients, in contrast to the Binding method defined in [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy,

R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.). Therefore, a TURN server may require the authentication and authorization of STUN requests. This authentication is provided by mechanisms defined in the STUN specification itself, in particular digest authentication.

Because TURN servers allocate resources, they can be susceptible to denial-of-service attacks. All Allocate transactions are authenticated, so that an unknown attacker cannot launch an attack. An authenticated attacker can generate multiple Allocate requests, however. To prevent a single malicious user from allocating all of the resources on the server, it is RECOMMENDED that a server implement a per user limit on the number of allocations that can active at one time. Such a mechanism does not prevent a large number of malicious users from each requesting a small number of allocations. Attacks such as these are possible using botnets, and are difficult to detect and prevent. Implementors of TURN should keep up with best practices around detection of anomalous botnet attacks.

A client will use the transport address learned from the RELAYED-ADDRESS attribute of the Allocate response to tell other users how to reach them. Therefore, a client needs to be certain that this address is valid, and will actually route to them. Such validation occurs through the message integrity checks provided in the Allocate response. They can guarantee the authenticity and integrity of the allocated addresses. Note that TURN is not susceptible to the attacks described in Section 12.2.3, 12.2.4, 12.2.5 or 12.2.6 of [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for (NAT) (STUN),” July 2008.) [[TODO: Update section number references to 3489bis]]. These attacks are based on the fact that a STUN server mirrors the source IP address, which cannot be authenticated. STUN does not use the source address of the Allocate request in providing the RELAYED-ADDRESS, and therefore, those attacks do not apply.

TURN attempts to adhere as closely as possible to common firewall policies, consistent with allowing data to flow. TURN has fairly limited applicability, requiring a user to explicitly authorize permission to receive data from a peer, one IP address at a time. Thus, it does not provide a general technique for externalizing sockets. Rather, it has similar security properties to the placement of an address-restricted NAT in the network, allowing messaging in from a peer only if the internal client has sent a packet out towards the IP address of that peer. This limitation means that TURN cannot be used to run, for example, SIP servers, NTP servers, FTP servers or other network servers that service a large number of clients. Rather, it facilitates rendezvous of NATted clients that use some other protocol, such as SIP, to communicate IP addresses and ports for communications.

Confidentiality of the transport addresses learned through Allocate transactions does not appear to be that important. If required, it can be provided by running TURN over TLS.

TURN does not and cannot guarantee that UDP data is delivered in sequence or to the correct address. As most TURN clients will only communicate with a single peer, the use of a single channel number will be very common. Consider an enterprise where Alice and Bob are involved in separate calls through the enterprise NAT to their corporate TURN server. If the corporate NAT reboots, it is possible that Bob will obtain the exact NAT binding originally used by Alice. If Alice and Bob were using identical channel numbers, Bob will receive unencapsulated data intended for Alice and will send data accidentally to Alice's peer. This is not a problem with TURN. This is precisely what would happen if there was no TURN server and Bob and Alice instead provided a (STUN) reflexive transport address to their peers. If detecting this misdelivery is a problem, the client and its peer need to use message integrity on their data.

Relay servers are useful even for users not behind a NAT. They can provide a way for truly anonymous communications. A user can cause a call to have its media routed through a TURN server, so that the user's IP addresses are never revealed.

Any relay addresses learned through an Allocate request will not operate properly with IPSec Authentication Header (AH) (Kent, S., "IP Authentication Header," December 2005.) [RFC4302] in transport or tunnel mode. However, tunnel-mode IPSec ESP (Kent, S., "IP Encapsulating Security Payload (ESP)," December 2005.) [RFC4303] should still operate.

---

16. IANA Considerations

TOC

Since TURN is an extension to STUN [I-D.ietf-behave-rfc3489bis] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)," July 2008.), the methods, attributes and error codes defined in this specification are new method, attributes, and error codes for STUN. This section directs IANA to add these new protocol elements to the IANA registry of STUN protocol elements.

The codepoints for the new STUN methods defined in this specification are listed in Section 12 (New STUN Methods).

The codepoints for the new STUN attributes defined in this specification are listed in Section 13 (New STUN Attributes).

The codepoints for the new STUN error codes defined in this specification are listed in Section 14 (New STUN Error Response Codes).

Extensions to TURN can be made through IETF consensus.

TOC

17. IAB Considerations

The IAB has studied the problem of "Unilateral Self Address Fixing", which is the general process by which a client attempts to determine its address in another realm on the other side of a NAT through a collaborative protocol reflection mechanism [RFC3424] (Daigle, L. and IAB, “IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation,” November 2002.). The TURN extension is an example of a protocol that performs this type of function. The IAB has mandated that any protocols developed for this purpose document a specific set of considerations.
TURN is an extension of the STUN protocol. As such, the specific usages of STUN that use the TURN extensions need to specifically address these considerations. Currently the only STUN usage that uses TURN is ICE (Rosenberg, J., “Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols,” October 2007.) [I-D.ietf-mmusic-ice].

---

18. Example

TOC

TBD

---

19. Open Issues

TOC

This section lists the known issues in this version of the specification.

1. Detecting in-use channels. Do we need a way for a client to determine if a channel is currently bound? Right now, the only way is to try to bind it to an address.

2. Public TURN servers. The spec currently hints (but does not say anything solid) that the way to run a publicly-accessable TURN server is to not require authentication. But perhaps a better way is to require authentication but have some unspecified method to allow any user to create an account on the server.

3. IPv6. Currently, TURN supports IPv4-to-IPv4 relaying, and IPv6-to-IPv6 relaying, but does not support IPv4-to-IPv6 relaying. To ensure this, a server requires that the family of the relayed address match that of the 5-tuple as seen by the server. However, some people would like to see a different rule.

4. ALTERNATE-SERVER and Anycast. The details of ALTERNATE-SERVER support are still under discussion. In particular, some people would like to use ALTERNATE-SERVER to support anycast discovery of a TURN server.

5. Authenticated Permission Refresh. Currently, permissions can be refreshed by unauthenticated Send indications and ChannelData messages. Some have suggested that this is a security issue.

6. PMTUD for non-preserving allocations. Some people would like a way to do PMTUD even if the allocation is non-preserving, and have suggested that a way for the client to indicate to the server (in a Send indication) that the DF bit should be set when sending to the peer might allow this.

7. Security. The security consideration section is out-of-date with the changes to the rest of the draft, and it has been suggested that TURN might require TLS to provide proper security. Updating the security consideration section will answer this question.

---

20. Changes from Previous Versions

TOC

Note to RFC Editor: Please remove this section prior to publication of this document as an RFC.
This section lists the changes between the various versions of this specification.

20.1. Changes from -08 to -09

TOC

*Added text to properly define the ICMP attribute. This attribute was introduced in TURN-08, but not fully defined due to an oversight. Clarified that the attribute can appear in a Data indication, but not a Send indication. Added text to the section on receiving a Data indication that points out that this attribute may be present.

*Changed the wording around the handling of the DSCP field to allow the server to set the DSCP to an arbitrary value if the next hop is a Diff-Serv classifier and marker.

*When the server generates a 508 response due to an unsupported flag in the REQUESTED-PROPS attribute, the server now includes

the REQUESTED-PROPS attribute in the response with all the flags it supports set to 1. This allows the client to see if the server does not understand one of its flags. Similarly, the client is now allowed to immediately retry the request if it modifies the included REQUESTED-PROPS attribute.

*Clarified that the REQUESTED-PROPS attribute can be used in conjunction with the RESERVATION-TOKEN attribute as long as both the E and R bits are 0. The spec previously contradicted itself on this point.

*Clarified that when the server receives a ChannelData message with a length field of 0, it sends a UDP Datagram to the peer that contains no application data.

*Rewrote some text around relaying incoming UDP Datagrams to avoid duplication of text in the Data indication and Channel sections.

*Added a note that points out that the on-going work on randomizing port allocations [I-D.ietf-tsvwg-port-randomization] (Larsen, M. and F. Gont, “Transport Protocol Port Randomization Recommendations,” April 2010.) may be applicable to TURN.

*Clarified that the Allocate request containing a RESERVATION-TOKEN attribute can use any 5-tuple, and that 5-tuple need not have any specific relationship to the 5-tuple of the Allocate request that created the reservation.

*Added a note that discusses retransmitted Allocate requests over UDP where the first request receives a failure response, but the second receives a success response. The server may elect to remember transmitted failure responses to avoid this situation.

*Added text about the usage of the SOFTWARE-TYPE attribute (formerly known as the SERVER attribute) in TURN messages.

*Rewrote the text in the Overview that motivates why TURN supports TCP and TLS between the client and the server. The previous text had been identified by various readers as inadequate and misleading.

*Rewrote the section how a server handles a Refresh request to clarify processing in various error conditions. The new text makes it clear that it is OK to delete a non-existent allocation. It also clarifies how to handle retransmissions of Refresh requests over UDP.

*Renamed the "RELAY-ADDRESS" attribute to "RELAYED-ADDRESS", since the text consistently uses the term "relayed transport address" for the concept and ICE uses the term "relayed candidate".

*Changed the codepoint assigned to the error code "Wrong Credentials" from 438 to 441 to avoid a conflict with the "Stale Nonce" error code of STUN.

*Changed the text to consistently use non-capitalized "request", "response" and "indication", except in headings, error code names, etc.

*Added a note mentioning that TURN packets can be demuxed from other packets arriving on the same socket by looking at the 5-tuple of the arriving packet.

*Clarified that there are no required attributes is a ChannelBind success response.

20.2. Changes from -07 to -08

TOC

*Removed the BANDWIDTH attribute and all associated text (including error code 507 "Insufficient Bandwidth Capacity"), as the requirements for this feature were not clear and it was felt the feature could be easily added later.

*Changed the format of the REQUESTED-PROPS attribute from a one-byte field to a set of bit flags. Changed the semantics of the unused portion of the value from RFFU to "MUST be 0" to give a more desirable behavior when new flags are defined.

*Introduced the concept of Preserving vs. Non-Preserving allocations. As a result, completely revamped the rules for how to set the fields in the IP header, and added rules for relaying ICMP messages when the allocation is Preserving.

20.3. Changes from -06 to -07

TOC

*Rewrote the General Behavior section, making various changes in the process.

*Changed the usage of authentication from MUST to SHOULD.

*Changed the requirement that subsequent requests use the same username and password from MUST to SHOULD to allow for the possibility of changing the credentials using some unspecified mechanism.

*Introduced a 438 (Wrong Credentials) error which is used when a non-Allocate request authenticates but does not use the same username and password as the Allocate request. Having a separate error code for this case avoids the client being confused over what the error actually is.

*The server must now prevent the relayed transport address and the 5-tuple from being reused in different allocations for 2 minutes after the allocation expires.

*Changed the usage of FINGERPRINT from MUST NOT to MAY, to allow for the possible multiplexing of TURN with some other protocol.

*Rewrote much of the section on Allocations, splitting it into three new sections (one on allocations in general, one on creating an allocation, and one on refreshing an allocation).

*Replaced the mechanism for requesting relayed transport addresses with specific properties. The new mechanism is less powerful: a client can request an even port, or a pair of ports, but cannot request a single odd port or a specific port as was possible under the old mechanism. Nor can the client request a specific IP address.

*Changed the rules for handling ALTERNATE-SERVER, removing the requirement that the referring server have "positive knowledge" about the state of the alternate server. The new rules instead rely on text in STUN to prevent referral loops.

*Changed the rules for allocation lifetimes. Allocations lifetimes are now a minimum of 10 minutes; the client can ask for longer values, but requests for shorter values are ignored. The text now recommends that the client refresh an allocation one minute before it expires.

*Put in temporary procedures for handling the BANDWIDTH attribute, modelled on the LIFETIME attribute. These procedures are mostly placeholders and likely to change in the next revision.

*Added a detailed description of how a client reacts to the various errors it can receive in reply to an Allocate request. This replaces the various descriptions that were previously scattered throughout the document, which were inconsistent and sometimes contradictory.

*Added a new section that gives the normative rules for permissions.

*Changed the rules around permission lifetimes. The text used to recommend a value of one minute; it MUST now be 5 minutes.

*Removed the errors "Channel Missing or Invalid", "Peer Address Missing or Invalid" and "Lifetime Malformed or Invalid" and used 400 "Bad Request" instead.

*Rewrote portions of the section on Send and Data indications and the section on Channels to try to make the client vs. server behavior clearer.

*Channel bindings now expire after 10 minutes, and must be refreshed to keep them alive.

*Binding a channel now installs or refreshes a permission for the IP address of corresponding peer.

*Changed the wording describing the situation when the client sends a ChannelData message before receiving the ChannelBind success response. -06 said that client SHOULD NOT do this; -07 now says that a client MAY, but describes the consequences of doing it.

*Added a section discussing the setting of fields in the IP header.

*Replaced the REQUESTED-PORT-PROPS attribute with the REQUESTED-PROPS attribute that has a different format and semantics, but reuses the same code point.

*Replaced the REQUESTED-IP attribute with the RESERVATION-TOKEN attribute, which has a different format and semantics, but reuses the same code point.

*Removed error codes 443 and 444, and replaced them with 508 (Insufficient Port Capacity). Also changed the error text for code 507 from "Insufficient Capacity" to "Insufficient Bandwidth Capacity".

20.4. Changes from -05 to -06

TOC

*Changed the mechanism for allocating channels to the one proposed by Eric Rescorla at the Dec 2007 IETF meeting.

*Removed the framing mechanism (which was used to frame all messages) and replaced it with the ChannelData message. As part of this change, noted that the demux of ChannelData messages from TURN messages can be done using the first two bits of the message.

*Rewrote the sections on transmitted and receiving data as a result of the above to changes, splitting it into a section on Send and Data indications and a separate section on channels.

*Clarified the handling of Allocate request messages. In particular, subsequent Allocate request messages over UDP with the same transaction id are not an error but a retransmission.

*Restricted the range of ports available for allocation to the Dynamic and/or Private Port range, and noted when ports outside this range can be used.

*Changed the format of the REQUESTED-TRANSPORT attribute. The previous version used 00 for UDP and 01 for TCP; the new version uses protocol numbers from the IANA protocol number registry. The format of the attribute also changed.

*Made a large number of changes to the non-normative portion of the document to reflect technical changes and improve the presentation.

*Added the Issues section.

20.5. Changes from -04 to -05

TOC

*Removed the ability to allocate addresses for TCP relaying. This is now covered in a separate document. However, communication between the client and the server can still run over TCP or TLS/TCP. This resulted in the removal of the Connect method and the TIMER-VAL and CONNECT-STAT attributes.

*Added the concept of channels. All communication between the client and the server flows on a channel. Channels are numbered 0..65535. Channel 0 is used for TURN messages, while the remaining channels are used for sending unencapsulated data to/from a remote peer. This concept adds a new Channel Confirmation method and a new CHANNEL-NUMBER attribute. The new attribute is also used in the Send and Data methods.

*The framing mechanism formally used just for stream-oriented transports is now also used for UDP, and the former Type and Reserved fields in the header have been replaced by a Channel Number field. The length field is zero when running over UDP.

*TURN now runs on its own port, rather than using the STUN port. The use of channels requires this.

*Removed the SetActiveDestination concept. This has been replaced by the concept of channels.

*Changed the allocation refresh mechanism. The new mechanism uses a new Refresh method, rather than repeating the Allocation transaction.

*Changed the syntax of SRV requests for secure transport. The new syntax is "_turns._tcp" rather than the old "_turn._tls". This change mirrors the corresponding change in STUN SRV syntax.

*Renamed the old REMOTE-ADDRESS attribute to PEER-ADDRESS, and changed it to use the XOR-MAPPED-ADDRESS format.

*Changed the RELAY-ADDRESS attribute to use the XOR-MAPPED-ADDRESS format (instead of the MAPPED-ADDRESS format)).

*Renamed the 437 error code from "No Binding" to "Allocation Mismatch".

*Added a discussion of what happens if a client's public binding on its outermost NAT changes.

*The document now consistently uses the term "peer" as the name of a remote endpoint with which the client wishes to communicate.

*Rewrote much of the document to describe the new concepts. At the same time, tried to make the presentation clearer and less repetitive.

---

21. Open Issues

TOC

NOTE to RFC Editor: Please remove this section prior to publication of this document as an RFC.
Bandwidth: How should bandwidth be specified? What are the right rules around bandwidth?
Alternate Server: Do we still want this mechanism? Is the current proposal acceptable? Note that the usage of the ALTERNATE-SERVER attribute in this document is inconsistent with its usage in STUN. In STUN, if the ALTERNATE-SERVER attribute is used, then the error that the server would otherwise generate is replaced by a 300 (Try Alternate) code. In this document, the 300 error code is not used, and the server returns an appropriate error code and then includes the ALTERNATE-SERVER attribute in the response. In this way, the client can see the actual error code, rather than always seeing error code 300,

and can thus make a more intelligent decision on whether it wishes to try the alternate server.
Public TURN servers: The text currently says that a server "SHOULD" use the Long-Term Credential mechanism, with the unstated idea that a public TURN server would not use it. But this really weakens the security of TURN. Is there a better way to allow public servers? Or should we just drop the notion of a public server entirely?

---

22. Acknowledgements

TOC

The authors would like to thank the various participants in the BEHAVE working group for their many comments on this draft. Marc Petit-Huguenin, Remi Denis-Courmont, Derek MacDonald, Cullen Jennings, Lars Eggert, Magnus Westerlund, and Eric Rescorla have been particularly helpful, with Eric also suggesting the channel allocation mechanism, and Cullen suggesting the REQUESTED-PROPS mechanism. Christian Huitema was an early contributor to this document and was a co-author on the first few drafts. Finally, the authors would like to thank Dan Wing for both his contributions to the text and his huge help in restarting progress on this draft after work had stalled.

---

23. References

TOC

23.1. Normative References

TOC

[I-D.ietf-behave-rfc3489bis]	Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)," draft-ietf-behave-rfc3489bis-18 (work in progress), July 2008 (TXT).
[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2474]	Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, December 1998 (TXT, HTML, XML).
[RFC3697]	Rajahalme, J., Conta, A., Carpenter, B., and S. Deering, "IPv6 Flow Label Specification," RFC 3697, March 2004 (TXT).

23.2. Informative References

TOC

[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," BCP 5, RFC 1918, February 1996 (TXT).
[RFC1981]	McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6," RFC 1981, August 1996 (TXT).
[RFC3264]	Rosenberg, J. and H. Schulzrinne, "An Offer/ Answer Model with Session Description Protocol (SDP)," RFC 3264, June 2002 (TXT).
[RFC4302]	Kent, S., "IP Authentication Header," RFC 4302, December 2005 (TXT).
[RFC4303]	Kent, S., "IP Encapsulating Security Payload (ESP)," RFC 4303, December 2005 (TXT).
[RFC3424]	Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation," RFC 3424, November 2002 (TXT).
[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 (TXT).
[RFC4787]	Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP," BCP 127, RFC 4787, January 2007 (TXT).
[I-D.ietf-behave-turn-tcp]	Perreault, S. and J. Rosenberg, "Traversal Using Relays around NAT (TURN) Extensions for TCP Allocations," draft-ietf-behave-turn-tcp-06 (work in progress), March 2010 (TXT).
[I-D.ietf-behave-turn-ipv6]	Camarillo, G., Novo, O., and S. Perreault, "Traversal Using Relays around NAT (TURN) Extension for IPv6," draft-ietf-behave-turn-ipv6-09 (work in progress), March 2010 (TXT).
[I-D.ietf-tsvwg-udp-guidelines]	Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers," draft-ietf-tsvwg-udp-guidelines-11 (work in progress), October 2008 (TXT).
[I-D.ietf-tsvwg-port-randomization]	Larsen, M. and F. Gont, "Transport Protocol Port Randomization Recommendations," draft-ietf-tsvwg-port-randomization-07 (work in progress), April 2010 (TXT).
[RFC1191]	Mogul, J. and S. Deering, "Path MTU discovery," RFC 1191, November 1990 (TXT).

[RFC4821]	Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” RFC 4821, March 2007 (TXT).
[RFC1928]	Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, “SOCKS Protocol Version 5,” RFC 1928, March 1996 (TXT).
[Port-Numbers]	“IANA Port Numbers Registry.”
[Protocol-Numbers]	“IANA Protocol Numbers Registry,” 2005.

Authors' Addresses

TOC

	Jonathan Rosenberg
	Cisco Systems, Inc.
	Edison, NJ
	USA
Email:	jdrosen@cisco.com
URI:	http://www.jdrosen.net
	Rohan Mahy
	Plantronics, Inc.
Email:	rohan@ekabal.com
	Philip Matthews
	(Unaffiliated)
Fax:	
Email:	philip_matthews@magma.ca
URI:	

Full Copyright Statement

TOC

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.