BEHAVE WG	J. Rosenberg
Internet-Draft	Cisco
Intended status: Standards Track	R. Mahy
Expires: October 14, 2009	(Unaffiliated)
	P. Matthews
	Alcatel-Lucent
	April 12, 2009

TOC

Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)
draft-ietf-behave-turn-14

Status of this Memo

Copyright Notice

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). 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. TURN differs from some other relay control protocols in that it allows a client to communicate with multiple peers using a single relay address. The TURN protocol was designed to be used as part of the ICE (Interactive Connectivity Establishment) approach to NAT traversal, though it can be also used without ICE.

Table of Contents

---

1. Introduction

TOC

A host behind a NAT may wish to exchange packets with other hosts, some of which may also be behind NATs. To do this, the hosts involved can use 'Hole Punching' techniques (see [RFC5128] (Srisuresh, P., Ford, B., and D. Kegel, “State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs),” March 2008.)) in an attempt discover a direct communication path; that is, a communication path that goes from host to another through intervening NATs and routers, but does not traverse any relays.

As described in [RFC5128] (Srisuresh, P., Ford, B., and D. Kegel, “State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs),” March 2008.) and [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.), hole punching techniques will fail if both hosts are behind NATs that are not well-behaved. For example, if both hosts are behind NATs that have a mapping behavior of "address dependent mapping" or "address and port dependent mapping", then hole punching techniques generally fail.

When a direct communication path cannot be found, it is necessary to use the services of an intermediate host that acts as a relay for the packets. This relay typically sits in the public Internet and relays packets between two hosts that both sit behind NATs.

This specification defines a protocol, called TURN, that allows a host behind a NAT (called the TURN client) to request that another host

(called the TURN server) act as a relay. The client can arrange for the server to relay packets to and from certain other hosts (called peers) and can control aspects of how the relaying is done. The client does this by obtaining an IP address and port on the server, called the relayed-transport-address. When a peer sends a packet to the relayed-transport-address, the server relays the packet to the client. When the client sends a data packet to the server, the server relays it to the appropriate peer using the relayed-transport-address as the source.

A client using TURN must have some way to communicate the relayed-transport-address to its peers, and to learn each peer's IP address and port (more precisely, each peer's server-reflexive transport address, see Section 2 (Overview of Operation)). How this is done is out of the scope of the TURN protocol. One way this might be done is for the client and peers to exchange e-mail messages. Another way is for the client and its peers to use a special-purpose 'introduction' or 'rendezvous' protocol (see [RFC5128] (Srisuresh, P., Ford, B., and D. Kegel, "State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)," March 2008.) for more details).

If TURN is used 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.), then the relayed-transport-address and the IP addresses and ports of the peers are included in the ICE candidate information which the rendezvous protocol must carry. For example, if TURN and ICE are used as part of a multimedia solution using SIP [RFC3261] (Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol," June 2002.), then SIP serves the role of the rendezvous protocol, carrying the ICE candidate information inside the body of SIP messages. If TURN and ICE are used with some other rendezvous protocol, then [I-D.rosenberg-mmusic-ice-nonsip] (Rosenberg, J., "Guidelines for Usage of Interactive Connectivity Establishment (ICE) by non Session Initiation Protocol (SIP) Protocols," July 2008.) provides guidance on the services the rendezvous protocol must perform.

Though the use of a TURN server to enable communication between two hosts behind NATs is very likely to work, it comes at a high cost to the provider of the TURN server, since the server typically needs a high bandwidth connection to the Internet . As a consequence, it is best to use a TURN server only when a direct communication path cannot be found. When the client and a peer use ICE to determine the communication path, ICE will use hole punching techniques to search for a direct path first and only use a TURN server when a direct path cannot be found.

TURN was originally invented to support multimedia sessions signaled using SIP. Since SIP supports forking, TURN supports multiple peers per relayed-transport-address; 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 to make sure that TURN is suitable for other types of applications.
TURN was designed as one piece in the larger ICE approach to NAT traversal. Implementors of TURN are urged to investigate ICE and seriously consider using it for their application. However, it is possible to use TURN without ICE.
TURN is an extension to the STUN (Session Traversal Utilities for NAT [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.)) protocol. Most, though not all, TURN messages are STUN-formatted messages. A reader of this document should be familiar with STUN.

---

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. The client uses the server as a relay to send packets to these peers and to receive packets from these peers.

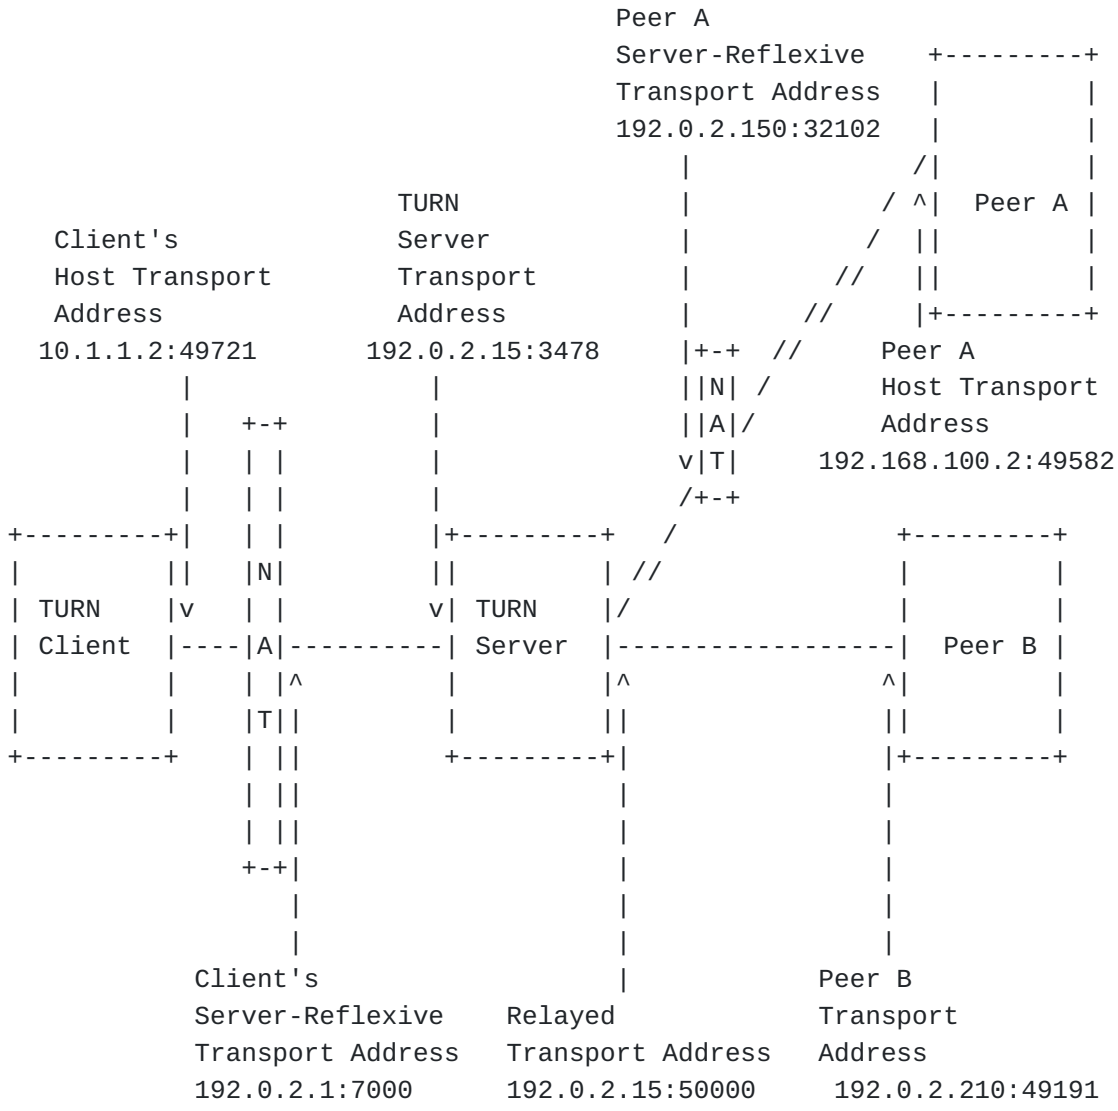

Figure 1

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 talks to the server from a (IP address, port) combination called the client's HOST TRANSPORT ADDRESS. (The combination of an IP address and port is called a 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

TRANSPORT ADDRESS. The client learns the server's transport address through some unspecified means (e.g., configuration), and this address is typically used by many clients simultaneously.

Since the client is behind a NAT, the server sees packets from the client 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 create and manipulate an ALLOCATION on the server. An allocation is a data structure on the server, an important component of which is a RELAYED TRANSPORT ADDRESS. The relayed transport address for the allocation is a transport address on the server which is used to send and receive packets to the peers.

Once an allocation is created, the client can send application data to the server along with an indication of which peer the data is to be sent to, and the server will relay this data to the appropriate peer. The client sends the application data to the server inside a TURN message; at the server, the data is extracted from the TURN message and sent to the peer in a UDP datagram. In the reverse direction, a peer can send application data in a UDP datagram to the relayed transport address for the allocation; the server will then encapsulate this data inside a TURN message and send it to the client along with an indication of which peer sent the data. Since the TURN message always contains an indication of which peer the client is communicating with, the client can use a single allocation to communicate with multiple peers.

When the peer is behind a NAT, then the client must identify the peer using its server-reflexive transport address rather than its host transport address. For example, to send application data to peer A in the example above, the client must specify 192.0.2.150:32102 (peer A's server-reflexive transport address) rather than 192.168.100.2:49582 (peer A's host transport address).

Each allocation on the server belongs to a single client and has exactly one relayed transport address which is used only by that allocation. Thus when a packet arrives at a relayed transport address on the server, the server knows which client the data is intended for. However, the client may have multiple allocations on a server at the same time.

2.1. Transports

TOC

TURN as defined in this specification always uses 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.

Since this version of TURN only supports UDP between the server and the peer, it is expected that most clients will prefer to also use UDP between the client and the server. That being the case, some readers may wonder: Why also support TCP and TLS over TCP?

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.

TURN as defined in this specification only supports IPv4. All IP addresses in this specification must be IPv4 addresses. However, there is a planned extension to TURN to add support for IPv6 and for relaying between IPv4 and IPv6 [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 some applications for TURN, the client may send and receive packets other than TURN packets on the host transport address it uses to communicate with the server. This can happen, for example, when using

TURN with ICE. In these cases, the client can distinguish TURN packets from other packets by examining the source address of the arriving packet: those arriving from the TURN server will be TURN packets.

2.2. Allocations

TOC

To create an allocation on the server, the client uses an Allocate transaction. The client sends a Allocate request to the server, and the server replies with an Allocate success 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). 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. 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 default lifetime of an allocation is 10 minutes -- this value was chosen to be long enough so that refreshing is not typically a burden on the client, while expiring allocations where the client has unexpectedly quit in a timely manner. However, the client can request a longer 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. Once a client no longer wishes to use an Allocation, it should delete the allocation using a Refresh request with a requested lifetime of 0.

Both the server and client keep track of a value known as the 5-TUPLE. At the client, the 5-tuple consists of the client's host transport address, the server transport address, and the transport protocol used by the client to communicate with the server. At the server, the 5-tuple value is the same except that the client's host transport address is replaced by the client's server-reflexive address, since that is the client's address as seen by the server.

Both the client and the server remember the 5-tuple used in the Allocate request. Subsequent messages between the client and the server uses the same 5-tuple. In this way, the client and server know which allocation is being referred to. If the client wishes to allocate a second relayed transport address, it must create a second allocation using a different 5-tuple (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, an implementation may use a file-descriptor in place of a 5-tuple to represent a TCP connection

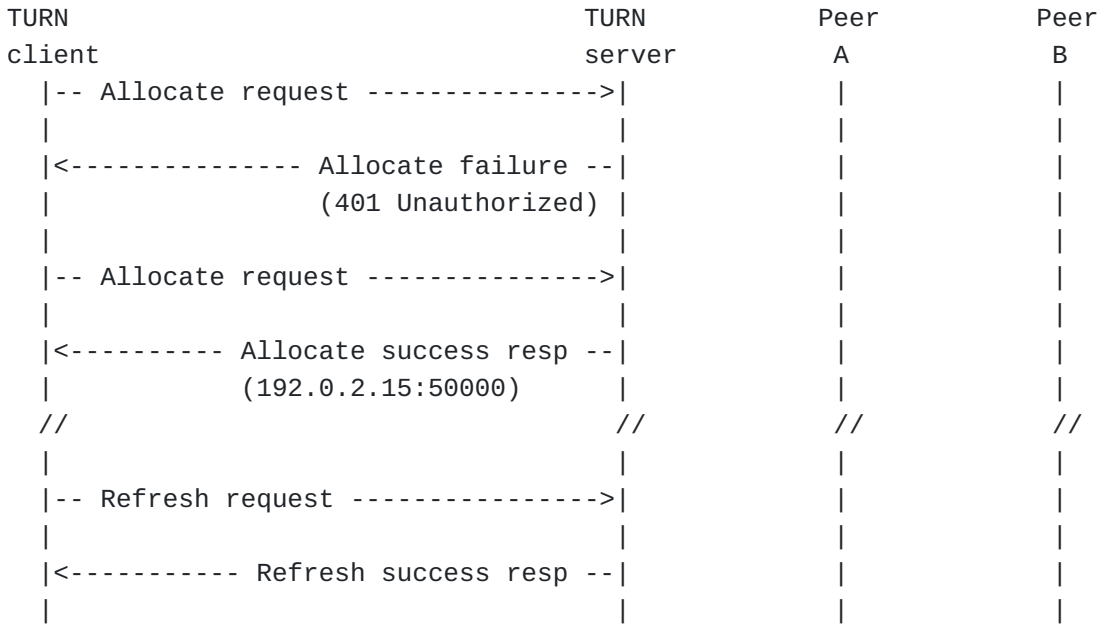

Figure 2

In Figure 2, the client sends an Allocate request to the server without credentials. Since the server requires that all requests be authenticated using STUN's long-term credential mechanism, the server rejects the request with a 401 (Unauthorized) error code. The client then tries again, this time including credentials (not shown). This time, the server accepts the Allocate request and returns an Allocate success response containing (amongst other things) the relayed transport address assigned to the allocation. Sometime later the client decides to refresh the allocation and thus sends a Refresh request to the server. The refresh is accepted and the server replies with a Refresh success response.

2.3. 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.).

An allocation can have zero or more permissions. Each permission consists of an IP address and a lifetime. When the server receives a UDP datagram on the allocation's relayed transport address, it first checks the list of permissions. If the source IP address of the datagram matches a permission, the application data is relayed to the client, otherwise the UDP datagram is silently discarded.

A permission expires after 5 minutes if it is not refreshed, and there is no way to explicitly delete a permission. This behavior was selected to match the behavior of a NAT that complies with [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.).

The client can install or refresh a permission using either a CreatePermission request or a ChannelBind request. Using the CreatePermission request, multiple permissions can be installed or refreshed with a single request -- this is important for applications that use ICE. For security reasons, permissions can only be installed or refreshed by transactions that can be authenticated; thus Send indications and ChannelData messages (which are used to send data to peers) do not install or refresh any permissions.

Note that permissions are within the context of an allocation, so adding or expiring a permission in one allocation does not affect other allocations.

2.4. Send Mechanism

TOC

There are two mechanisms for the client and peers to exchange application data using the TURN server. The first mechanism uses the Send and Data methods, 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.

The Send mechanism uses Send and Data indications. Send indications are used to send application data from the client to the server, while Data indications are used to send application data from the server to the client.

When using the Send mechanism, the client sends a Send indication to the TURN server containing (a) an XOR-PEER-ADDRESS attribute specifying the (server-reflexive) transport address of the peer and (b) a DATA attribute holding the application data. When the TURN server receives the Send indication, it extracts the application data from the DATA

attribute and sends it in a UDP datagram to the peer, using the allocated relay address as the source address. Note that there is no need to specify the relayed transport address, since it is implied by the 5-tuple used for the Send indication.

In the reverse direction, UDP datagrams arriving at the relayed transport address on the TURN server are converted into Data indications and sent to the client, with the server-reflexive transport address of the peer included in an XOR-PEER-ADDRESS attribute and the data itself in a DATA attribute. Since the relayed transport address uniquely identified the allocation, the server knows which client to relay the data to.

Send and Data indications cannot be authenticated, since the Long-Term Credential Mechanism of STUN does not support authenticating indications. This is not as big an issue as it might first appear, since the client-to-server leg is only half of the total path to the peer; applications that want proper security need to use encryption or similar to protect their data in the UDP datagrams between the server and the peer. However, to prevent attackers from injecting rogue Send indications to arbitrary destinations, TURN requires that a client install a permission to a peer before sending data to it using a Send indication.

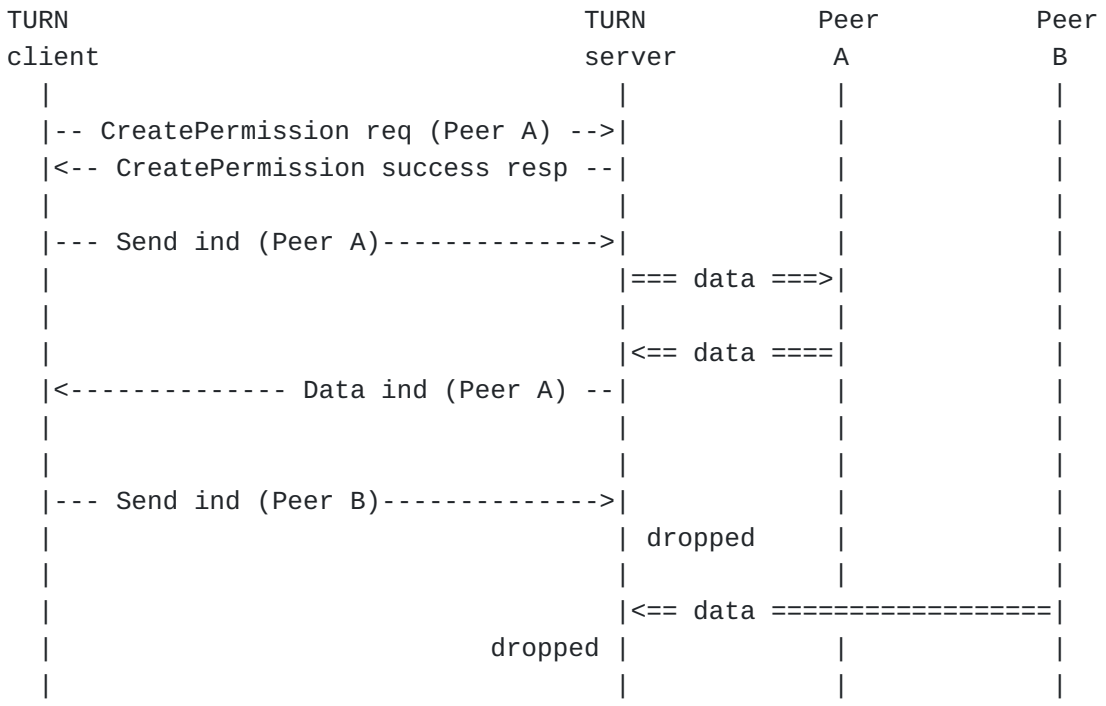

Figure 3

In Figure 3, the client has already created an allocation and now wishes to send data to its peers. The client first creates a permission by sending the server a CreatePermission request specifying peer A's (server reflexive) IP address in the XOR-PEER-ADDRESS attribute; if this was not done, the server would not relay data between the client and the server. The client then sends data to Peer A using a Send indication; at the server, the application 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. Later, the client attempts to exchange data with Peer B, however no permission has been installed for Peer B, so the Send indication from the client and the UDP datagram from the peer are both dropped by the server.

2.5. Channels

TOC

For some applications (e.g. Voice over IP), the 36 bytes of overhead that a Send indication 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 host transport 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 -- this lifetime was chosen to be longer than the permission lifetime. Channel bindings are refreshed by sending another ChannelBind request rebinding the channel to the peer. Like permissions (but unlike allocations), there is no way to explicitly delete a channel binding; the client must simply wait for it to time out.

Figure 4

Figure 4 shows the channel mechanism in use. The client has already created an allocation and now wishes to bind a channel to peer A. To do this, the client 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. When the ChannelData message arrives at the server, the server transfers the data to a UDP datagram and sends it to the peer A, as indicated by the channel number. When peer A sends a UDP datagram to the relayed transport address, the data is placed inside a ChannelData message and sent to the client.

Once a channel has been bound, the client is free to intermix ChannelData messages and Send indications. In the figure, the client later decides to use a Send indication rather than a ChannelData message to send additional data to peer A. The client might decide to do this, for example, so it can use the DONT-FRAGMENT attribute (see the next section). However, once a channel is bound, the server will always use a ChannelData message, as shown in the call flow.

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 would use the Send mechanism.

2.6. Unprivileged TURN Servers

TOC

This version of TURN is designed so that the server can be implemented as an application that runs in user space under commonly available operating systems without requiring special privileges. This design decision was taken to make it easy to deploy a TURN server: for example, to allow a TURN server to be integrated into a peer-to-peer application so that one peer can offer NAT traversal services to another peer.

This design decision has the following implications for data relayed by a TURN server:

- *The value of the Diff-Serv field may not be preserved across the server;
- *The TTL field may be reset, rather than decremented, across the server;
- *The ECN field may be reset by the server;
- *ICMP messages are not relayed by the server;
- *There is no end-to-end fragmentation, since the packet is re-assembled at the server.

Future work may specify alternate TURN semantics that address these limitations.

2.7. Avoiding IP Fragmentation

TOC

For reasons described in [Frag-Harmful] (Kent and Mogul, “Fragmentation Considered Harmful,” .), applications, especially those sending large volumes of data, should try hard to avoid having their packets fragmented. Applications using TCP can more-or-less ignore this issue because fragmentation avoidance is now a standard part of TCP, but applications using UDP (and thus any application using this version of TURN) must handle fragmentation avoidance themselves.

The application running on the client and the peer can take one of two approaches to avoid IP fragmentation.

The first approach is to avoid sending large amounts of application data in the TURN messages/UDP datagrams exchanged between the client and the peer. This is the approach taken by most VoIP (Voice-over-IP) applications. In this approach, the application exploits the fact that the IP specification [RFC0791] (Postel, J., “Internet Protocol,”

September 1981.) specifies that IP packets up to 576 bytes should never need to be fragmented.

The exact amount of application data that can be included while avoiding fragmentation depends the details of the TURN session between the client and the server: whether UDP, TCP, or TLS transport is used, whether ChannelData messages or Send/Data indications are used, and whether any additional attributes (such as the DONT-FRAGMENT attribute) are included. Another factor, which is hard to determine, is whether the MTU is somewhere along the path is reduced for other reasons, such as the use of IP-in-IP tunneling.

As a guideline, sending a maximum of 500 bytes of application data in a single TURN message (by the client on the client-to-server leg) or a UDP datagram (by the peer on the peer-to-server leg) will generally avoid IP fragmentation. To further reduce the chance of fragmentation, it is recommended that the client use ChannelData messages when transferring significant volumes of data, since the overhead of the ChannelData message is less than Send and Data indications.

The second approach the client and peer can take to avoid fragmentation is to use a path MTU discovery algorithm to determine the maximum amount of application data than can be sent without fragmentation. Unfortunately, because servers implementing this version of TURN do not relay ICMP messages, the classic Path MTU Discovery algorithm defined in [RFC1191] (Mogul, J. and S. Deering, “Path MTU discovery,” November 1990.) is not able to discover the MTU of the transmission path between the client and the peer. (Even if they did relay ICMP messages, the algorithm would not always work since ICMP messages are often filtered out by combined NAT/firewall devices).

So the client and server need to use a path MTU discovery algorithm that does not require ICMP messages. The Packetized Path MTU Discovery algorithm defined in [RFC4821] (Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” March 2007.) is one such algorithm.

The details of how to use the algorithm of [RFC4821] (Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” March 2007.) with TURN are still under investigation. However, as a step towards this goal, this version of TURN supports a DONT-FRAGMENT attribute. When the client includes this attribute in a Send indication, this tells the server to set the DF bit in the resulting UDP datagram that it sends to the peer. Since some servers may be unable to set the DF bit, the client should also include this attribute in the Allocate request -- any server that does not support the DONT-FRAGMENT attribute will indicate this by rejecting the Allocate request.

TOC

2.8. RTP Support

One of the envisioned uses of TURN is as a relay for clients and peers wishing to exchange real-time data (e.g. voice or video) using RTP. To facilitate the use of TURN for this purpose, TURN includes some special support for older versions of RTP.
Old versions of RTP [RFC3550] (Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, “RTP: A Transport Protocol for Real-Time Applications,” July 2003.) required that the RTP stream be on an even port number and the associated RTCP stream, if present, be on the next highest port. To allow clients to work with peers that still require this, TURN allows the client to request that the server allocate a relayed-transport-address with an even port number, and to optionally request the server reserve the next-highest port number for a subsequent allocation.

2.9. Anycast Discovery of Servers

TOC

This version of TURN has been designed to permit the future specification of a method of doing anycast discovery of a TURN server over UDP.
Specifically, a TURN server can reject an Allocate request with the suggestion that the server try an alternate server. To avoid certain types of attacks, the client must use the same credentials with the alternate server as it would have with the initial server.

---

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 familiar with [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.) and the terms defined there.
The following terms are used in this document:

> **TURN:** The protocol spoken between a TURN client and a TURN server. It is an extension to the STUN protocol [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 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.

Transport Address: The combination of an IP address and a port.

Host Transport Address: A transport address on a client or a peer.

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 on the TURN server that is used for relaying packets between the client and a peer. A peer sends to this address on the TURN server, and the packet is then relayed to the client.

TURN Server Transport Address: A transport address on the TURN server that is used for sending TURN messages to the server. This is the transport address that the client uses to communicate with the server.

Peer Transport Address: The transport address of the peer as seen by the server. When the peer is behind a NAT, this is the peer's server-reflexive transport address.

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 (currently one of UDP, TCP, or TLS)) 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.

Channel: A channel number and associated peer transport address. Once a channel number is bound to a peer's transport address, the client and server can use the more bandwidth-efficient ChannelData message to exchange data.

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.

Realm A string used to describe the server or a context within the server. The realm tells the client which username and password combination to use to authenticate requests.

Nonce A string chosen at random by the server and included in the message-digest. To prevent reply attacks, the server should change the nonce regularly.

---

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 [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 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 [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.).

In addition, the server SHOULD demand that all requests from the client be authenticated, using the Long-Term Credential mechanism described in [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.), and the client MUST be prepared to authenticate requests if required. Note that this authentication mechanism applies only to requests and cannot be used to authenticate indications, thus indications in TURN are never authenticated. If the server requires requests to be authenticated, then 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). For each allocation, the server SHOULD generate a new random nonce when the allocation is first attempted following the randomness recommendations in [RFC4086] (Eastlake, D., Schiller, J., and S. Crocker, “Randomness Requirements for Security,”

June 2005.) and SHOULD expire the nonce at least once every hour during the lifetime of the allocation.

All requests after the initial Allocate must 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.

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 client SHOULD include the SOFTWARE attribute in all Allocate and Refresh requests and MAY include it in any other requests or indications. The server SHOULD include the SOFTWARE attribute in all Allocate and Refresh responses (either success or failure) and MAY include it in other responses or indications. The client and the server MAY include the FINGERPRINT attribute in any STUN-formatted messages defined in this document.

TURN does not use the backwards-compatibility mechanism described in [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.).

By default, TURN runs on the same ports as STUN: 3478 for TURN over UDP and TCP, and 5349 for TURN over TLS. However, TURN has its own set of SRV service names: "turn" for UDP and TCP, and "turns" for TLS. Either the SRV procedures or the ALTERNATE-SERVER procedures, both described in Section 6 (Creating an Allocation), can be used to run TURN on a different port.

TURN as defined in this specification only supports IPv4. The client's IP address, the server's IP address and all IP addresses appearing in a relayed-transport-address MUST be IPv4 addresses.

When UDP transport is used between the client and the server, the client will retransmit a request if it does not receive a response within a certain timeout period. Because of this, the server may receive two (or more) requests with the same 5-tuple and same transaction id. STUN requires that the server recognize this case and treat the request as idempotent (see [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.)). Some implementations may choose to meet this requirement by remembering all received requests and the corresponding responses for 40 seconds. Other implementations may choose to reprocess the request and arrange that such reprocessing returns essentially the

same response. To aid implementors who choose the latter approach (the so-called "stateless stack approach"), this specification includes some implementation notes on how this might be done. Implementations are free to choose either approach or choose some other approach that gives the same results.

When TCP transport is used between the client and the server, it is possible that a bit error will cause a length field in a TURN packet to become corrupted, causing the receiver to lose synchronization with the incoming stream of TURN messages. A client or server which detects a long sequence of invalid TURN messages over TCP transport SHOULD close the corresponding TCP connection to help the other end detect this situation more rapidly.

To mitigate either intentional or unintentional denial-of-service attacks against the server by clients with valid usernames and passwords, it is RECOMMENDED that the server impose limits on both the number of allocations active at one time for a given username and on the amount of bandwidth those allocations can use. The server should reject new allocations that would exceed the limit on the allowed number of allocations active at one time with a 486 (Allocation Quota Exceeded) (see Section 6.2 (Receiving an Allocate Request)), and should discard application data traffic that exceeds the bandwidth quota.

---

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's IP address, client's port, server IP address, server port, transport protocol)
- the authentication information
- the time-to-expiry
- A list of permissions
- A list of channel to peer bindings

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. On the client, the 5-tuple uses the client's host transport address, while on the server the 5-tuple uses the client's server-reflexive transport address.

Both the relayed-transport-address and the 5-tuple MUST be unique across all allocations, so either one can be used to uniquely identify the allocation.
The authentication information (e.g., username, password, realm, and nonce) are used to both verify subsequent requests and to compute the message integrity of responses. The username, realm, and nonce values are initially those used in the authenticated Allocate request that creates the allocation, though the server can change the nonce value during the lifetime of the allocation using a 438 (Stale Nonce) reply. Note that rather than storing the password explicitly, it may be desirable for security reasons for the server to store the key value which is an MD5 hash over the username, realm and password (see [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.)).
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 the default lifetime value of 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.
The list of permissions is described in Section 8 (Permissions) and the list of channels is described in Section 11 (Channels).

---

6. Creating an Allocation

TOC

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 picks a host transport address. It is RECOMMENDED that the client pick a currently-unused transport address, typically by allowing the underlying OS to pick a currently-unused port for a new socket.
The client then picks 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 also picks a server transport address, which SHOULD be done as follows. The client receives (perhaps through configuration) a domain name for a TURN server. The client then uses the DNS procedures described in [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 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.
If the client wishes the server to initialize the time-to-expiry 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 wishes to later use the DONT-FRAGMENT attribute in one or more Send indications on this allocation, then the client SHOULD include the DONT-FRAGMENT attribute in the Allocate request. This allows the client to test whether this attribute is supported by the server.
If the client requires the port number of the relayed-transport address be even, the client includes the EVEN-PORT attribute. If this attribute is not included, then the port can be even or odd. By setting the R bit in the EVEN-PORT attribute to 1, the client can request that the server reserve the next highest port number (on the same IP address) for a subsequent allocation. If the R bit is 0, no such request is made.
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 omit the EVEN-PORT attribute.
Once constructed, the client sends the Allocate request on the 5-tuple.

TOC

6.2. Receiving an Allocate Request

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

1. The server SHOULD require that the request be authenticated using the Long-Term Credential mechanism of [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)," October 2008.).

2. The server checks if the 5-tuple is currently in use by an existing allocation. If yes, the server rejects the request with a 437 (Allocation Mismatch) error.

3. The server checks if the request contain a REQUESTED-TRANSPORT 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 442 (Unsupported Transport Protocol) error.

4. The request may contain a DONT-FRAGMENT attribute. If it does, but the server does not support sending UDP datagrams with the DF bit set to 1 (see Section 12 (IP Header Fields)), then the server treats the DONT-FRAGMENT attribute in the Allocate request as an unknown comprehension-required attribute.

5. The server checks if the request contains a RESERVATION-TOKEN attribute. If yes, and the request also contains a EVEN-PORT attribute, then the server rejects 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. The server checks if the request contains an EVEN-PORT attribute. If yes, then the server checks that it can satisfy the request (i.e., can allocate a relayed-transport-address as described below). If the server cannot satisfy the request, then the server rejects the request with a 508 (Insufficient Port Capacity) error.

7. At any point, the server MAY 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.

8. Also at any point, the server MAY choose to reject the request with a 300 (Try Alternate) error if it wishes to redirect the client to a different server. The use of this error code and attribute follow the specification in [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)," October 2008.), with the modification that a TURN server MAY return this error code and attribute in unauthenticated error responses as well as in authenticated error responses.

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.
The server chooses a relayed-transport-address for the allocation as follows:

*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 that the reservation is a server-wide reservation and is not specific to a particular allocation, since the Allocate request containing the RESERVATION-TOKEN uses a different 5-tuple than the Allocate request that made the reservation. The 5-tuple for the Allocate request containing the RESERVATION-TOKEN attribute can be any allowed 5-tuple; it can use 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).

*If the request contains an EVEN-PORT attribute with the R bit set to 0, then the server allocates a relayed-transport-address with an even port number.

*If the request contains an EVEN-PORT attribute with the R bit set to 1, 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.

*Otherwise, the server allocates any available relayed-transport-address.

In all cases, 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 strongly 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 EVEN-PORT attribute.

The server determines the initial value of the time-to-expiry field as follows. If the request contains a LIFETIME attribute, then the server computes MIN(client's proposed lifetime, server's maximum allowed lifetime). If this computed lifetime is greater than the default lifetime, then the server uses that value. Otherwise, the server uses the default lifetime. It is RECOMMENDED that the server use a maximum allowed lifetime value 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 allowed 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-expiry 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 XOR-RELAYED-ADDRESS attribute containing the relayed transport address;

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

> NOTE: Implementations may implement the idempotency of the Allocate request over UDP using the so-called "stateless stack approach" as follows. To detect retransmissions when the original request was successful in creating an allocation, the server can store the transaction id that created the request with the allocation data and compare it with incoming Allocate requests on the same 5-tuple. Once such a request is detected, the server can stop parsing the request and immediately generate a success response. When building this response, the value of the LIFETIME attribute can be taken from the time-to-expiry field in the allocate state data, even though this value may differ slightly from the LIFETIME value originally returned. In addition, the server may need to store an indication of any reservation token returned in the original response, so that this may be returned in any retransmitted responses.
>
> For the case where the original request was unsuccessful in creating an allocation, the server may choose to do nothing special. Note, however, that there is a rare case where the server rejects the original request but accepts the retransmitted request (because conditions 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 will cause it to retry with a different 5-tuple. The server may use a smaller maximum lifetime value to minimize the lifetime of allocations "orphaned" in this manner.

6.3. Receiving an Allocate Success Response

TOC

If the client receives an Allocate success response, then it MUST check that the mapped address and the relayed transport address are in an

address family that the client understands and is prepared to deal with. This specification only covers the case where these two addresses are IPv4 addresses. If these two addresses are 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 IPv6, but the request would arrive at the server over IPv4, 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.

6.4. Receiving an Allocate Error Response

TOC

If the client receives an Allocate 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 considers the current transaction as having failed but MAY choose to retry the Allocate request using a different transport (e.g., TCP instead of UDP).

* 300 (Try Alternate): The server would like the client to use the server specified in the ALTERNATE-SERVER attribute instead. The client considers the current transaction as having failed, but SHOULD try the Allocate request with the alternate server before trying any other servers (e.g., other servers discovered using the SRV procedures). When trying the Allocate request with the alternate server, the client follows the ALTERNATE-SERVER procedures specified in [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT

(STUN)," October 2008.) with the following changes: the client SHOULD accept unauthenticated error responses containing the 300 (Try Alternate) error code, the client MUST ensure that the realm value received from the alternate server is as expected, the client MUST use the same transport protocol to the alternate server as it used to the original server, and the client MUST use the same username and password as it would have with the original server. The latter checks protect against an attacker sending the client an unauthenticated Allocate error response that redirects the client to some totally different and unexpected server.

*400 (Bad Request): The server believes the client's request is malformed for some reason. The client considers the current transaction as having failed. The client MAY notify the user or operator and SHOULD NOT retry the request with this server until it believes the problem has been fixed.

*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. In this case, the client considers the current transaction as having failed and SHOULD notify the user or operator. The client SHOULD NOT send any further requests to this server until it believes the problem has been fixed.

*403 (Forbidden): The request is valid, but the server is refusing to perform it, likely due to administrative restrictions. The client considers the current transaction as having failed. 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.

*420 (Unknown Attribute): If the client included a DONT-FRAGMENT attribute in the request and the server rejected the request with a 420 error code and listed the DONT-FRAGMENT attribute in the UNKNOWN-ATTRIBUTES attribute in the error response, then the client now knows that the server does not support the DONT-FRAGMENT attribute. The client considers the current transaction as having failed but MAY choose to retry the Allocate request without the DONT-FRAGMENT attribute.

*437 (Allocation Mismatch): This indicates that the client has picked a 5-tuple which the server sees as already in use. One way this could happen is if an intervening NAT assigned a mapped transport address that was used by another client which recently crashed. The client considers the current transaction as having failed. 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.

*438 (Stale Nonce): See the procedures for the Long-Term Credential mechanism [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)," October 2008.).

*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.

*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 reattempt the request with this server until it believes the problem has been fixed.

*486 (Allocation Quota Reached): The server is currently unable to create any more allocations with this username. The client considers the current transaction as having failed. The client SHOULD wait at least 1 minute before trying to create any more allocations on the server.

*508 (Insufficient Port Capacity): The server has no more relayed transport addresses available, or has none with the requested properties, or the one that was reserved is no longer available. The client considers the current operation as having failed. If the client is using either the EVEN-PORT 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.

---

7. Refreshing an Allocation

TOC

A Refresh transaction can be used to either (a) refresh an existing allocation and update its time-to-expiry, 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 refresh 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-expiry timer of an allocation. If the client wishes the server to set the time-to-expiry 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-expiry 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.

7.2. Receiving a Refresh Request

TOC

When the server receives a Refresh request, it processes as per Section 4 (General Behavior) plus the specific rules mentioned here.
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, then the server computes MIN(client's requested lifetime, server's maximum allowed lifetime). If this computed value is greater than the default lifetime, then the "desired lifetime" is the computed value. Otherwise the "desired lifetime" is the default lifetime.
Subsequent processing depends on the desired lifetime value:

- If desired lifetime is 0, then the request succeeds and the allocation is deleted.

- If 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-expiry timer.

> NOTE: A server need not do anything special to implement idempotency of Refresh requests over UDP using the "stateless stack approach". Retransmitted Refresh requests with a non-zero desired lifetime will simply refresh the allocation. A retransmitted Refresh request with a zero desired lifetime will cause a 437 (Allocation Mismatch) response if the allocation has already been deleted, but the client will treat this as equivalent to a success response (see below).

7.3. Receiving a Refresh Response

TOC

If the client receives a success response to its Refresh request with a non-zero lifetime, it updates its copy of the allocation data structure with the time-to-expiry value contained in the response.
If the client receives a 437 (Allocation Mismatch) error response to a request to delete the allocation, then the allocation no longer exists and it should consider its request as having effectively succeeded.

---

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 set of peers that are allowed to send data to the client, and the time-to-expiry is the number of seconds until the permission expires.
By sending either CreatePermission requests or ChannelBind requests, the client can cause the server to install or refresh a permission for a given IP address. 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 Permission Lifetime.

- If a permission for that IP address already exists, then the time-to-expiry for that permission is reset to Permission Lifetime.

The Permission Lifetime MUST be 300 seconds (= 5 minutes).
Each permission’s time-to-expiry decreases down once per second until it reaches 0, at which point the permission expires and is deleted.
CreatePermission and ChannelBind requests may be freely intermixed on a permission. A given permission may be installed or refreshed at one

point in time with a CreatePermission request, and then refreshed with a ChannelBind request at a different point in time, or vice-versa.
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. CreatePermission

TOC

TURN supports two ways for the client to install or refresh permissions on the server. This section describes one way: the CreatePermission request.
A CreatePermission request may be used in conjunction with either the Send mechanism in Section 10 (Send and Data Methods) or the Channel mechanism in Section 11 (Channels).

9.1. Forming a CreatePermission request

TOC

The client who wishes to install or refresh one or more permissions can send a CreatePermission request to the server.
When forming a CreatePermission request, the client MUST include at least one XOR-PEER-ADDRESS attribute, and MAY include more than one such attribute. The IP address portion of each XOR-PEER-ADDRESS attribute contains the IP address for which a permission should be installed or refreshed. The port portion of each XOR-PEER-ADDRESS attribute will be ignored and can be any arbitrary value. The various XOR-PEER-ADDRESS attributes can appear in any order.

9.2. Receiving a CreatePermission request

TOC

When the server receives the CreatePermission request, it processes as per Section 4 (General Behavior) plus the specific rules mentioned here.
The message is checked for validity. The CreatePermission request MUST contain at least XOR-PEER-ADDRESS attribute and MAY contain multiple such attributes. If no such attribute exists, or if any of these attributes are invalid, then a 400 (Bad Request) error is returned. If the request is valid, but the server is unable to satisfy the request due to some capacity limit or similar, then a 508 (Insufficient Capacity) error is returned.
The server MAY impose restrictions on the IP address and port values allowed in the XOR-PEER-ADDRESS attribute -- if a value is not allowed, the server rejects the request with a 403 (Forbidden) error.
If the message is valid and the server is capable of carrying out the request, then the server installs or refreshes a permission for the IP address contained in each XOR-PEER-ADDRESS attribute as described in Section 8 (Permissions). The port portion of each attribute is ignored and may be any arbitrary value.
The server then responds with a CreatePermission success response. There are no mandatory attributes in the success response.

> NOTE: A server need not do anything special to implement idempotency of CreatePermission requests over UDP using the "stateless stack approach". Retransmitted CreatePermission requests will simply refresh the permissions.

9.3. Receiving a CreatePermission response

TOC

If the client receives a valid CreatePermission success response, then the client updates its data structures to indicate that the permissions have been installed or refreshed.

10. Send and Data Methods

TOC

TURN supports two mechanisms for sending and receiving data from peers. This section describes the use of the Send and Data mechanism, while Section 11 (Channels) describes the use of the Channel mechanism.

10.1. Forming a Send Indication

TOC

The client can use a Send indication to pass data to the server for relaying to a peer. A client may use a Send indication even if a channel is bound to that peer. However the client MUST ensure that there is a permission installed for the IP address of the peer to which the Send indication is being sent; this prevents a third party from using a TURN server to send data to arbitrary destinations.
When forming a Send indication, the client MUST include a XOR-PEER-ADDRESS attribute and a DATA attribute. The XOR-PEER-ADDRESS attribute contains the transport address of the peer to which the data is to be sent, and the DATA attribute contains the actual application data to be sent to the peer.
The client MAY include a DONT-FRAGMENT attribute in the Send indication if it wishes the server to set the DF bit on the UDP datagram sent to the peer.

10.2. Receiving a Send Indication

TOC

When the server receives a Send indication, it processes as per Section 4 (General Behavior) plus the specific rules mentioned here.
The message is first checked for validity. The Send indication MUST contain both a XOR-PEER-ADDRESS attribute and a DATA attribute. If one of these attributes is missing or invalid, then the message is discarded. Note that the DATA attribute is allowed to contain zero bytes of data.
The Send indication may also contain the DONT-FRAGMENT attribute. If the server is unable to set the DF bit on outgoing UDP datagrams when this attribute is present, then the server acts as if the DONT-FRAGMENT attribute is an unknown comprehension-required attribute (and thus the Send indication is discarded).
The server also checks that there is a permission installed for the IP address contained in the XOR-PEER-ADDRESS attribute. If no such permission exists, the message is discarded. Note that a Send indication never causes the server to refresh the permission.
The server MAY impose restrictions on the IP address and port values allowed in the XOR-PEER-ADDRESS attribute -- if a value is not allowed, the server silently discards the Send indication.
If everything is OK, then the server 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 XOR-PEER-ADDRESS attribute;

*the data following the UDP header is the contents of the value field of the DATA attribute.

The handling of the DONT-FRAGMENT attribute (if present), is described in Section 12 (IP Header Fields).
The resulting UDP datagram is then sent to the peer.

10.3. 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 11 (Channels)). If a channel is bound, then processing proceeds as described in Section 11.7 (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 XOR-PEER-ADDRESS and a DATA attribute. The DATA attribute is set to the value of the ‘data octets’ field from the datagram, and the XOR-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.

10.4. Receiving a Data Indication

TOC

When the client receives a Data indication, it checks that the Data indication contains both a XOR-PEER-ADDRESS and a DATA attribute, and discards the indication if it does not. The client SHOULD also check that the XOR-PEER-ADDRESS attribute value contains an IP address with which the client believes there is an active permission, and discard the Data indication otherwise. Note that the DATA attribute is allowed to contain zero bytes of data.

> NOTE: The latter check protects the client against an attacker who somehow manages to trick the server into installing permissions not desired by the client.

If the Data indication passes the above checks, 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 XOR-PEER-ADDRESS attribute.

---

11. 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.
The ChannelData message (see Section 11.4 (The ChannelData Message)) starts with a two-byte field that carries the channel number. The values of this field are allocated as follows:

> 0x0000 through 0x3FFF: These values can never be used for channel numbers.
>
> 0x4000 through 0x7FFF: These values are the allowed channel numbers (16,383 possible values)
>
> 0x8000 through 0xFFFF: These values are reserved for future use.

Because of this division, ChannelData messages can be distinguished from STUN-formatted messages (e.g., Allocate request, Send indication, etc) by examining the first two bits of the message:

> 0b00: STUN-formatted message (since the first two bits of a STUN-formatted message are always zero)
>
> 0b01: ChannelData message (since the channel number is the first field in the ChannelData message and channel numbers fall in the range 0x4000 - 0x7FFF)
>
> 0b10: Reserved
>
> 0b11: Reserved

The reserved values may be used in the future to extend the range of channel numbers. Thus an implementation MUST NOT assume that a TURN message always starts with a 0 bit.
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 to 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 the use of a channel number or peer transport address in a channel binding in one

allocation has no impact on their use in a different 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 peer's 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 a ChannelBind request rebinding the channel to the same peer) resets the time-to-expiry timer back to 10 minutes.
When the channel binding expires, the channel becomes unbound. Once unbound, the channel number can be bound to a different transport address, and the transport address can be bound to a different channel number. To prevent race conditions, the client MUST wait 5 minutes after the channel binding expires before attempting to bind the channel number to a different transport address or the transport address to a different channel number.
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.

11.1. Sending a ChannelBind Request

TOC

A channel binding is created or refreshed using a ChannelBind transaction. A ChannelBind transaction also creates or refreshes a permission towards the peer (see Section 8 (Permissions)).
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 XOR-PEER-ADDRESS attribute. Section 11.2 (Receiving a ChannelBind Request) describes the restrictions on these attributes.
Rebinding a channel to the same transport address that it is already bound to provides a way to refresh a channel binding and the corresponding permission without sending data to the peer. Note however, that permissions need to be refreshed more frequently than channels.

11.2. Receiving a ChannelBind Request

TOC

When the server receives a ChannelBind request, it processes as per Section 4 (General Behavior) plus the specific rules mentioned here.
The server checks the following:

- The request contains both a CHANNEL-NUMBER and a XOR-PEER-ADDRESS attribute;
- The channel number is in the range 0x4000 through 0x7FFE (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 a 400 (Bad Request) error.
The server MAY impose restrictions on the IP address and port values allowed in the XOR-PEER-ADDRESS attribute -- if a value is not allowed, the server rejects the request with a 403 (Forbidden) error.
If the request is valid, but the server is unable to fulfill the request due to some capacity limit or similar, the server replies with a 508 (Insufficient Capacity) error.
Otherwise, the server replies with a ChannelBind success response.
There are no required attributes in a successful ChannelBind response.
If the server can satisfy the request, then the server creates or refreshes the channel binding using the channel number in the CHANNEL-NUMBER attribute and the transport address in the XOR-PEER-ADDRESS attribute. The server also installs or refreshes a permission for the IP address in the XOR-PEER-ADDRESS attribute as described in Section 8 (Permissions).

> NOTE: A server need not do anything special to implement idempotency of ChannelBind requests over UDP using the "stateless stack approach". Retransmitted ChannelBind requests will simply refresh

> the channel binding and the corresponding permission. Furthermore, the client must wait 5 minutes before binding a previously bound channel number or peer address to a different channel, eliminating the possibility that the transaction would initially fail but succeed on a retransmission.

11.3. Receiving a ChannelBind Response

TOC

When the client receives a ChannelBind success response, it updates its data structures to record that the channel binding is now active. It also updates its data structures to record that the corresponding permission has been installed or refreshed.
If the client receives a ChannelBind failure response that indicates that the channel information is out-of-sync between the client and the server (e.g., an unexpected 400 "Bad Request" response), then it is RECOMMENDED that the client immediately delete the allocation and start afresh with a new allocation.

11.4. The ChannelData Message

TOC

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

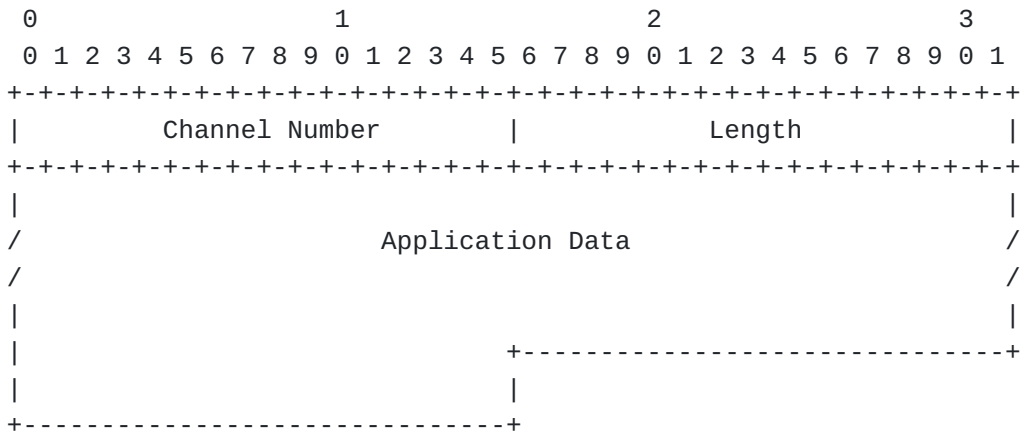

The Channel Number field specifies the number of the channel on which the data is traveling, and thus the address of the peer that is sending or is to receive the data.
The Length field specifies the length in bytes of the application data field (i.e., it does not include the size of the ChannelData 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.

11.5. 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 11.4 (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.

11.6. Receiving a ChannelData Message

TOC

The receiver of the ChannelData message uses the first two bits to distinguish it from STUN-formatted messages, as described above. If the message uses a value in the reserved range (0x8000 through 0xFFFF), then the message is silently discarded.
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 11.5 (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.
On the client, it is RECOMMENDED that the client discard the ChannelData message if the client believes there is no active permission towards the peer. On the server, the receipt of a ChannelData message MUST NOT refresh either the channel binding or the permission towards the peer.

On the server, 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.

11.7. 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 10.3 (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 11.5 (Sending a ChannelData Message).

---

12. IP Header Fields

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. The descriptions in this section apply: (a) when the server sends a UDP datagram to the peer, or (b) when the server sends a Data indication or ChannelData message to the client over UDP transport. The descriptions in this section do not apply to TURN messages sent over TCP or TLS transport from the server to the client.

The descriptions below have two parts: a preferred behavior and an alternate behavior. The server SHOULD implement the preferred behavior, but if that is not possible for a particular field, then it SHOULD implement the alternative behavior.

Time to Live (TTL) field

> Preferred Behavior: If the incoming value is 0, then the drop the incoming packet. 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 (DSCP) field [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.)

> 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.

Explicit Congestion Notification (ECN) field [RFC3168] (Ramakrishnan, K., Floyd, S., and D. Black, “The Addition of Explicit Congestion Notification (ECN) to IP,” September 2001.)

> Preferred Behavior: Set the outgoing value to the incoming value, UNLESS the server is doing Active Queue Management, the incoming ECN field is ECT(1) (=0b01) or ECT(0) (=0b10), and the server wishes to indicate that congestion has been experienced, in which case set the outgoing value to CE (=0b11).
>
> Alternate Behavior: Set the outgoing value to Not-ECT (=0b00).

IPv4 Fragmentation fields

> Preferred Behavior:
>
> > When the server sends a packet to a peer in response to a Send indication containing the DONT-FRAGMENT attribute, then set the DF bit in the outgoing IP header to 1. In all other cases when sending an outgoing packet containing application data (e.g.,

> Data indication, ChannelData message, or DONT-FRAGMENT attribute not included in the Send indication), copy the DF bit from the DF bit of the incoming packet that contained the application data.
>
> Set the other fragmentation fields (Identification, MF, Fragment Offset) as appropriate for a packet originating from the server.

Alternate Behavior: As described in the Preferred Behavior, except always assume the incoming DF bit is 0.

In both the Preferred and Alternate Behaviors, the resulting packet may be too large for the outgoing link. If this is the case, then the normal fragmentation rules apply [RFC1122] (Braden, R., "Requirements for Internet Hosts - Communication Layers," October 1989.).

IPv4 Options

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

---

13. 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.

```
      0x003  :  Allocate          (only request/response semantics defined)
      0x004  :  Refresh           (only request/response semantics defined)
      0x006  :  Send              (only indication semantics defined)
      0x007  :  Data              (only indication semantics defined)
      0x008  :  CreatePermission  (only request/response semantics defined
      0x009  :  ChannelBind       (only request/response semantics defined)
```

---

14. New STUN Attributes

TOC

This STUN extension defines the following new attributes:

```
0x000C: CHANNEL-NUMBER
0x000D: LIFETIME
0x0010: Reserved (was BANDWIDTH)
0x0012: XOR-PEER-ADDRESS
0x0013: DATA
0x0016: XOR-RELAYED-ADDRESS
0x0018: EVEN-PORT
0x0019: REQUESTED-TRANSPORT
0x001A: DONT-FRAGMENT
0x0021: Reserved (was TIMER-VAL)
0x0022: RESERVATION-TOKEN
```

14.1. 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.

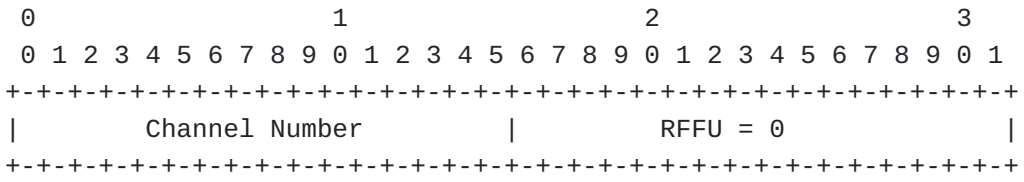

14.2. LIFETIME

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.

14.3. XOR-PEER-ADDRESS

TOC

The XOR-PEER-ADDRESS specifies the address and port of the peer as seen from the TURN server. (In other words, the peer's server-reflexive transport address if the peer is behind a NAT). It is encoded in the same way as XOR-MAPPED-ADDRESS [RFC5389] (Rosenberg, J., Mahy, R.,

Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.).

14.4. DATA

TOC

The DATA attribute is present in all Send and Data 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).

14.5. XOR-RELAYED-ADDRESS

TOC

The XOR-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 [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.).

14.6. EVEN-PORT

TOC

This attribute allows the client to request that the port in the relayed-transport-address be even, and (optionally) that the server reserve the next-higher port number. The attribute is 8 bits long. Its format is:

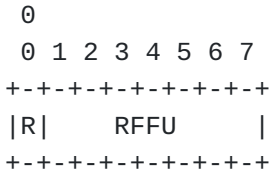

The attribute contains a single 1-bit flag:

> **R:** If 1, the server is requested to reserve the next higher port number (on the same IP address) for a subsequent allocation. If 0, no such reservation is requested.

The other 7 bits of the attribute must be set to zero on transmission and ignored on reception.

14.7. 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:

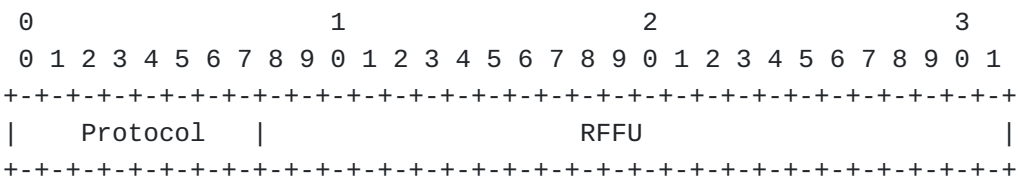

The Protocol field specifies the desired protocol. The codepoints used in this field are taken from those allowed in the Protocol field in the IPv4 header and the NextHeader field in the IPv6 header [Protocol-Numbers] (, “IANA Protocol Numbers Registry,” 2005.). This specification only allows the use 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.

14.8. DONT-FRAGMENT

TOC

This attribute is used by the client to request that the server set the DF (Don't Fragment) bit in the IP header when relaying the application data onward to the peer. This attribute has no value part and thus the attribute length field is 0.

14.9. 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.

---

15. New STUN Error Response Codes

TOC

This document defines the following new error response codes:

403 (Forbidden): The request was valid, but cannot be performed due to administrative or similar restrictions.

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 Capacity): The server is unable to carry out the request due to some capacity limit being reached. In an Allocate response, this could be due to the server having no more relayed transport addresses available right now, or having none with the requested properties, or the one that corresponds to the specified reservation token is not available.

---

16. Detailed Example

TOC

This section gives a example of the use of TURN, showing in detail the contents of the messages exchanged. The example uses the network diagram shown in the Overview (Figure 1).
For each message, the attributes included in the message and their values are shown. For convenience, values are shown in a human-readable format rather than showing the actual octets; for example "XOR-RELAYED-ADDRESS=192.0.2.15:9000" shows that the XOR-RELAYED-ADDRESS attribute is included with an address of 192.0.2.15 and a port of 9000, here the address and port are shown before the xor-ing is done. For attributes with string-like values (e.g. SOFTWARE="Example client, version 1.03" and NONCE="adl7W7PeDU4hKE72jdaQvbAMcr6h39sm"), the value of the attribute is shown in quotes for readability, but these quotes do not appear in the actual value.

The client begins by selecting a host transport address to use for the TURN session; in this example the client has selected 10.1.1.2:49721 as shown in Figure 1. The client then sends an Allocate request to the server at the server transport address. The client randomly selects a 96-bit transaction id of 0xA56250D3F17ABE679422DE85 for this transaction; this is encoded in the transaction id field in the fixed header. The client includes a SOFTWARE attribute that gives information about the client's software; here the value is "Example client, version 1.03" to indicate that this is version 1.03 of something called the Example client. The client includes the LIFETIME attribute because it wishes the allocation to have a longer lifetime than the default of 10 minutes; the value of this attribute is 3600 seconds, which corresponds

to 1 hour. The client must always include a REQUESTED-TRANSPORT attribute in an Allocate request and the only value allowed by this specification is 17, which indicates UDP transport between the server and the peers. The client also includes the DONT-FRAGMENT attribute because it wishes to use the DONT-FRAGMENT attribute later in Send indications; this attribute consists of only an attribute header, there is no value part. We assume the client has not recently interacted with the server, thus the client does not include USERNAME, REALM, NONCE, or MESSAGE-INTEGRITY attribute. Finally, note that the order of attributes in a message is arbitrary (except for the MESSAGE-INTEGRITY and FINGERPRINT attributes) and the client could have used a different order.

The server follows the recommended practice in this specification of requiring all requests to be authenticated. Thus when the server receives the initial Allocate request, it rejects the request because the request does not contain the authentication attributes. Following the procedures of the Long-Term Credential Mechanism of STUN [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT (STUN),” October 2008.), the server includes an ERROR-CODE attribute with a value of 401 (Unauthorized), a REALM attribute that specifies the authentication realm used by the server (in this case, the server's domain "example.com"), and a nonce value in a NONCE attribute. The server also includes a SOFTWARE attribute that gives information about the server's software.

The client, upon receipt of the 401 error, re-attempts the Allocate request, this time including the authentication attributes. The client selects a new transaction id, and then populates the new Allocate request with the same attributes as before. The client includes a USERNAME attribute and uses the realm value received from the server to help it determine which value to use; here the client is configured to use the username "George" for the realm "example.com". The client also includes the REALM and NONCE attributes, which are just copied from the 401 error response. Finally, the client includes a MESSAGE-INTEGRITY attribute as the last attribute in the message, whose value is an HMAC-SHA1 hash over the contents of the message (shown as just "..." above); this HMAC-SHA1 computation also covers a password value, thus an attacker cannot compute the message integrity value without somehow knowing the secret password.

The server, upon receipt of the authenticated Allocate request, checks that everything is OK, then creates an allocation. The server replies with an Allocate success response. The server includes a LIFETIME attribute giving the lifetime of the allocation; here, the server as reduced the client's requested 1 hour lifetime to just 20 minutes, because this particular server doesn't allow lifetimes longer than 20 minutes. The server includes an XOR-RELAYED-ADDRESS attribute whose value is the relayed transport address of the allocation. The server includes an XOR-MAPPED-ADDRESS attribute whose value is the server-reflexive address of the client; this value is not used otherwise in TURN but is returned as a convenience to the client. The server

includes a MESSAGE-INTEGRITY attribute to authenticate the response and to insure its integrity; note that the response does not contain the USERNAME, REALM, and NONCE attributes. The server also includes a SOFTWARE attribute.

The client then creates a permission towards peer A in preparation for sending it some application data. This is done through a CreatePermission request. The XOR-PEER-ADDRESS attribute contains the IP address for which a permission is established (the IP address of peer A); note that the port number in the attribute is ignored when used in a CreatePermission request, and here it has been set to 0; also note how the client uses Peer A's server-reflexive IP address and not its (private) host address. The client uses the same username, realm, and nonce values as in the previous request on the allocation. Though it is allowed to do so, the client has chosen not to include a SOFTWARE attribute in this request.

The server receives the CreatePermission request, creates the corresponding permission, and then replies with a CreatePermission success response. Like the client, the server chooses not to include the SOFTWARE attribute in its reply. Again, note how success responses contain a MESSAGE-INTEGRITY attribute (assuming the server uses the Long-Term Credential Mechanism), but no USERNAME, REALM, and NONCE attributes.

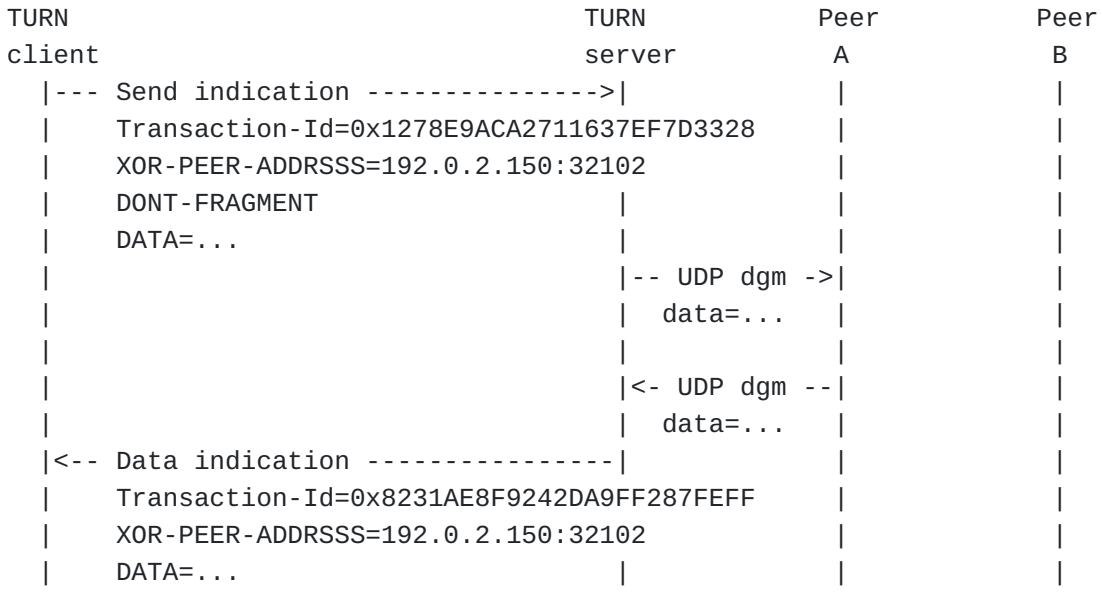

The client now sends application data to Peer A using a Send indication. Peer A's server-reflexive transport address is specified in the XOR-PEER-ADDRESS attribute, and the application data (shown here as just "...") is specified in the DATA attribute. The client is doing a form of path MTU discovery at the application layer and thus specifies (by including the DONT-FRAGMENT attribute) that the server should set the DF bit in the UDP datagram send to the peer. Indications cannot be authenticated using the Long-Term Credential Mechanism of STUN, so no MESSAGE-INTEGRITY attribute is included in the message. An application wishing to ensure that its data is not altered or forged must integrity-protect its data at the application level.

Upon receipt of the Send indication, the server extracts the application data and sends it in a UDP datagram to Peer A, with the relayed-transport-address as the source transport address of the datagram, and with the DF bit set as requested. Note that, had the client not previously established a permission for Peer A's server-reflexive IP address, then the server would have silently discarded the Send indication instead.

Peer A then replies with its own UDP datagram containing application data. The datagram is sent to the relayed-transport-address on the server. When this arrives, the server creates a Data indication containing the source of the UDP datagram in the XOR-PEER-ADDRESS attribute, and the data from the UDP datagram in the DATA attribute. The resulting Data indication is then sent to the client.

The client now binds a channel to Peer B, specifying a free channel number (0x4000) in the CHANNEL-NUMBER attribute, and Peer B's transport address in the XOR-PEER-ADDRESS attribute. As before, the client re-uses the username, realm, and nonce from its last request in the message.
Upon receipt of the request, the server binds the channel number to the peer, installs a permission for Peer B's IP address, and then replies with ChannelBind success response.

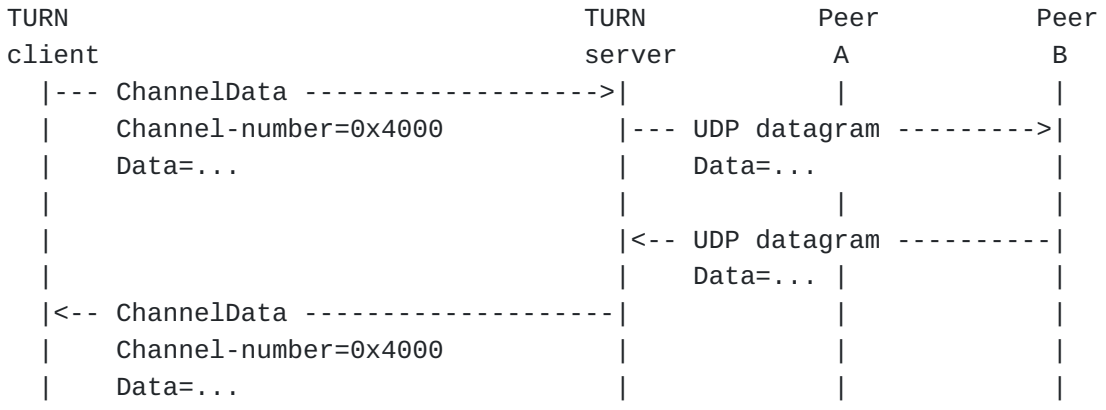

The client now sends a ChannelData message to the server with data destined for Peer B. The ChannelData message is not a STUN message, and thus has no transaction id. Instead, its fixed header has only two fields: channel number and data; here the channel number field is 0x4000 (the channel the client just bound to Peer B). When the server receives the ChannelData message, it checks that the channel is currently bound (which it is) and then sends the data onward to Peer B in a UDP datagram, using the relayed-transport-address as the source transport address and 192.0.2.210:49191 (the value of the XOR-PEER-ADDRESS attribute in the ChannelBind request) as the destination transport address.

Later, Peer B sends a UDP datagram back to the relayed-transport-address. This causes the server to send a ChannelData message to the client containing the data from the UDP datagram. The server knows which client to send the ChannelData message to because of the relayed-transport-address the UDP datagram arrived at, and knows to use channel 0x4000 because this is the channel bound to 192.0.2.210:49191. Note that if there had not been any channel number bound to that address, the server would have used a Data indication instead.

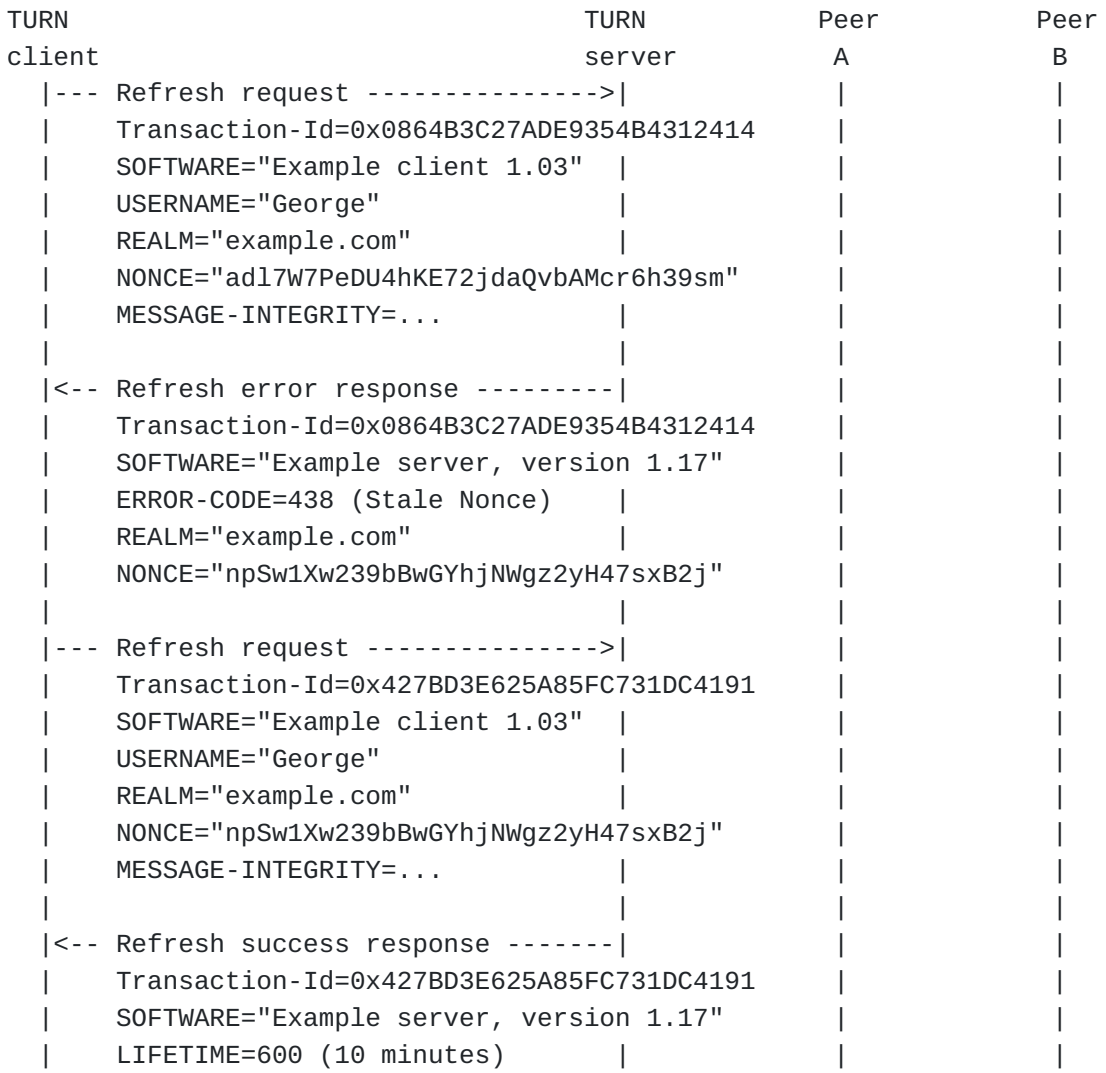

Sometime before the 20 minute lifetime is up, the client refreshes the allocation. This is done using a Refresh request. As before, the client includes the latest username, realm, and nonce values in the request. The client also includes the SOFTWARE attribute, following the recommended practice of always including this attribute in Allocate and Refresh messages. When the server receives the Refresh request, it notices that the nonce value has expired, and so replies with 438 (Stale Nonce) error given a new nonce value. The client then reattempts the request, this time with the new nonce value. This second attempt is accepted, and the server replies with a success response. Note that the

client did not include a LIFETIME attribute in the request, so the server refreshes the allocation for the default lifetime of 10 minutes (as can be seen by the LIFETIME attribute in the success response).

---

17. Security Considerations

TOC

This section considers attacks that are possible in a TURN deployment, and discusses how they are mitigated by mechanisms in the protocol or recommended practices in the implementation.

17.1. Outsider Attacks

TOC

Outsider attacks are ones where the attacker has no credentials in the system, and is attempting to disrupt the service seen by the client or the server.

17.1.1. Obtaining Unauthorized Allocations

TOC

An attacker might wish to obtain allocations on a TURN server for any number of nefarious purposes. A TURN server provides a mechanism for sending and receiving packets while cloaking the actual IP address of the client. This makes TURN servers an attractive target for attackers who wish to use it to mask their true identity.
An attacker might also wish to simply utilize the services of a TURN server without paying for them. Since TURN services require resources from the provider, it is anticipated that their usage will come with a cost.
These attacks are prevented using the digest authentication mechanism which allows the TURN server to determine the identity of the requestor and whether the requestor is allowed to obtain the allocation.

17.1.2. Offline Dictionary Attacks

TOC

The digest authentication mechanism used by TURN is subject to offline dictionary attacks. An attacker that is capable of eavesdropping on a message exchange between a client and server can determine the password by trying a number of candidate passwords and seeing if one of them is correct. This attack works when the passwords are low entropy, such as a word from the dictionary. This attack can be mitigated by using

strong passwords with large entropy. In situations where even stronger mitigation is required, TLS transport between the client and the server can be used.

17.1.3. Faked Refreshes and Permissions

TOC

An attacker might wish to attack an active allocation by sending it a Refresh request with an immediate expiration, in order to delete it and disrupt service to the client. This is prevented by authentication of refreshes. Similarly, an attacker wishing to send CreatePermission requests to create permissions to undesirable destinations is prevented from doing so through authentication. The motivations for such an attack are described in Section 17.2 (Firewall Considerations).

17.1.4. Fake Data

TOC

An attacker might wish to send data to the client or the peer, as if they came from the peer or client respectively. To do that, the attacker can send the client a faked Data Indication or ChannelData message, or send the TURN server a faked Send Indication or ChannelData message.

Indeed, since indications and ChannelData messages are not authenticated, this attack is not prevented by TURN. However, this attack is generally present in IP-based communications and is not substantially worsened by TURN. Consider an normal, non-TURN IP session between hosts A and B. An attacker can send packets to B as if they came from A by sending packets towards A with a spoofed IP address of B. This attack requires the attacker to know the IP addresses of A and B. With TURN, an attacker wishing to send packets towards a client using a Data indication needs to know its IP address (and port), the IP address and port of the TURN server, and the IP address and port of the peer (for inclusion in the XOR-PEER-ADDRESS attribute). To send a fake ChannelData message to a client, an attacker needs to know the IP address and port of the client, the IP address and port of the TURN server, and the channel number. This particular combination is mildly more guessable than in the non-TURN case.

These attacks are more properly mitigated by application layer authentication techniques. In the case of real time traffic, usage of SRTP [RFC3711] (Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, “The Secure Real-time Transport Protocol (SRTP),” March 2004.) prevents these attacks.

In some situations, the TURN server may be situated in the network such that it is able to send to hosts that the client cannot directly send to. This can happen, for example, if the server is located behind a

firewall that allows packets from outside the firewall to be delivered to the server, but not to other hosts behind the firewall. In these situations, an attacker could send the server a Send indication with an XOR-PEER-ADDRESS attribute containing the transport address of one of the other hosts behind the firewall. If the server was to allow relaying of traffic to arbitrary peers, then this would provide a way for the attacker to attack arbitrary hosts behind the firewall.

To mitigate this attack, TURN requires that the client establish a permission to a host before sending it data. Thus an attacker can only attack hosts that the client is already communicating with, unless the attacker is able to create authenticated requests. Furthermore, the server administrator may configure the server to restrict the range of IP addresses and ports that it will relay data to. To provide even greater security, the server administrator can require that the client use TLS for all communication between the client and the server.

17.1.5. Impersonating a Server

TOC

When a client learns a relayed address from a TURN server, it uses that relayed address in application protocols to receive traffic. Therefore, an attacker wishing to intercept or redirect that traffic might try to impersonate a TURN server and provide the client with a faked relayed address.

This attack is prevented through the digest authentication mechanism, which provides message integrity for responses in addition to verifying that they came from the server. Furthermore, an attacker cannot replay old server responses as the transaction ID in the STUN header prevents this. Replay attacks are further thwarted through frequent changes to the nonce value.

17.1.6. Eavesdropping Traffic

TOC

TURN concerns itself primarily with authentication and message integrity. Confidentiality is only a secondary concern, as TURN control messages do not include information that is particularly sensitive. The primary protocol content of the messages is the IP address of the peer. If it is important to prevent an eavesdropper on a TURN connection from learning this, TURN can be run over TLS.

Confidentiality for the application data relayed by TURN is best provided by the application protocol itself, since running TURN over TLS does not protect application data between the server and the peer. If confidentiality of application data is important, then the application should encrypt or otherwise protect its data. For example, for real time media, confidentiality can be provided by using SRTP.

17.1.7. TURN loop attack

TOC

An attacker might attempt to cause data packets to loop indefinitely between two TURN servers. The attack goes as follows. First, the attacker sends an Allocate request to server A, using the source address of server B. Server A will send its response to server B, and for the attack to succeed, the attacker must have the ability to either view or guess the contents of this response, so that the attacker can learn the allocated relayed-transport-address. The attacker then sends an Allocate request to server B, using the source address of server A. Again, the attacker must be able to view or guess the contents of the response, so it can send learn the allocated relayed-transport-address. Using the same spoofed source address technique, the attacker then binds a channel number on server A to the relayed-transport-address on server B, and similarly binds the same channel number on server B to the relayed-transport-address on server A. Finally, the attacker sends a ChannelData message to server A.

The result is a data packet that loops from the relayed-transport-address on server A to the relayed-transport-address on server B, then from server B's transport address to server A's transport address, and then around the loop again.

This attack is mitigated as follows. By requiring all requests to be authenticated and/or by randomizing the port number allocated for the relayed-transport-address, the server forces the attacker to either intercept or view responses sent to a third party (in this case, the other server) so that the attacker can authenticate the requests and learn the relayed-transport-address. Without one of these two measures, an attacker can guess the contents of the responses without needing to see them, which makes the attack much easier to perform. Furthermore, by requiring authenticated requests, the server forces the attacker to have credentials acceptable to the server, which turns this from an outsider attack into an insider attack and allows the attack to be traced back to the client initiating it.

The attack can be further mitigated by imposing a per-username limit on the bandwidth used to relay data by allocations owned by that username, to limit the impact of this attack on other allocations. More mitigation can be achieved by decrementing the TTL when relaying data packets (if the underlying OS allows this).

17.2. Firewall Considerations

TOC

A key aspect of TURN's security considerations is that it should not weaken the protections afforded by firewalls deployed between a client

and a TURN server. It is anticipated that TURN servers will often be present on the public Internet, and clients may often be inside enterprise networks with corporate firewalls. If TURN servers provide a 'backdoor' for reaching into the enterprise, TURN will be blocked by these firewalls.
TURN servers therefore emulate the behavior of NAT devices which implement address-dependent filtering [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.), a property common in many firewalls as well. When a NAT or firewall implements this behavior, packets from an outside IP address are only allowed to be sent to an internal IP address and port if the internal IP address and port had recently sent a packet to that outside IP address. TURN servers introduce the concept of permissions, which provide exactly this same behavior on the TURN server. An attacker cannot send a packet to a TURN server and expect it to be relayed towards the client, unless the client has tried to contact the attacker first.
It is important to note that some firewalls have policies which are even more restrictive than address-dependent filtering. Firewalls can also be configured with address and port dependent filtering, or can be configured to disallow inbound traffic entirely. In these cases, if a client is allowed to connect the TURN server, communications to the client will be less restrictive than what the firewall would normally allow.

17.2.1. Faked Permissions

TOC

In firewalls and NAT devices, permissions are granted implicitly through the traversal of a packet from the inside of the network towards the outside peer. Thus, a permission cannot, by definition, be created by any entity except one inside the firewall or NAT. With TURN, this restriction no longer holds. Since the TURN server sits outside the firewall, at attacker outside the firewall can now send a message to the TURN server and try to create a permission for itself.
This attack is prevented because all messages which create permissions (i.e., ChannelBind and CreatePermission) are authenticated.

17.2.2. Blacklisted IP Addresses

TOC

Many firewalls can be configured with blacklists which prevent a client behind the firewall from sending packets to, or receiving packets from, ranges of blacklisted IP addresses. This is accomplished by inspecting the source and destination addresses of packets entering and exiting the firewall, respectively.

If a client connects to a TURN server, it will be able to bypass such blacklisting policies and communicate with IP addresses which the firewall would otherwise restrict. This is a problem for other protocols that provide tunneling functions, such as VPNs. It is possible to build TURN-aware firewalls which inspect TURN messages, and check the IP address of the correspondent. TURN messages to offending destinations can then be rejected. TURN is designed so that this inspection can be done statelessly.

17.2.3. Running Servers on Well-Known Ports

TOC

A malicious client behind a firewall might try to connect to a TURN server and obtain an allocation which it then uses to run a server. For example, a client might try to run a DNS server or FTP server.
This is not possible in TURN. A TURN server will never accept traffic from a peer for which the client has not installed a permission. Thus, peers cannot just connect to the allocated port in order to obtain the service.

17.3. Insider Attacks

TOC

In insider attacks, a client has legitimate credentials but defies the trust relationship that goes with those credentials. These attacks cannot be prevented by cryptographic means but need to be considered in the design of the protocol.

17.3.1. DoS Against TURN Server

TOC

A client wishing to disrupt service to other clients might obtain an allocation and then flood it with traffic, in an attempt to swamp the server and prevent it from servicing other legitimate clients. This is mitigated by the recommendation that the server limit the amount of bandwidth it will relay for a given username. This won't prevent a client from sending a large amount of traffic, but it allows the server to immediately discard traffic in excess.
Since each allocation uses a port number on the IP address of the TURN server, the number of allocations on a server is finite. An attacker might attempt to consume all of them by requesting a large number of allocations. This is prevented by the recommendation that the server impose a limit of the number of allocations active at a time for a given username.

17.3.2. Anonymous Relaying of Malicious Traffic

TOC

TURN servers provide a degree of anonymization. A client can send data to correspondent peers without revealing their own IP addresses. TURN servers may therefore become attractive vehicles for attackers to launch attacks against targets without fear of detection. Indeed, it is possible for a client to chain together multiple TURN servers, such that any number of relays can be used before a target receives a packet.

Administrators who are worried about this attack can maintain logs which capture the actual source IP and port of the client, and perhaps even every permission that client installs. This will allow for forensic tracing to determine the original source, should it be discovered that an attack is being relayed through a TURN server.

17.3.3. Manipulating other Allocations

TOC

An attacker might attempt to disrupt service to other users of the TURN server by sending Refresh requests or CreatePermission requests which (through source address spoofing) appear to be coming from another user of the TURN server. TURN prevents this by requiring that the credentials used in CreatePermission, Refresh, and ChannelBind messages match those used to create the initial allocation. Thus, the fake requests from the attacker will be rejected.

17.4. Other Considerations

TOC

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

---

18. IANA Considerations

TOC

Since TURN is an extension to STUN [RFC5389] (Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, “Session Traversal Utilities for NAT

(STUN)," October 2008.), the methods, attributes and error codes defined in this specification are new methods, attributes, and error codes for STUN. This section requests 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 13 (New STUN Methods).
The codepoints for the new STUN attributes defined in this specification are listed in Section 14 (New STUN Attributes).
The codepoints for the new STUN error codes defined in this specification are listed in Section 15 (New STUN Error Response Codes).
IANA is requested to allocate the SRV service name of "turn" for TURN over UDP or TCP, and the service name of "turns" for TURN over TLS.
IANA is requested to create a registry for TURN channel numbers, initially populated as follows:

> 0x0000 through 0x3FFF: Not available for use, since they conflict with the STUN header.
>
> 0x4000 through 0x7FFF: A TURN implementation is free to use channel numbers in this range.
>
> 0x8000 through 0xFFFF: Reserved.

Any change to this registry must be made through an IETF Standards Action.

---

19. IAB Considerations

TOC

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. These considerations and the responses for TURN are documented in this section.
Consideration 1: Precise definition of a specific, limited-scope problem that is to be solved with the UNSAF proposal. A short term fix should not be generalized to solve other problems. Such generalizations lead to the prolonged dependence on and usage of the supposed short term fix -- meaning that it is no longer accurate to call it "short term".
Response: TURN is a protocol for communication between a relay (= TURN server) and its client. The protocol allows a client that is behind a NAT to obtain and use a public IP address on the relay. As a

convenience to the client, TURN also allows the client to determine its server-reflexive transport address.
Consideration 2: Description of an exit strategy/transition plan. The better short term fixes are the ones that will naturally see less and less use as the appropriate technology is deployed.
Response: TURN will no longer be needed once there are no longer any NATs. Unfortunately, as of the date of publication of this document, it no longer seems very likely that NATs will go away any time soon. However, the need for TURN will also decrease as the number of NATs with the mapping property of Endpoint-Independent Mapping [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.) increases.
Consideration 3: Discussion of specific issues that may render systems more "brittle". For example, approaches that involve using data at multiple network layers create more dependencies, increase debugging challenges, and make it harder to transition.
Response: TURN is "brittle" in that it requires the NAT bindings between the client and the server to be maintained unchanged for the lifetime of the allocation. This is typically done using keep-alives. If this is not done, then the client will lose its allocation and can no longer exchange data with its peers.
Consideration 4: Identify requirements for longer term, sound technical solutions; contribute to the process of finding the right longer term solution.
Response: The need for TURN will be reduced once NATs implement the recommendations for NAT UDP behavior documented in [RFC4787] (Audet, F. and C. Jennings, “Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” January 2007.). Applications are also strongly urged to use 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.) to communicate with peers; though ICE uses TURN, it does so only as a last resort, and uses it in a controlled manner.
Consideration 5: Discussion of the impact of the noted practical issues with existing deployed NATs and experience reports.
Response: Some NATs deployed today exhibit a mapping behavior other than Endpoint-Independent mapping. These NATs are difficult to work with, as they make it difficult or impossible for protocols like ICE to use server-reflexive transport addresses on those NATs. A client behind such a NAT is often forced to use a relay protocol like TURN because "UDP hole punching" techniques [RFC5128] (Srisuresh, P., Ford, B., and D. Kegel, “State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs),” March 2008.) do not work.

TOC

20. Open Issues

Note to RFC Editor: Please remove this section prior to publication of this document as an RFC.
This section lists the known issues in this version of the specification.
(No known issues at this time).

---

21. 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 technical and major editorial changes between the various versions of this specification. Minor editorial changes are not described.

21.1. Changed from -13 to -14

TOC

*Reworded the text in Section 6.2 (Receiving an Allocate Request) and Section 7.2 (Receiving a Refresh Request) to more clearly describe how the allocation lifetime is computed in the case where a client requests a lifetime that is greater than both the default lifetime and the server's maximum allowed lifetime.

*In Section 8 (Permissions), changed the term "default permission lifetime" to "Permission Lifetime" to make it clearer that the lifetime of a permission is not configurable.

*In Section 6.2 (Receiving an Allocate Request), swapped the steps that check the RESERVATION-TOKEN and EVEN-PORT attributes to correctly handle the case where an Allocate request contains both attributes. The new text correctly returns 400 Bad Request in this case.

*Added text in the Overview section to describe why various timer values were chosen.

*Added a sentence to the IAB consideration section saying that the disappearance of NATs in the near-term seems unlikely.

*The former "Other Features" section of the Overview has been replaced with a series of sections describing various secondary features of TURN, and the text describing and motivating these secondary features has been expanded. As a part of this rewrite,

there is now a section that describes how to avoid IP fragmentation when using TURN.

*Added some additional text in the Overview to explain how a client would select between UDP, TCP, and TLS transport.

*Fixed various minor typos.

21.2. Changes from -12 to -13

TOC

*Added a new error code: 403 (Forbidden).

*When processing a CreatePermission or ChannelBind request containing a XOR-PEER-ADDRESS attribute, the server is allow to reject certain IP address and port combinations for administrative or other reasons by returning a 403 (Forbidden) error.

*Added a request to IANA to establish a registery for channel numbers.

*Clarified the usage of the nonce value: a new random nonce SHOULD be selected for each Allocate attempt, and the nonce SHOULD be expired at least once an hour. Referenced [RFC4086] (Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security," June 2005.) for guidelines on selecting the nonce value.

*Made a number of minor editoral changes.

21.3. Changes from -11 to -12

TOC

*Changed the port numbers used in the examples for the client, the peers, and the relayed-transport-address to put them in the Dynamic port range. They were previously in the Registered port range, which was arguably unrealistic.

*Noted that the XOR-MAPPED-ADDRESS attribute is defined in RFC 5389.

*Used the codepoint names (Not-ECT, ECT(0), ECT(1), and CE) when talking about the ECN field.

*Updated the Introduction to note that the client must not only communicate its relayed-transport-address to the peers, but also learn the peers' server-reflexive transport addresses. As a result, removed the suggestion that the client could use a webpage to communicate with its peers.

*Added a description of the "TURN Loop attack" and its mitigation to the Security Considerations section.

*Fixed some errors in the examples in the Overview section. They had not been updated to be consistent with the change introduced in version -11 that a permission must be created before a client can send data to a peer.

*In the Additional Features subsection of the Overview, reworded the discussion of what end-to-end features are preserved by TURN. The previous text said that a number of features did not work, but as of version -11, these features _may_ work. At the same time, added a sentence noting that any Path MTU Discovery mechanism using the DONT-FRAGMENT attribute will not receive ICMP messages and will thus have to use techniques like those described in [RFC4821] (Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” March 2007.).

*Added the recommendation that, when TCP transport is used between the client and the server, both ends should close the connection if they notice a long sequence of invalid TURN messages. A likely cause of this is an undetected bit error corrupting a length field somewhere.

*Reworded the paragraph explaining that channel bindings are per-allocation to further stress this point.

*In the discussion on setting the fragmentation fields, added a sentence saying that the client or server should follow the normal rules for fragmentation as described in [RFC1122] (Braden, R., “Requirements for Internet Hosts - Communication Layers,” October 1989.).

21.4. Changes from -10 to -11

TOC

*Clarified that, when the client is redirected to an alternate server, the client uses the same transport protocol to the alternate server as it did to the original server.

*Clarified the information that the server needs to store to authenticate requests and to compute the message-integrity on responses. Noted that the server need not store the password explicitly, but can instead store the key value, which may be desirable for security reasons.

*Clarified that TURN runs on the same ports as TURN by default, but noted that a server can use a different port because TURN has its own SRV service names. Strengthened the language for using the SRV procedures from "typically" to "SHOULD". Also added a sentence in the IANA considerations section requesting that IANA reserve the service names for TURN; previously they were described in the text but not mentioned in the IANA considerations section.

*Added a detailed example, complete with attributes and their values, of the use of TURN.

*Reduced the range of channel numbers. Channel numbers now range from 0x4000 through 0x7FFF. Values in the range 0x8000 through 0xFFFF are now reserved.

*Rewrote the IAB Considerations section to directly address the considerations listed in [RFC3424] (Daigle, L. and IAB, “IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation,” November 2002.).

*Generalized the 508 error code so it can be used for any sort of capacity-related problem. This error code was previously allowed only in Allocate responses, but is now also allowed in CreatePermission and ChannelBind responses to indicate that the server is unable to carry out the request due to some capacity problem.

*Changed the syntax of the CreatePermission request to allow multiple XOR-PEER-ADDRESS attributes to appear in the message, so that multiple permissions can be created or refreshed at the same time.

*Added the restriction that the server must already have a permission installed for the IP address in the XOR-PEER-ADDRESS attribute of a Send indication, otherwise the Send indication is ignored by the server.

*Put back the preferred behaviors into Section 12 (IP Header Fields), reversing the change made in version -10.

*Explicitly allow the server to restrict the range of IP addresses and ports it is willing to relay data too.

21.5. Changes from -09 to -10

TOC

*Changed the recommendation for using the SOFTWARE attribute. Previously its use was recommended in all requests and responses; now it is only recommended in Allocate and Refresh requests and responses, though it may appear elsewhere. Also, version -09 incorrectly referred to this attribute as "SOFTWARE-TYPE".

*Changed the name of the PEER-ADDRESS and RELAYED-ADDRESS attributes to XOR-PEER-ADDRESS and XOR-RELAYED-ADDRESS respectively for consistency with other specifications.

*Removed the concept of a "preserving" allocation. All allocations are now non-preserving. This simplifies the base specification and allows it to advance more rapidly; see the discussion in the BEHAVE meeting of 29 July 2008. The concept of a preserving allocation will be advanced as an extension to TURN. As part of this change, the P bit in the REQUESTED-PROPS attribute, the ICMP attribute, and ICMP message relaying was removed. Further, in Section 12 (IP Header Fields), the preferred behaviors were removed, leaving the alternate behaviors as the specified behaviors.

*Replaced the REQUESTED-PROPS attribute with the EVEN-PORT attribute. The new attribute lacks the feature of the old attribute of being an alternate way to specify new allocation properties. As a consequence, the only way to specify a new allocation property is to define a new attribute.

*Added text recommending that the client check that the IP address in XOR-PEER-ADDRESS attribute in a received Data indication is one with which the client believes there is an active permission. Similarly, it is recommended that the client check that a permission exist when receiving a ChannelData message.

*Added text recommending that the client delete the allocation if it receives a ChannelBind failure response on an unbound channel.

*Added the CreatePermission request/response transaction which adds or updates permissions, and removed the ability for Send indications and ChannelBind messages to install or update permissions. The net effect is that only authenticate-able messages (i.e., CreatePermission requests and ChannelBind requests) can install or refresh permissions; unauthenticate-able Send indications and ChannelData messages do not.

*Removed all support for IPv6. All IPv6 support, including ways of relaying between IPv4 and IPv6, will now be covered in [I-D.ietf-behave-turn-ipv6] (Camarillo, G., Novo, O., and S. Perreault, “Traversal Using Relays around NAT (TURN) Extension for IPv6,” March 2010.).

*Reserved attribute code point 0x0021. This was previously used for the TIMER-VAL attribute, which was removed when the SetActiveDestination feature was removed.

*Added the DONT-FRAGMENT attribute which allows the client to request that the server set the DF bit when sending the UDP datagram to the peer. This attribute may appear in both Allocate requests and Send indications.

*Changed how the ALTERNATE-SERVER attribute is used. The attribute can no longer be used with any error code, but must be used with 300 (Try Alternative). It can now appear in unauthenticated responses, however there are restrictions around how the subsequent Allocate request is authenticated.

*Reworked the details of how idempotency of requests is handled, making it clear that the stack can either remember all transactions for 40 seconds, or can handle this using the so-called "stateless stack approach". Made some changes to the semantics of the Allocate, Refresh, and ChannelBind requests as a consequence.

*Added the requirement that a client cannot re-use previously bound channel number or transport address until 5 minutes after the channel binding expires. This avoids various race conditions.

*Removed the requirement that an allocation cannot be re-used within 2 minutes of having been deleted. This requirement was put in place to prevent mis-delivered packets but is no longer seen as having any real value.

*Added a recommendation that the server impose quotas on both the number of allocations and the amount of bandwidth a given username can use at one time. These quotas help protect against denial-of-service attacks.

*Completely rewrote the security considerations section.

*Made quite a few changes to the descriptive text in both the Overview and the normative text to try to further clarify concepts.

21.6. 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.

21.7. 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.

21.8. 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".

21.9. 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.

21.10. 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.

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, Jason Fischl, Derek MacDonald, Scott Godin, Cullen Jennings, Lars Eggert, Magnus Westerlund, Benny Prijono, and Eric Rescorla have been particularly helpful, with Eric also suggesting the channel allocation mechanism, and Cullen suggesting the REQUESTED-PORT-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

[RFC5389]	Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)," RFC 5389, October 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).
[RFC3168]	Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP," RFC 3168, September 2001 (TXT).
[RFC1122]	Braden, R., "Requirements for Internet Hosts - Communication Layers," STD 3, RFC 1122, October 1989 (TXT).

23.2. Informative References

TOC

[RFC1191]	Mogul, J. and S. Deering, "Path MTU discovery," RFC 1191, November 1990 (TXT).
[RFC0791]	Postel, J., "Internet Protocol," STD 5, RFC 791, September 1981 (TXT).
[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets," BCP 5, RFC 1918, February 1996 (TXT).
[RFC3424]	Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation," RFC 3424, November 2002 (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-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).
[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-port-randomization]	Larsen, M. and F. Gont, "Transport Protocol Port Randomization Recommendations," draft-ietf-tsvwg-port-randomization-07 (work in progress), April 2010 (TXT).
[RFC5128]	Srisuresh, P., Ford, B., and D. Kegel, "State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)," RFC 5128, March 2008 (TXT).
[RFC1928]	Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5," RFC 1928, March 1996 (TXT).
[RFC3550]	Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications," STD 64, RFC 3550, July 2003 (TXT, PS, PDF).
[RFC3711]	

	Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, “The Secure Real-time Transport Protocol (SRTP),” RFC 3711, March 2004 (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).
[RFC4821]	Mathis, M. and J. Heffner, “Packetization Layer Path MTU Discovery,” RFC 4821, March 2007 (TXT).
[RFC3261]	Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, “SIP: Session Initiation Protocol,” RFC 3261, June 2002 (TXT).
[I-D.rosenberg-mmusic-ice-nonsip]	Rosenberg, J., “Guidelines for Usage of Interactive Connectivity Establishment (ICE) by non Session Initiation Protocol (SIP) Protocols,” draft-rosenberg-mmusic-ice-nonsip-01 (work in progress), July 2008 (TXT).
[RFC4086]	Eastlake, D., Schiller, J., and S. Crocker, “Randomness Requirements for Security,” BCP 106, RFC 4086, June 2005 (TXT).
[Frag-Harmful]	Kent and Mogul, “Fragmentation Considered Harmful.” Proc. SIGCOMM ‘87, vol. 17, No. 5, October 1987
[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
	(Unaffiliated)
Email:	rohan@ekabal.com
	Philip Matthews
	Alcatel-Lucent
	600 March Road
	Ottawa, Ontario

	Canada
Phone:	
Fax:	
Email:	philip_matthews@magma.ca
URI:	