

BEHAVE Working Group
Internet-Draft
Obsoletes: [3489](#) (if approved)
Intended status: Standards Track
Expires: January 9, 2008

J. Rosenberg
Cisco
C. Huitema
Microsoft
R. Mahy
Plantronics
P. Matthews
Avaya
D. Wing
Cisco
July 8, 2007

Session Traversal Utilities for (NAT) (STUN)
draft-ietf-behave-rfc3489bis-07

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 9, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Session Traversal Utilities for NAT (STUN) is a protocol that serves

Internet-Draft

STUN

July 2007

as a tool for other protocols in dealing with NAT traversal. It can be used by an endpoint to determine the IP address and port allocated to it by a NAT. It can also be used to check connectivity between two endpoints, and as a keep-alive protocol to maintain NAT bindings. STUN works with many existing NATs, and does not require any special behavior from them.

STUN is not a NAT traversal solution by itself. Rather, it is a tool to be used in the context of a NAT traversal solution. This is an important change from the previous version of this specification ([RFC 3489](#)), which presented STUN as a complete solution.

This document obsoletes [RFC 3489](#).

Table of Contents

1.	Introduction	4
2.	Evolution from RFC 3489	4
3.	Overview of Operation	5
4.	Terminology	8
5.	Definitions	9
6.	STUN Message Structure	10
7.	Base Protocol Procedures	12
7.1.	Forming a Request or an Indication	12
7.2.	Sending the Request or Indication	13
7.2.1.	Sending over UDP	13
7.2.2.	Sending over TCP or TLS-over-TCP	14
7.3.	Receiving a STUN Message	15
7.3.1.	Processing a Request	16
7.3.1.1.	Forming a Success or Error Response	17
7.3.1.2.	Sending the Success or Error Response	17
7.3.2.	Processing an Indication	18
7.3.3.	Processing a Success Response	18
7.3.4.	Processing an Error Response	18
8.	FINGERPRINT Mechanism	19
9.	DNS Discovery of a Server	19
10.	Authentication and Message-Integrity Mechanisms	20
10.1.	Short-Term Credential Mechanism	21
10.1.1.	Forming a Request or Indication	21
10.1.2.	Receiving a Request or Indication	21
10.1.3.	Receiving a Response	22
10.2.	Long-term Credential Mechanism	23

10.2.1.	Forming a Request	24
10.2.1.1.	First Request	24
10.2.1.2.	Subsequent Requests	24
10.2.2.	Receiving a Request	24
10.2.3.	Receiving a Response	25

11.	ALTERNATE-SERVER Mechanism	26
12.	Backwards Compatibility with RFC 3489	26
12.1.	Changes to Client Processing	27
12.2.	Changes to Server Processing	27
13.	STUN Usages	28
14.	STUN Attributes	29
14.1.	MAPPED-ADDRESS	30
14.2.	XOR-MAPPED-ADDRESS	31
14.3.	USERNAME	32
14.4.	MESSAGE-INTEGRITY	32
14.5.	FINGERPRINT	33
14.6.	ERROR-CODE	33
14.7.	REALM	34
14.8.	NONCE	35
14.9.	UNKNOWN-ATTRIBUTES	35
14.10.	SERVER	35
14.11.	ALTERNATE-SERVER	36
15.	Security Considerations	36
15.1.	Attacks against the Protocol	36
15.1.1.	Outside Attacks	36
15.1.2.	Inside Attacks	36
15.2.	Attacks Affecting the Usage	36
15.2.1.	Attack I: DDoS Against a Target	37
15.2.2.	Attack II: Silencing a Client	37
15.2.3.	Attack III: Assuming the Identity of a Client	37
15.2.4.	Attack IV: Eavesdropping	38
15.3.	Hash Agility Plan	38
16.	IAB Considerations	38
17.	IANA Considerations	39
17.1.	STUN Methods Registry	39
17.2.	STUN Attribute Registry	39
17.3.	STUN Error Code Registry	40
18.	Changes Since RFC 3489	40
19.	Acknowledgements	42
20.	References	42
20.1.	Normative References	42

20.2.	Informational References	43
Appendix A.	C Snippet to Determine STUN Message Types	44
	Authors' Addresses	44
	Intellectual Property and Copyright Statements	46

[1.](#) Introduction

The protocol defined in this specification, Session Traversal Utilities for NAT, provides a toolkit of functions for dealing with NATs. It provides a means for an endpoint to determine the IP address and port allocated by a NAT that corresponds to its private IP address and port. It also provides a way for an endpoint to keep a NAT binding alive. With some extensions, the protocol can be used to do connectivity checks between two endpoints [[I-D.ietf-mmusic-ice](#)], or to relay packets between two endpoints [[I-D.ietf-behave-turn](#)].

In keeping with its toolkit nature, this specification defines an extensible packet format, defines operation over several transport protocols, and provides for two forms of authentication.

STUN is intended to be used in context of one or more NAT traversal solutions. These solutions are known as STUN usages. Each usage describes how STUN is utilized to achieve the NAT traversal solution. Typically, a usage indicates when STUN messages get sent, which optional attributes to include, what server is used, and what authentication mechanism is to be used. Interactive Connectivity Establishment (ICE) [[I-D.ietf-mmusic-ice](#)] is one usage of ICE. SIP Outbound [[I-D.ietf-sip-outbound](#)] is another usage of ICE. In some cases, a usage will require extensions to STUN. A STUN extension can be in the form of new methods, attributes, or error response codes. More information on STUN usages can be found in [Section 13](#).

2. Evolution from [RFC 3489](#)

STUN was originally defined in [RFC 3489](#) [[RFC3489](#)]. That specification, sometimes referred to as "classic STUN", represented itself as a complete solution to the NAT traversal problem. In that solution, a client would discover whether it was behind a NAT, determine its NAT type, discover its IP address and port on the public side of the outermost NAT, and then utilize that IP address and port within the body of protocols, such as the Session Initiation Protocol (SIP) [[RFC3261](#)]. However, experience since the publication of [RFC 3489](#) has found that classic STUN simply does not work sufficiently well to be a deployable solution. The address and port learned through classic STUN are sometimes usable for communications with a peer, and sometimes not. Classic STUN provided no way to discover whether it would, in fact, work or not, and it provided no remedy in cases where it did not. Furthermore, classic STUN's algorithm for classification of NAT types was found to be faulty, as many NATs did not fit cleanly into the types defined there. Classic STUN also had security vulnerabilities which required an extremely

complicated mechanism to address, and despite the complexity of the mechanism, were not fully remedied.

For these reasons, this specification obsoletes [RFC 3489](#), and instead describes STUN as a tool that is utilized as part of a complete NAT traversal solution. ICE is a complete NAT traversal solutions for protocols based on the offer/answer [[RFC3264](#)] methodology, such as SIP. SIP Outbound is a complete solution for traversal of SIP signaling, and it uses STUN in a very different way. Though it is possible that a protocol may be able to use STUN by itself (classic STUN) as a traversal solution, such usage is not described here and is strongly discouraged for the reasons described above.

The on-the-wire protocol described here is changed only slightly from classic STUN. The protocol now runs over TCP in addition to UDP. Extensibility was added to the protocol in a more structured way. A magic-cookie mechanism for demultiplexing STUN with application protocols was added by stealing 32 bits from the 128 bit transaction ID defined in [RFC 3489](#), allowing the change to be backwards compatible. Mapped addresses are encoded using a new exclusive-or format. There are other, more minor changes. See [Section 18](#) for a more complete listing.

Due to the change in scope, STUN has also been renamed from "Simple Traversal of UDP Through NAT" to "Session Traversal Utilities for NAT". The acronym remains STUN, which is all anyone ever remembers anyway.

3. Overview of Operation

This section is descriptive only.

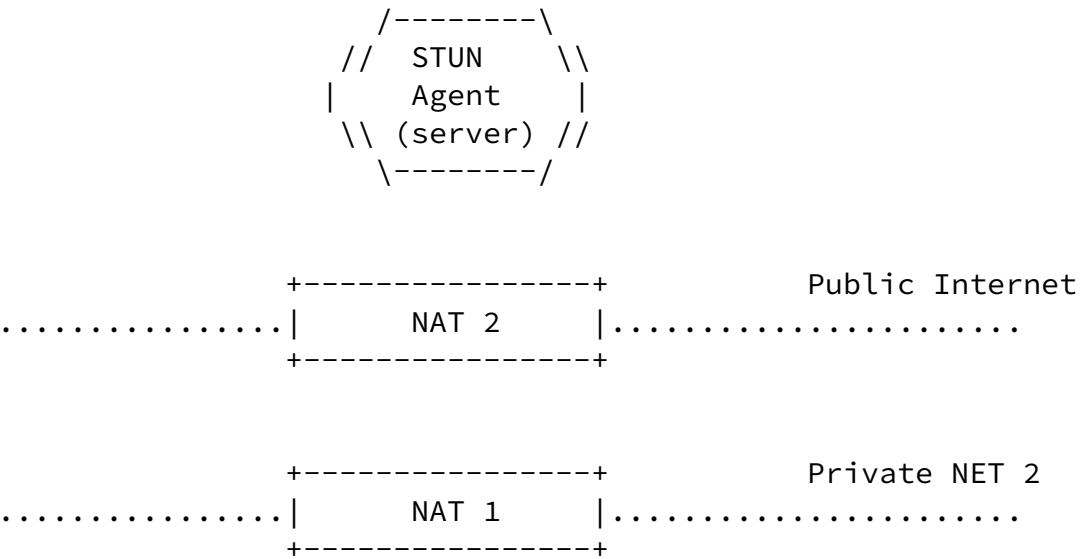




Figure 1: One possible STUN Configuration

One possible STUN configuration is shown in Figure 1. In this configuration, there are two entities (called STUN agents) that implement the STUN protocol. The lower agent in the figure is connected to private network 1. This network connects to private network 2 through NAT 1. Private network 2 connects to the public Internet through NAT 2. The upper agent in the figure resides on the public Internet.

STUN is a client-server protocol. It supports two types of transactions. One is a request/response transaction in which a client sends a request to a server, and the server returns a response. The second is an indication transaction in which a client sends an indication to the server and the server does not respond. Both types of transactions include a transaction ID, which is a randomly selected 96-bit number. For request/response transactions, this transaction ID allows the client to associate the response with the request that generated it; for indications, this simply serves as a debugging aid.

All STUN messages start with a fixed header that includes a method, a class, and the transaction ID. The method indicates which of the various requests or indications this is; this specification defines just one method, Binding, but other methods are expected to be

defined in other documents. The class indicates whether this is a request, a (success) response, an error response, or an indication. Following the fixed header comes zero or more attributes, which are type-length-value extensions that convey additional information for the specific message.

This document defines a single method called Binding. The Binding method can be used either in request/response transactions or in

indication transactions. When used in request/response transactions, the Binding method can be used to determine the particular "binding" a NAT has allocated to a STUN client. When used in either request/response or in indication transactions, the Binding method can also be used to keep these "bindings" alive.

In the Binding request/response transaction, a Binding Request is sent from a STUN client to a STUN server. When the Binding Request arrives at the STUN server, it may have passed through one or more NATs between the STUN client and the STUN server (in Figure 1, there were two such NATs). As the Binding Request message passes through a NAT, the NAT will modify the source transport address (that is, the source IP address and the source port) of the packet. As a result, the source transport address of the request received by the server will be the public IP address and port created by the NAT closest to the server. This is called a reflexive transport address. The STUN server copies that source transport address into an XOR-MAPPED-ADDRESS attribute in the STUN Binding Response and sends the Binding Response back to the the STUN client. As this packet passes back through a NAT, the NAT will modify the destination transport address, but the transport address in the XOR-MAPPED-ADDRESS attribute within the body of the STUN response will remain untouched. In this way, the client can learn its reflexive transport address allocated by the outermost NAT with respect to the STUN server.

In some usages, STUN must be multiplexed with other protocols (e.g., [[I-D.ietf-mmusic-ice](#)], [[I-D.ietf-sip-outbound](#)]). In these usages, there must be a way to inspect a packet and determine if it is a STUN packet or not. STUN provides two fields in the STUN header with fixed values that can be used for this purpose. If this is not sufficient, then STUN packets can also contain a FINGERPRINT value which can further be used to distinguish the packets.

STUN has optional mechanisms for providing authentication and message integrity when required. These mechanisms revolve around the use of a username, password, and message-integrity value. Two of these mechanisms, the long-term credential mechanism and the short-term credential mechanism, are defined in this specification. Each usage specifies the mechanisms allowed with that usage.

In the short-term credential mechanism, the client and the server

exchange a username and password through some out-of-band method prior to the STUN exchange. For example, in the ICE usage [[I-D.ietf-mmusic-ice](#)] the two endpoints use out-of-band signaling to exchange a username and password. The client then includes the username and a message-integrity value in the request message, where the message-integrity value is computed as a cryptographic hash of the message contents and the password. If the server replies with a success response, then the response will include a message-integrity value (computed using the same username and password), but the username is not included. Error responses are not message-integrity protected.

In the long-term credential mechanism, the client and server share a pre-provisioned username and password and perform a digest challenge/response exchange inspired by (but differing in details) to the one defined for HTTP [[RFC2617](#)]. Initially, the client sends a request message (e.g., a Binding Request) without any username or message-integrity value included. The server replies with an error response indicating that the request must be authenticated. This error response includes a realm value and a nonce value. The client then uses the realm value to help it select a username and password (for example, the client might have a number of username and password combinations stored, each one keyed by a different realm value). The client then retries the request, this time including the realm, the username, the nonce, and a message-integrity value in the request, where the message-integrity value is computed as a cryptographic hash of the message contents and the password. The nonce is provided by the server and merely echoed by the client into the request. It is chosen by the server such that it encodes information about the client, the time-of-day, or other parameters. In this way, if an attacker should attempt to replay the request, the server would find the nonce invalid, and then reject the request. If the server replies with a success response, then the response will include a message-integrity value (computed using the same username and password), but realm, username, and nonce are not included. Error responses are not message-integrity protected. If the client has further requests to send, it can try to reuse the same username, realm, and nonce values. If the server does not accept them, it will reply with an error response giving a realm and nonce value again.

[4.](#) Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[RFC2119](#)] and indicate requirement levels for compliant STUN

implementations.

5. Definitions

STUN Agent: An entity that implements the STUN protocol. Agents can act as STUN clients for some transactions and as STUN servers for other transactions.

STUN Client: A logical role in the STUN protocol. A STUN client sends STUN requests or STUN indications, and receives STUN responses. The term "STUN client" is also used colloquially to refer to a STUN agent that only acts as a STUN client.

STUN Server: A logical role in the STUN protocol. A STUN server receives STUN requests or STUN indications and sends STUN responses. The term "STUN server" is also used colloquially to refer to a STUN agent that only acts as a STUN server.

Transport Address: The combination of an IP address and port number (such as a UDP or TCP port number).

Reflexive Transport Address: A transport address learned by a client that identifies that client as seen by another host on an IP network, typically a STUN server. When there is an intervening NAT between the client and the other host, the reflexive transport address represents the mapped address allocated to the client on the public side of the NAT. Reflexive transport addresses are learned from the mapped address attribute (MAPPED-ADDRESS or XOR-MAPPED-ADDRESS) in STUN responses.

Mapped Address: Same meaning as Reflexive Address. This term is retained only for historic reasons and due to the naming of the MAPPED-ADDRESS and XOR-MAPPED-ADDRESS attributes.

Long Term Credential: A username and associated password that represent a shared secret between client and server. Long term credentials are generally granted to the client when a subscriber enrolls in a service and persist until the subscriber leaves the service or explicitly changes the credential.

Long Term Password: The password from a long term credential.

Short Term Credential: A temporary username and associated password which represent a shared secret between client and server. Short term credentials are obtained through some kind of protocol

mechanism between the client server, preceding the STUN exchange. A short term credential has an explicit temporal scope, which may

Internet-Draft

STUN

July 2007

be based on a specific amount of time (such as 5 minutes) or on an event (such as termination of a SIP dialog). The specific scope of a short term credential is defined by the application usage.

Short Term Password: The password component of a short term credential.

STUN Indication: A STUN message that does not receive a response

Attribute: The STUN term for a Type-Length-Value (TLV) object that can be added to a STUN message. Attributes are divided into two types: comprehension-required and comprehension-optional. STUN agents can safely ignore comprehension-optional attributes they don't understand, but cannot successfully process a message if it contains comprehension-required attributes that are not understood.

RT0: Retransmission TimeOut

6. STUN Message Structure

STUN messages are encoded in binary using network-oriented format (most significant byte or octet first, also commonly known as big-endian). The transmission order is described in detail in [Appendix B of RFC791](#) [RFC0791]. Unless otherwise noted, numeric constants are in decimal (base 10).

All STUN messages MUST start with a 20-byte header followed by zero or more Attributes. The STUN header contains a STUN message type, magic cookie, transaction ID, and message length.

[illegible]

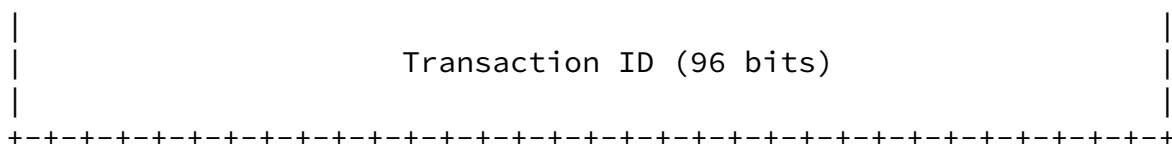


Figure 2: Format of STUN Message Header

The most significant two bits of every STUN message MUST be zeroes. This can be used to differentiate STUN packets from other protocols

when STUN is multiplexed with other protocols on the same port.

The message type defines the message class (request, success response, failure response, or indication) and the message method (the primary function) of the STUN message. Although there are four message classes, there are only two types of transactions in STUN: request/response transactions (which consist of a request message and a response message), and indication transactions (which consists a single indication message). Response classes are split into error and success responses to aid in quickly processing the STUN message.

The message type field is decomposed further into the following structure:

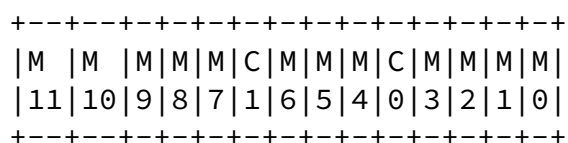


Figure 3: Format of STUN Message Type Field

Here the bits in the message type field are shown as most-significant (M11) through least-significant (M0). M11 through M0 represent a 12-bit encoding of the method. C1 and C0 represent a 2 bit encoding of the class. A class of 0b00 is a Request, a class of 0b01 is an indication, a class of 0b10 is a success response, and a class of 0b11 is an error response. This specification defines a single method, Binding. The method and class are orthogonal, so that four each method, a request, success response, error response and indication are defined for that method.

For example, a Binding Request has class=0b00 (request) and

method=0b00000000000001 (Binding), and is encoded into the first 16 bits as 0x0001. A Binding response has class=0b10 (success response) and method=0b00000000000001, and is encoded into the first 16 bits as 0x0101.

Note: This unfortunate encoding is due to assignment of values in [\[RFC3489\]](#) which did not consider encoding Indications, Success, and Errors using bit fields.

The magic cookie field MUST contain the fixed value 0x2112A442 in network byte order. In [RFC 3489](#) [\[RFC3489\]](#), this field was part of the transaction ID; placing the magic cookie in this location allows a server to detect if the client will understand certain attributes that were added in this revised specification. In addition, it aids in distinguishing STUN packets from packets of other protocols when STUN is multiplexed with those other protocols on the same port.

The transaction ID is a 96 bit identifier, used to uniquely identify STUN transactions. The transaction ID is chosen by the STUN client. It primarily serves to correlate requests with responses, though it also plays a small role in helping to prevent certain types of attacks. As such, the transaction ID MUST be uniformly and randomly chosen from the interval 0 .. 2**96-1. Resends of the same request reuse the same transaction ID, but the client MUST choose a new transaction ID for new transactions unless the new request is bit-wise identical to the previous request and sent from the same transport address to the same IP address. Success and error responses MUST carry the same transaction ID as their corresponding request. When an agent is acting as a STUN server and STUN client on the same port, the transaction IDs in requests sent by the agent have no relationship to the transaction IDs in requests received by the agent.

The message length MUST contain the size, in bytes, of the message not including the 20 byte STUN header. Since all STUN attributes are padded to a multiple of four bytes, the last two bits of this field are always zero. This provides another way to distinguish STUN packets from packets of other protocols.

Following the STUN fixed portion of the header are zero or more attributes. Each attribute is TLV (type-length-value) encoded. The details of the encoding, and of the attributes themselves is given in

[Section 14.](#)

[7.](#) Base Protocol Procedures

This section defines the base procedures of the STUN protocol. It describes how messages are formed, how they are sent, and how they are processed when they are received. It also defines the detailed processing of the Binding method. Other sections in this document describe optional procedures that a usage may elect to use in certain situations. Other documents may define other extensions to STUN, by adding new methods, new attributes, or new error response codes.

[7.1.](#) Forming a Request or an Indication

When formulating a request or indication message, the client **MUST** follow the rules in [Section 6](#) when creating the header. In addition, the message class **MUST** be either "Request" or "Indication" (as appropriate), and the method must be either Binding or some method defined in another document.

The client then adds any attributes specified by the method or the usage. For example, some usages may specify that the client use an

authentication method ([Section 10](#)) or the FINGERPRINT attribute ([Section 8](#)).

For the Binding method with no authentication, no attributes are required unless the usage specifies otherwise.

[7.2.](#) Sending the Request or Indication

The client then sends the request to the server. This document specifies how to send STUN messages over UDP, TCP, or TLS-over-TCP; other transport protocols may be added in the future. The STUN usage must specify which transport protocol is used, and how the client determines the IP address and port of the server. [Section 9](#) describes a DNS-based method of determining the IP address and port of a server which a usage may elect to use.

At any time, a client **MAY** have multiple outstanding STUN requests with the same STUN server (that is, multiple transactions in

progress, with different transaction ids).

7.2.1. Sending over UDP

When running STUN over UDP it is possible that the STUN message might be dropped by the network. Reliability of STUN request/response transactions is accomplished through retransmissions of the request message by the client application itself. STUN indications are not retransmitted; thus indication transactions over UDP are not reliable.

A client SHOULD retransmit a STUN request message starting with an interval of RTO ("Retransmission TimeOut"), doubling after each retransmission. The RTO is an estimate of the round-trip-time, and is computed as described in [RFC 2988](#) [RFC2988], with two exceptions. First, the initial value for RTO SHOULD be configurable (rather than the 3s recommended in [RFC 2988](#)). In fixed- line access links, a value of 100ms is RECOMMENDED. Secondly, the value of RTO MUST NOT be rounded up to the nearest second. Rather, a 1ms accuracy MUST be maintained. As with TCP, the usage of Karn's algorithm is RECOMMENDED. When applied to STUN, it means that RTT estimates SHOULD NOT be computed from STUN transactions which result in the retransmission of a request.

The value for RTO SHOULD be cached by an client after the completion of the transaction, and used as the starting value for RTO for the next transaction to the same server (based on equality of IP address). The value SHOULD be considered stale and discarded after 10 minutes.

Retransmissions continue until a response is received, or until a total of 7 requests have been sent. If, after the last request, a duration equal to 16 times the RTO has passed without a response, the client SHOULD consider the transaction to have failed. A STUN transaction over UDP is also considered failed if there has been a transport failure of some sort, such as a fatal ICMP error. For example, assuming an RTO of 100ms, requests would be sent at times 0ms, 100ms, 300ms, 700ms, 1500ms, 3100ms, and 6300ms. If the client has not received a response after 7900ms, the client will consider the transaction to have timed out.

7.2.2. Sending over TCP or TLS-over-TCP

For TCP and TLS-over-TCP, the client opens a TCP connection to the server.

In some usage of STUN, STUN is sent as the only protocol over the TCP connection. In this case, it can be sent without the aid of any additional framing or demultiplexing. In other usages, or with other extensions, it may be multiplexed with other data over a TCP connection. In that case, STUN MUST be run ontop of some kind of framing protocol, specified by the usage or extension, which allows for the agent to extract complete STUN messages and complete application layer messages.

For TLS-over-TCP, the TLS_RSA_WITH_AES_128_CBC_SHA ciphersuite MUST be supported at a minimum. Implementations MAY also support any other ciphersuite. When it receives the TLS Certificate message, the client SHOULD verify the certificate and inspect the site identified by the certificate. If the certificate is invalid, revoked, or if it does not identify the appropriate party, the client MUST NOT send the STUN message or otherwise proceed with the STUN transaction. The client MUST verify the identity of the server. To do that, it follows the identification procedures defined in Section 3.1 of [RFC 2818](#) [[RFC2818](#)]. Those procedures assume the client is dereferencing a URI. For purposes of usage with this specification, the client treats the domain name or IP address used in [Section 8.1](#) as the host portion of the URI that has been dereferenced. If DNS was not used, the client MUST be configured with a set of authorized domains whose certificates will be accepted.

Reliability of STUN over TCP and TLS-over-TCP is handled by TCP itself, and there are no retransmissions at the STUN protocol level. However, for a request/response transaction, if the client has not received a response after 7900ms, it considers the transaction to have timed out. This value has been chosen to equalize the TCP and UDP timeouts for the default initial RT0.

In addition, if the client is unable to establish the TCP connection, or the TCP connection is reset or fails before a response is received, any request/response transaction in progress is considered to have failed

The client MAY send multiple transactions over a single TCP (or TLS-over-TCP) connection, and it MAY send another request before receiving a response to the previous. The client SHOULD keep the connection open until it

- o has no further STUN requests or indications to send over that connection, and;
- o has no plans to use any resources (such as a mapped address (MAPPED-ADDRESS or XOR-MAPPED-ADDRESS) or relayed address [[I-D.ietf-behave-turn](#)]) that were learned through STUN requests sent over that connection, and;
- o if multiplexing other application protocols over that port, has finished using that other application, and;
- o if using that learned port with a remote peer, has established communications with that remote peer, as is required by some TCP NAT traversal techniques (e.g., [[I-D.ietf-mmusic-ice-tcp](#)]).

At the server end, the server SHOULD keep the connection open, and let the client close it. If a server becomes overloaded and needs to close connections to free up resources, it SHOULD close an existing connection rather than reject new connection requests. The server SHOULD NOT close a connection if a request was received over that connection for which a response was not sent. A server MUST NOT ever open a connection back towards the client in order to send a response.

[7.3.](#) Receiving a STUN Message

This section specifies the processing of a STUN message. The processing specified here is for STUN messages as defined in this specification; additional rules for backwards compatibility are defined in in [Section 12](#). Those additional procedures are optional, and usages can elect to utilize them. First, a set of processing operations are applied that are independent of the class. This is followed by class-specific processing, described in the subsections which follow.

When a STUN agent receives a STUN message, it first checks that the message obeys the rules of [Section 6](#). It checks that the first two bits are 0, that the magic cookie field has the correct value, that

the message length is sensible, and that the method value is a supported method. If the message-class is Success Response or Error Response, the agent checks that the transaction ID matches a transaction that is still in progress. If the FINGERPRINT extension is being used, the agent checks that the FINGERPRINT attribute is present and contains the correct value. If any errors are detected, the message is silently discarded. In the case when STUN is being multiplexed with another protocol, an error may indicate that this is not really a STUN message; in this case, the agent should try to parse the message as a different protocol.

The STUN agent then does any checks that are required by a authentication mechanism that the usage has specified (see [Section 10](#)).

Once the authentication checks are done, the STUN agent checks for unknown attributes and known-but-unexpected attributes in the message. Unknown comprehension-optional attributes **MUST** be ignored by the agent. Known-but-unexpected attributes **SHOULD** be ignored by the agent. Unknown comprehension-required attributes cause processing that depends on the message-class and is described below.

At this point, further processing depends on the message class of the request.

[7.3.1](#). Processing a Request

If the request contains one or more unknown comprehension-required attributes, the server replies with an error response with an error code of 420 Unknown attributes, and includes an UNKNOWN-ATTRIBUTES attribute in the response that lists the unknown comprehension-required attributes.

The server then does any additional checking that the method or the specific usage requires. If all the checks succeed, the server formulates a success response as described below.

If the request uses UDP transport and is a retransmission of a request for which the server has already generated a success response within the last 10 seconds, the server **MUST** retransmit the same success response. One way for a server to do this is to remember all transaction IDs received over UDP and their corresponding responses in the last 10 seconds. Another way is to reprocess the request and recompute the response. The latter technique **MUST** only be applied to requests which are idempotent and result in the same success response for the same request. The Binding method is considered to idempotent in this way (even though certain rare network events could cause the

SHOULD state whether their request types have this property or not.

[7.3.1.1](#). Forming a Success or Error Response

When forming the response (success or error), the server follows the rules of [section 6](#). The method of the response is the same as that of the request, and the message class is either "Success Response" or "Error Response".

For an error response, the server MUST add an ERROR-CODE attribute containing the error code specified in the processing above. The reason phrase is not fixed, but SHOULD be something suitable for the error code. For certain errors, additional attributes are added to the message. These attributes are spelled out in the description where the error code is specified. For example, for an error code of 420 Unknown Attribute, the server MUST include an UNKNOWN-ATTRIBUTES attribute. Certain authentication errors also cause attributes to be added (see [Section 10](#)). Extensions may define other errors and/or additional attributes to add in error cases.

If the server authenticated the request using an authentication mechanism, then the server SHOULD add the appropriate authentication attributes to the response (see [Section 10](#)).

The server also adds any attributes required by the specific method or usage. In addition, the server SHOULD add a SERVER attribute to the message.

For the Binding method, no additional checking is required unless the usage specifies otherwise. When forming the success response, the server adds a XOR-MAPPED-ADDRESS attribute to the response, where the contents of the attribute are the source transport address of the request message. For UDP, this is the source IP address and source UDP port of the request message. For TCP and TLS-over-TCP, this is the source IP address and source TCP port of the TCP connection as seen by the server.

[7.3.1.2](#). Sending the Success or Error Response

The response (success or error) is sent over the same transport as

the request was received on. If the request was received over UDP, the destination IP address and port of the response is the source IP address and port of the received request message, and the source IP address and port of the response is equal to the destination IP address and port of the received request message. If the request was received over TCP or TLS-over-TCP, the response is sent back on the same TCP connection as the request was received on.

[7.3.2.](#) Processing an Indication

If the indication contains unknown comprehension-required attributes, the indication is discarded and processing ceases.

The server then does any additional checking that the method or the specific usage requires. If all the checks succeed, the server then processes the indication. No response is generated for an indication.

For the Binding method, no additional checking or processing is required, unless the usage specifies otherwise. The mere receipt of the message by the server has refreshed the "bindings" in the intervening NATs.

Since indications are not re-transmitted over UDP (unlike requests), there is no need to handle re-transmissions of indications at the server.

[7.3.3.](#) Processing a Success Response

If the success response contains unknown comprehension-required attributes, the response is discarded and the transaction is considered to have failed.

The client then does any additional checking that the method or the specific usage requires. If all the checks succeed, the client then processes the success response.

For the Binding method, the client checks that the XOR-MAPPED-ADDRESS attribute is present in the response. The client checks the address family specified. If it is an unsupported address family, the attribute SHOULD be ignored. If it is an unexpected but supported

address family (for example, the Binding transaction was sent over IPv4, but the address family specified is IPv6), then the client MAY accept and use the value.

[7.3.4.](#) Processing an Error Response

If the error response contains unknown comprehension-required attributes, or if the error response does not contain an ERROR-CODE attribute, then the transaction is simply considered to have failed.

The client then does any processing specified by the authentication mechanism (see [Section 10](#)). This may result in a new transaction attempt.

The processing at this point depends on the error-code, the method,

and the usage; the following are the default rules:

- o If the error code is 300 through 399, the client SHOULD consider the transaction as failed unless the ALTERNATE-SERVER extension is being used. See [Section 11](#).
- o If the error code is 400 through 499, the client declares the transaction failed; in the case of 420, the response should contain a UNKNOWN-ATTRIBUTES attribute that gives additional information.
- o If the error code is 500 through 599, the client MAY resend the request; clients that do so MUST limit the number of times they do this. Any other error code causes the client to consider the transaction failed.

[8.](#) FINGERPRINT Mechanism

This section describes an optional mechanism for STUN that aids in distinguishing STUN messages from packets of other protocols when the two are multiplexed on the same transport address. This mechanism is optional, and a STUN usage must describe if and when it is used.

In some usages, STUN messages are multiplexed on the same transport address as other protocols, such as RTP. In order to apply the

processing described in [Section 7](#), STUN messages must first be separated from the application packets. [Section 6](#) describes three fixed fields in the STUN header that can be used for this purpose. However, in some cases, these three fixed fields may not be sufficient.

When the FINGERPRINT extension is used, an agent includes the FINGERPRINT attribute in messages it sends to another agent. [Section 14.5](#) describes the placement and value of this attribute. When the agent receives what it believes is a STUN message, then, in addition to other basic checks, the agent also checks that the message contains a FINGERPRINT attribute and that the attribute contains the correct value (see [Section 7.3](#)). This additional check helps the agent detect messages of other protocols that might otherwise seem to be STUN messages.

[9](#). DNS Discovery of a Server

This section describes an optional procedure for STUN that allows a client to use DNS to determine the IP address and port of a server. A STUN usage must describe if and when this extension is used. To

use this procedure, the client must have a domain name and a service name; the usage must also describe how the client obtains these.

When a client wishes to locate a STUN server in the public Internet that accepts Binding Request/Response transactions, the SRV service name is "stun". STUN usages MAY define additional DNS SRV service names.

The domain name is resolved to a transport address using the SRV procedures specified in [[RFC2782](#)]. The DNS SRV service name is the service name provided as input to this procedure. The protocol in the SRV lookup is the transport protocol the client will run STUN over: "udp" for UDP, "tcp" for TCP, and "tls" for TLS-over-TCP. If, in the future, additional SRV records are defined for TLS over other transport protocols, those will need to utilize an SRV transport token of the form "tls-foo" for transport protocol "foo".

The procedures of [RFC 2782](#) are followed to determine the server to contact. [RFC 2782](#) spells out the details of how a set of SRV records

are sorted and then tried. However, [RFC2782](#) only states that the client should "try to connect to the (protocol, address, service)" without giving any details on what happens in the event of failure. When following these procedures, if the STUN transaction times out without receipt of a response, the client SHOULD retry the request to the next server in the list of servers from the DNS SRV response. Such a retry is only possible for request/response transmissions, since indication transactions generate no response or timeout.

The default port for STUN requests is 3478, for both TCP and UDP. Administrators SHOULD use this port in their SRV records for UDP and TCP, but MAY use others. There is no default port for STUN over TLS, however a STUN server SHOULD use a port number for TLS different from 3478 so that the server can determine whether the first message it will receive after the TCP connection is set up, is a STUN message or a TLS message.

If no SRV records were found, the client performs an A or AAAA record lookup of the domain name. The result will be a list of IP addresses, each of which can be contacted at the default port using UDP or TCP, independent of the STUN usage. For usages that require TLS, lack of SRV records is equivalent to a failure of the transaction, since the request or indication MUST NOT be sent unless SRV records provided a transport address specifically for TLS.

[10.](#) Authentication and Message-Integrity Mechanisms

This section defines two mechanisms for STUN that a client and server

can use to provide authentication and message-integrity; these two mechanisms are known as the short-term credential mechanism and the long-term credential mechanism. These two mechanisms are optional, and each usage must specify if and when these mechanisms are used. An overview of these two mechanisms is given in .

Each mechanism specifies the additional processing required to use that mechanism, extending the processing specified in [Section 7](#). The additional processing occurs in three different places: when forming a message; when receiving a message immediately after the the basic checks have been performed; and when doing the detailed processing of error responses.

[10.1.](#) Short-Term Credential Mechanism

The short-term credential mechanism assumes that, prior to the STUN transaction, the client and server have used some other protocol to exchange a credential in the form of a username and password. This credential is time-limited. The time-limit is defined by the usage. As an example, in the ICE usage [[I-D.ietf-mmusic-ice](#)], the two endpoints use out-of-band signaling to agree on a username and password, and this username and password is applicable for the duration of the media session.

This credential is used to form a message integrity check in each request and in many responses. There is no challenge and response as in the long term mechanism; consequently, replay is prevented by virtue of the time-limited nature of the credential.

[10.1.1.](#) Forming a Request or Indication

For a request or indication message, the agent MUST include the USERNAME and MESSAGE-INTEGRITY attributes in the message. The HMAC for the MESSAGE-INTEGRITY attribute is computed as described in [Section 14.4](#). The key for the HMAC is the password. Note that the password is never included in the request or indication.

[10.1.2.](#) Receiving a Request or Indication

After the agent has done the basic processing of a message, the agent performs the checks listed below in order specified:

- o If the message does not contain both a MESSAGE-INTEGRITY and a USERNAME attribute:
 - * If the message is a request, the server MUST reject the request with an error response. This response MUST use an error code of 400.

- * If the message is an indication, the server MUST silently discard the indication.
- o If the USERNAME does not contain a username value currently valid within the server:

- * If the message is a request, the server MUST reject the request with an error response. This response MUST use an error code of 401.
- * If the message is an indication, the server MUST silently discard the indication.
- o Using the password associated with the username, compute the value for the message-integrity as described in [Section 14.4](#). If the resulting value does not match the contents of the MESSAGE-INTEGRITY attribute:
 - * If the message is a request, the server MUST reject the request with an error response. This response MUST use an error code of 431.
 - * If the message is an indication, the server MUST silently discard the indication.

If these checks pass, the server continues to process the request or indication. Any response generated by the server MUST include the MESSAGE-INTEGRITY attribute, computed using the username and password utilized to authenticate the request.

If any of the checks fail, the server MUST NOT include a MESSAGE-INTEGRITY or USERNAME attribute in the error response.

[10.1.3](#). Receiving a Response

The processing here takes place prior to the processing in [Section 7.3.3](#) or [Section 7.3.4](#).

The client looks for the MESSAGE-INTEGRITY attribute in the response. If present, the client computes the message integrity over the response as defined in [Section 14.4](#), using the same password it utilized for the request. If the resulting value matches the contents of the MESSAGE-INTEGRITY attribute, the response is considered authenticated. If the value does not match, or if MESSAGE-INTEGRITY was absent, the response MUST be discarded, as if it was never received. This means that retransmits, if applicable, will continue.

10.2. Long-term Credential Mechanism

The long-term credential mechanism relies on a long term credential, in the form of a username and password, that are shared between client and server. The credential is considered long-term since it is assumed that it is provisioned for a user, and remains in effect until the user is no longer a subscriber of the system, or is changed. This is basically a traditional "log-in" username and password given to users.

Because these usernames and passwords are expected to be valid for extended periods of time, replay prevention is provided in the form of a digest challenge. In this mechanism, the client initially sends a request, without offering any credentials or any integrity checks. The server rejects this request, providing the user a realm (used to guide the user or agent in selection of a username and password) and a nonce. The nonce provides the replay protection. It is a cookie, selected by the server, and encoded in such a way as to indicate a duration of validity or client identity from which it is valid. The client retries the request, this time including its username, the realm, and echoing the nonce provided by the server. The client also includes a message-integrity, which provides an HMAC over the entire request, including the nonce. The server validates the nonce, and checks the message-integrity. If they match, the request is authenticated. If the nonce is no longer valid, it is considered "stale", and the server rejects the request, providing a new nonce.

In subsequent requests to the same server, the client reuses the nonce, username, realm and password it used previously. In this way, subsequent requests are not rejected until the nonce becomes invalid by the server, in which case the rejection provides a new nonce to the client.

Note that the long-term credential mechanism cannot be used to protect indications, since indications cannot be challenged. Usages utilizing indications must either use a short-term credential, or omit authentication and message integrity for them.

Since the long-term credential mechanism is susceptible to offline dictionary attacks, deployments SHOULD utilize strong passwords.

For STUN servers used in conjunction with SIP servers, it is desirable to use the same credentials for authentication to the SIP server and STUN server. Typically, SIP systems utilizing SIP's digest authentication mechanism do not actually store the password in the database. Rather, they store a value called $H(A1)$, which is computed as:

Internet-Draft

STUN

July 2007

$$H(A1) = \text{MD5}(\text{username} ":" \text{realm} ":" \text{password})$$

If a system wishes to utilize this credential, the STUN password would be computed by taking the user-entered username and password, and using $H(A1)$ as the STUN password. It is RECOMMENDED that clients utilize this construction for the STUN password.

[10.2.1.](#) Forming a Request

There are two cases when forming a request. In the first case, this is the first request from the client to the server (as identified by its IP address and port). In the second case, the client is submitting a subsequent request once a previous request/response transaction has completed successfully.

[10.2.1.1.](#) First Request

If the client has not completed a successful request/response transaction with the server, it SHOULD omit the USERNAME, MESSAGE-INTEGRITY, REALM, and NONCE attributes. In other words, the very first request is sent as if there were no authentication or message integrity applied.

[10.2.1.2.](#) Subsequent Requests

Once a request/response transaction has completed successfully, the client will have been presented a realm and nonce by the server, and selected a username and password with which it authenticated. The client SHOULD cache the username, password, realm, and nonce for subsequent communications with the server. When the client sends a subsequent request, it SHOULD include the USERNAME, REALM, and NONCE attributes with these cached values. It SHOULD include a MESSAGE-INTEGRITY attributed, computed as described in [Section 14.4](#) using the cached password as the key.

[10.2.2.](#) Receiving a Request

After the server has done the basic processing of a request, it performs the checks listed below in the order specified:

- o If the message:

- * does not contain a MESSAGE-INTEGRITY attribute,
 - * OR, it contains a USERNAME whose value is not a valid username,
- the server MUST generate an error response with an error code of 401. This response MUST include a REALM value. It is RECOMMENDED

that the REALM value by the domain name of the provider of the STUN server. The response MUST include a NONCE, selected by the server.

- o If the message contains a MESSAGE-INTEGRITY attribute, but is missing the USERNAME, REALM or NONCE attributes, the server MUST generate an error response with an error code of 400.
- o If the NONCE is no longer valid, the server MUST generate an error response with an error code of 438 (Stale Nonce). This response MUST include a NONCE and REALM attribute.
- o Using the password associated with the username in the USERNAME attribute, compute the value for the message-integrity as described in [Section 14.4](#). If the resulting value does not match the contents of the MESSAGE-INTEGRITY attribute, the server MUST reject the request with an error response. This response MUST use an error code of 401. It MUST include a REALM and NONCE attribute.

If these checks pass, the server continues to process the request or indication. Any response generated by the server MUST include the MESSAGE-INTEGRITY attribute, computed using the username and password utilized to authenticate the request. The REALM, NONCE, and USERNAME attributes SHOULD NOT be included.

[10.2.3](#). Receiving a Response

The processing here takes place prior to the processing in [Section 7.3.3](#) or [Section 7.3.4](#).

If the response is an error response, with an error code of 401 (Unauthorized), the client SHOULD retry the request with a new transaction. This request MUST contain a USERNAME, determined by the client as the appropriate username for the REALM from the error

response. The request MUST contain the REALM, copied from the error response. The request MUST contain the NONCE, copied from the error response. The request MUST contain the MESSAGE-INTEGRITY attribute, computed using the password associated with the username in the USERNAME attribute. The client MUST NOT perform this retry if it is not changing the USERNAME or REALM or its associated password, from the previous attempt.

If the response is an error response with an error code of 438, the client MUST retry the request, using the new NONCE supplied in the 438 response. This retry MUST also include the USERNAME, REALM and MESSAGE-INTEGRITY.

The client looks for the MESSAGE-INTEGRITY attribute in the response (either success or failure). If present, the client computes the message integrity over the response as defined in [Section 14.4](#), using the same password it utilized for the request. If the resulting value matches the contents of the MESSAGE-INTEGRITY attribute, the response is considered authenticated. If the value does not match, or if MESSAGE-INTEGRITY was absent, the response MUST be discarded, as if it was never received. This means that retransmits, if applicable, will continue.

[11.](#) ALTERNATE-SERVER Mechanism

This section describes a mechanism in STUN that allows a server to redirect a client to another server. This extension is optional, and a usage must define if and when this extension is used. To prevent denial-of-service attacks, this extension MUST only be used in situations where the client and server are using an authentication and message-integrity mechanism.

A server using this extension redirects a client to another server by replying to a request message with an error response message with an error code of 300 (Try Alternate). The server MUST include a ALTERNATE-SERVER attribute in the error response. The error response message MUST be authenticated, which in practice means the request message must have passed the authentication checks.

A client using this extension handles a 300 (Try Alternate) error

code as follows. If the error response has passed the authentication checks, then the client looks for a ALTERNATE-SERVER attribute in the error response. If one is found, then the client considers the current transaction as failed, and re-attempts the request with the server specified in the attribute. The client SHOULD reuse any authentication credentials from the old request in the new transaction.

12. Backwards Compatibility with [RFC 3489](#)

This section define procedures that allow a degree of backwards compatible with the original protocol defined in [RFC 3489](#) [[RFC3489](#)]. This mechanism is optional, meant to be utilized only in cases where a new client can connect to an old server, or vice-a-versa. A usage must define if and when this procedure is used.

[Section 18](#) lists all the changes between this specification and [RFC 3489](#) [[RFC3489](#)]. However, not all of these differences are important, because "classic STUN" was only used in a few specific ways. For the

purposes of this extension, the important changes are the following. In [RFC 3489](#):

- o UDP was the only supported transport;
- o The field that is now the Magic Cookie field was a part of the transaction id field, and transaction ids were 128 bits long;
- o The XOR-MAPPED-ADDRESS attribute did not exist, and the Binding method used the MAPPED-ADDRESS attribute instead
- o There were two comprehension-required attributes, RESPONSE-ADDRESS and CHANGE-REQUEST, that have been removed from this specification.
 - * These attributes are now part of the NAT Behavior Discovery usage.

12.1. Changes to Client Processing

A client that wants to interoperate with a [[RFC3489](#)] server SHOULD

send a request message that uses the Binding method, contains no attributes, and uses UDP as the transport protocol to the server. If successful, the success response received from the server will contain a MAPPED-ADDRESS attribute rather than an XOR-MAPPED-ADDRESS attribute; other than this change, the processing of the response is identical to the procedures described above.

12.2. Changes to Server Processing

A STUN server can detect when a given Binding Request message was sent from an [RFC 3489](#) [RFC3489] client by the absence of the correct value in the Magic Cookie field. When the server detects an [RFC 3489](#) client, it SHOULD copy the value seen in the Magic Cookie field in the Binding Request to the Magic Cookie field in the Binding Response message, and insert a MAPPED-ADDRESS attribute instead of an XOR-MAPPED-ADDRESS attribute.

The client might, in rare situations, include either the RESPONSE-ADDRESS or CHANGE-REQUEST attributes. In these situations, the server will view these as unknown comprehension-required attributes and reply with an error response. Since the mechanisms utilizing those attributes are no longer supported, this behavior is acceptable.

13. STUN Usages

STUN by itself is not a solution to the NAT traversal problem. Rather, STUN defines a toolkit of functions that can be used inside a larger solution. The term "STUN Usage" is used for any solution that uses STUN as a component.

At the time of writing, three STUN usages are defined: Interactive Connectivity Establishment (ICE) [[I-D.ietf-mmusic-ice](#)], Client-initiated connections for SIP [[I-D.ietf-sip-outbound](#)], and NAT Behavior Discovery [[I-D.ietf-behave-nat-behavior-discovery](#)]. Other STUN usages may be defined in the future.

A STUN usage defines how STUN is actually utilized - when to send

requests, what to do with the responses, and which optional procedures defined here (or in an extension to STUN) are to be used. A usage would also define:

- o Which STUN methods are used;
- o What authentication and message integrity mechanisms are used;
- o What mechanisms are used to distinguish STUN messages from other messages. When STUN is run over TCP, a framing mechanism may be required;
- o How a STUN client determines the IP address and port of the STUN server;
- o Whether backwards compatibility to [RFC 3489](#) is required;
- o What optional attributes defined here (such as FINGERPRINT and ALTERNATE-SERVER) or in other extensions are required.

In addition, any STUN usage must consider the security implications of using STUN in that usage. A number of attacks against STUN are known (see the Security Considerations section in this document) and any usage must consider how these attacks can be thwarted or mitigated.

Finally, a usage must consider whether its usage of STUN is an example of the Unilateral Self-Address Fixing approach to NAT traversal, and if so, address the questions raised in [RFC 3424](#).

[14.](#) STUN Attributes

After the STUN header are zero or more attributes. Each attribute MUST be TLV encoded, with a 16 bit type, 16 bit length, and value. Each STUN attribute MUST end on a 32 bit boundary. As mentioned above, all fields in an attribute are transmitted most significant bit first.

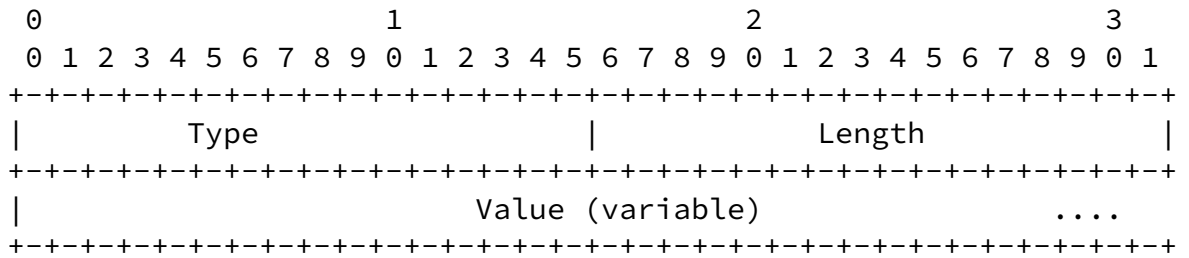


Figure 5: Format of STUN Attributes

The value in the Length field **MUST** contain the length of the Value part of the attribute, prior to padding, measured in bytes. Since STUN aligns attributes on 32 bit boundaries, attributes whose content is not a multiple of 4 bytes are padded with 1, 2 or 3 bytes of padding so that its value contains a multiple of 4 bytes. The padding bits are ignored, and may be any value.

Any attribute type **MAY** appear more than once in a STUN message. Unless specified otherwise, the order of appearance is significant: only the first occurrence needs to be processed by a receiver, and any duplicates **MAY** be ignored by a receiver.

To allow future revisions of this specification to add new attributes if needed, the attribute space is divided into two ranges. Attributes with type values between 0x0000 and 0x7FFF are comprehension-required attributes, which means that the STUN agent cannot successfully process the message unless it understands the attribute. Attributes with type values between 0x8000 and 0xFFFF are comprehension-optional attributes, which means that those attributes can be ignored by the STUN agent if it does not understand them.

The STUN Attribute types defined by this specification are:

Comprehension-required range (0x0000-0x7FFF):

0x0000: (Reserved)
 0x0001: MAPPED-ADDRESS
 0x0006: USERNAME
 0x0007: (Reserved; was PASSWORD)
 0x0008: MESSAGE-INTEGRITY
 0x0009: ERROR-CODE
 0x000A: UNKNOWN-ATTRIBUTES
 0x0014: REALM
 0x0015: NONCE
 0x0020: XOR-MAPPED-ADDRESS

Comprehension-optional range (0x8000-0xFFFF)

0x8022: SERVER
 0x8023: ALTERNATE-SERVER
 0x8028: FINGERPRINT

The rest of this section describes the format of the various attributes defined in this specification.

[14.1.](#) MAPPED-ADDRESS

The MAPPED-ADDRESS attribute indicates a reflexive transport address of the client. It consists of an eight bit address family, and a sixteen bit port, followed by a fixed length value representing the IP address. If the address family is IPv4, the address MUST be 32 bits. If the address family is IPv6, the address MUST be 128 bits. All fields must be in network byte order.

The format of the MAPPED-ADDRESS attribute is:

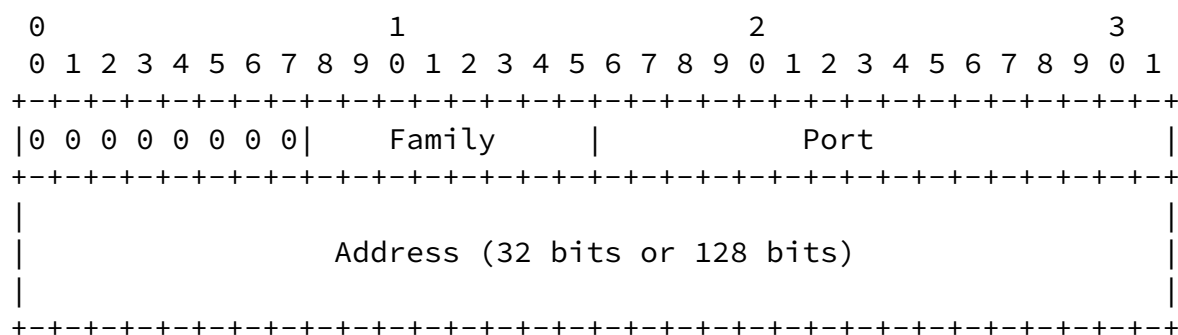


Figure 7: Format of MAPPED-ADDRESS attribute

Internet-Draft

STUN

July 2007

The address family can take on the following values:

0x01:IPv4

0x02:IPv6

The first 8 bits of the MAPPED-ADDRESS MUST be set to 0 and MUST be ignored by receivers. These bits are present for aligning parameters on natural 32 bit boundaries.

This attribute is used only by servers for achieving backwards compatibility with [RFC 3489](#) [[RFC3489](#)] clients.

[14.2.](#) XOR-MAPPED-ADDRESS

The XOR-MAPPED-ADDRESS attribute is identical to the MAPPED-ADDRESS attribute, except that the reflexive transport address is obfuscated through the XOR function.

The format of the XOR-MAPPED-ADDRESS is:

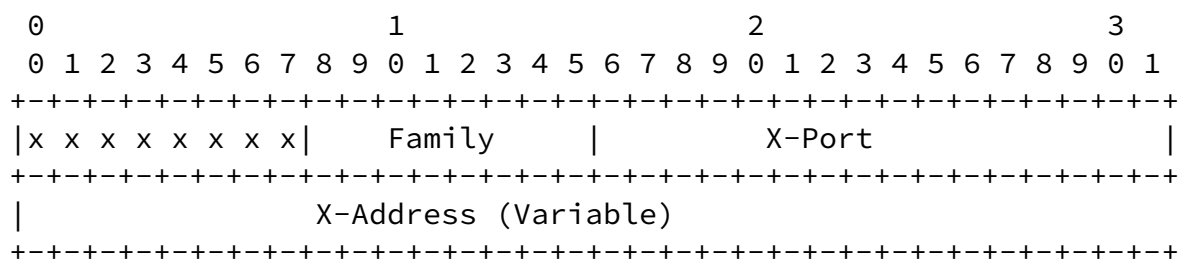


Figure 9: Format of XOR-MAPPED-ADDRESS Attribute

The Family represents the IP address family, and is encoded identically to the Family in MAPPED-ADDRESS.

X-Port is the mapped port, exclusive or'd with most significant 16 bits of the magic cookie. If the IP address family is IPv4, X-Address is the mapped IP address exclusive or'd with the magic cookie. If the IP address family is IPv6, the X-Address is the mapped IP address exclusively or'ed with the magic cookie and the 96-bit transaction ID.

For example, using the "^" character to indicate exclusive or, if the IP address is 192.168.1.1 (0xc0a80101) and the port is 5555 (0x15B3), the X-Port would be 0x15B3 ^ 0x2112 = 0x34A1, and the X-Address would

be $0xc0a80101 \wedge 0x2112A442 = 0xe1baa543$.

The rules for encoding and processing the first 8 bits of the attribute's value, the rules for handling multiple occurrences of the attribute, and the rules for processing addresses families are the

same as for MAPPED-ADDRESS.

NOTE: XOR-MAPPED-ADDRESS and MAPPED-ADDRESS differ only in their encoding of the transport address. The former encodes the transport address by exclusive-or'ing it with the magic cookie. The latter encodes it directly in binary. [RFC 3489](#) originally specified only MAPPED-ADDRESS. However, deployment experience found that some NATs rewrite the 32-bit binary payloads containing the NAT's public IP address, such as STUN's MAPPED-ADDRESS attribute, in the well-meaning but misguided attempt at providing a generic ALG function. Such behavior interferes with the operation of STUN and also causes failure of STUN's message integrity checking.

[14.3.](#) USERNAME

The USERNAME attribute is used for message integrity. It identifies the username and password combination used in the message integrity check.

The value of USERNAME is a variable length value. It MUST contain a UTF-8 encoded sequence of less than 128 characters.

[14.4.](#) MESSAGE-INTEGRITY

The MESSAGE-INTEGRITY attribute contains an HMAC-SHA1 [[RFC2104](#)] of the STUN message. The MESSAGE-INTEGRITY attribute can be present in any STUN message type. Since it uses the SHA1 hash, the HMAC will be 20 bytes. The text used as input to HMAC is the STUN message, including the header, up to and including the attribute preceding the MESSAGE-INTEGRITY attribute. With the exception of the FINGERPRINT attribute, which appears after MESSAGE-INTEGRITY, agents MUST ignore all other attributes that follow MESSAGE-INTEGRITY.

The key used as input to HMAC is the password.

Since the hash is computed over the entire STUN message, it includes

the length field from the STUN message header. This length indicates the length of the entire message, including the MESSAGE-INTEGRITY attribute itself. Consequently, the MESSAGE-INTEGRITY attribute MUST be inserted into the message (with dummy content) prior to the computation of the integrity check. Once the computation is performed, the value of the attribute can be filled in. This ensures the length has the correct value when the hash is performed. Similarly, when validating the MESSAGE-INTEGRITY, the length field should be adjusted to point to the end of the MESSAGE-INTEGRITY attribute prior to calculating the HMAC. Such adjustment is necessary when attributes, such as FINGERPRINT, appear after MESSAGE-INTEGRITY.

[14.5.](#) FINGERPRINT

The FINGERPRINT attribute may be present in all STUN messages. The value of the attribute is computed as the CRC-32 of the STUN message up to (but excluding) the FINGERPRINT attribute itself, xor-d with the 32 bit value 0x5354554e (the XOR helps in cases where an application packet is also using CRC-32 in it). The 32 bit CRC is the one defined in ITU V.42 [[ITU.V42.1994](#)], which has a generator polynomial of $x^{32}+x^{26}+x^{23}+x^{22}+x^{16}+x^{12}+x^{11}+x^{10}+x^8+x^7+x^5+x^4+x^2+x+1$. When present, the FINGERPRINT attribute MUST be the last attribute in the message, and thus will appear after MESSAGE-INTEGRITY.

The FINGERPRINT attribute can aid in distinguishing STUN packets from packets of other protocols. See [Section 8](#).

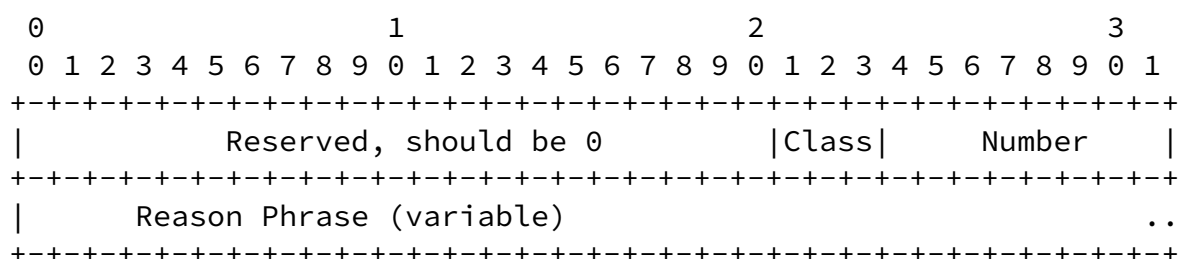
When using the FINGERPRINT attribute in a message, the attribute is first placed into the message with a dummy value, then the CRC is computed, and then the value of the attribute is updated. If the MESSAGE-INTEGRITY attribute is also present, then it must be present with the correct message-integrity value before the CRC is computed, since the CRC is done over the value of the MESSAGE-INTEGRITY attribute as well.

[14.6.](#) ERROR-CODE

The ERROR-CODE attribute is used in Error Response messages. It contains a numeric error code value in the range of 300 to 699 plus a textual reason phrase encoded in UTF-8, and is consistent in its code assignments and semantics with SIP [[RFC3261](#)] and HTTP [[RFC2616](#)]. The

reason phrase is meant for user consumption, and can be anything appropriate for the error code. Recommended reason phrases for the defined error codes are presented below. The reason phrase MUST be a UTF-8 encoded sequence of less than 128 characters.

To facilitate processing, the class of the error code (the hundreds digit) is encoded separately from the rest of the code.



The Reserved bits SHOULD be 0, and are for alignment on 32-bit boundaries. Receivers MUST ignore these bits. The Class represents the hundreds digit of the error code. The value MUST be between 3

and 6. The number represents the error code modulo 100, and its value MUST be between 0 and 99.

The following error codes, along with their recommended reason phrases (in brackets) are defined:

300 Try Alternate: The client should contact an alternate server for this request. This error response MUST only be sent if the request included a USERNAME attribute and a valid MESSAGE-INTEGRITY attribute; otherwise it MUST NOT be sent and error code 400 is suggested. This error response MUST be protected with the MESSAGE-INTEGRITY attribute, and receivers MUST validate the MESSAGE-INTEGRITY of this response before redirecting themselves to an alternate server.

Note: failure to generate and validate message-integrity for a 300 response allows an on-path attacker to falsify a 300 response thus causing subsequent STUN messages to be sent to a victim.

400 Bad Request: The request was malformed. The client SHOULD NOT retry the request without modification from the previous

attempt. The server may not be able to generate a valid MESSAGE-INTEGRITY for this error, so the client MUST NOT expect a valid MESSAGE-INTEGRITY attribute on this response.

- 401 Unauthorized: The request did not contain the expected MESSAGE-INTEGRITY attribute. The server MAY include the MESSAGE-INTEGRITY attribute in its error response.
- 420 Unknown Attribute: The server received STUN packet containing a comprehension-required attribute which it did not understand. The server MUST put this unknown attribute in the UNKNOWN-ATTRIBUTE attribute of its error response.
- 438 Stale Nonce: The NONCE used by the client was no longer valid. The client should retry, using the NONCE provided in the response.
- 500 Server Error: The server has suffered a temporary error. The client should try again.

[14.7.](#) REALM

The REALM attribute may be present in requests and responses. It contains text which meets the grammar for "realm-value" as described in [RFC 3261](#) [[RFC3261](#)] but without the double quotes and their surrounding whitespace. That is, it is an unquoted realm-value. It

MUST be a UTF-8 encoded sequence of less than 128 characters.

Presence of the REALM attribute in a request indicates that long-term credentials are being used for authentication. Presence in certain error responses indicates that the server wishes the client to use a long-term credential for authentication.

[14.8.](#) NONCE

The NONCE attribute may be present in requests and responses. It contains a sequence of qdtext or quoted-pair, which are defined in [RFC 3261](#) [[RFC3261](#)]. See [RFC 2617](#) [[RFC2617](#)], [Section 4.3](#), for guidance on selection of nonce values in a server. It MUST be less than 128 characters.

14.9. UNKNOWN-ATTRIBUTES

The UNKNOWN-ATTRIBUTES attribute is present only in an error response when the response code in the ERROR-CODE attribute is 420.

The attribute contains a list of 16 bit values, each of which represents an attribute type that was not understood by the server.

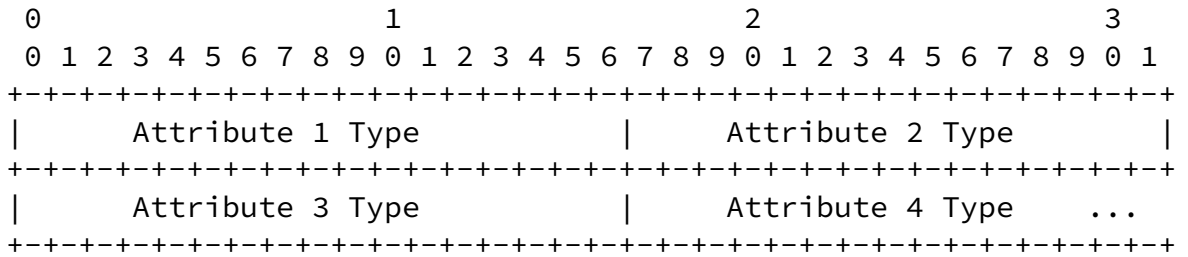


Figure 11: Format of UNKNOWN-ATTRIBUTES attribute

Note: In [RFC3489], this field was padded to 32 by duplicating the last attribute. In this version of the specification, the normal padding rules for attributes are used instead.

14.10. SERVER

The server attribute contains a textual description of the software being used by the server, including manufacturer and version number. The attribute has no impact on operation of the protocol, and serves only as a tool for diagnostic and debugging purposes. The value of SERVER is variable length. It MUST be a UTF-8 encoded sequence of less than 128 characters

14.11. ALTERNATE-SERVER

The alternate server represents an alternate transport address identifying a different STUN server which the STUN client should try.

It is encoded in the same way as MAPPED-ADDRESS, and thus refers to a single server by IP address. The IP address family **MUST** be identical to that of the source IP address of the request.

This attribute MUST only appear in an error response that contains a MESSAGE-INTEGRITY attribute. This prevents it from being used in denial-of-service attacks.

[15. Security Considerations](#)

[15.1. Attacks against the Protocol](#)

[15.1.1. Outside Attacks](#)

An attacker can try to modify STUN messages in transit, in order to cause a failure in STUN operation. These attacks are prevented for both requests and responses through the message integrity mechanism, using either a short term or long term credential.

An attacker that can observe, but not modify STUN messages in-transit (for example, an attacker present on a shared access medium, such as Wi-Fi), can see a STUN request, and then immediately send a STUN response, typically an error response, in order to disrupt STUN processing. This attack is also prevented for messages that utilize MESSAGE-INTEGRITY. However, some error responses, those related to authentication in particular, cannot be protected by MESSAGE-INTEGRITY. When STUN itself is run over a secure transport protocol (e.g., TLS), these attacks are completely mitigated.

[15.1.2. Inside Attacks](#)

A rogue client may try to launch a DoS attack against a server by sending it a large number of STUN requests. Fortunately, STUN requests can be processed statelessly by a server, making such attacks hard to launch.

[15.2. Attacks Affecting the Usage](#)

This section lists attacks that might be launched against a usage of STUN. Each STUN usage must consider whether these attacks are applicable to it, and if so, discuss counter-measures.

modifying the reflexive address learned by a STUN client through a Binding Request/Binding Response transaction. Since the usage of the reflexive address is a function of the usage, the applicability and remediation of these attacks is usage-specific. In common situations, modification of the reflexive address by an on-path attacker is easy to do. Consider, for example, the common situation where STUN is run directly over UDP. In this case, an on-path attacker can modify the source IP address of the Binding Request before it arrives at the STUN server. The STUN server will then return this IP address in the XOR-MAPPED-ADDRESS attribute to the client. Protecting against this attack by using a message-integrity check is impossible, since a message-integrity value cannot cover the source IP address, since the intervening NAT must be able to modify this value. Instead, one solution to preventing the attacks listed below is for the client to verify the reflexive address learned, as is done in ICE [[I-D.ietf-mmusic-ice](#)]. Other usages may use other means to prevent these attacks.

[15.2.1.](#) Attack I: DDoS Against a Target

In this attack, the attacker provides one or more clients with the same faked reflexive address that points to the intended target. This will trick the STUN clients into thinking that their reflexive addresses are equal to that of the target. If the clients hand out that reflexive address in order to receive traffic on it (for example, in SIP messages), the traffic will instead be sent to the target. This attack can provide substantial amplification, especially when used with clients that are using STUN to enable multimedia applications.

[15.2.2.](#) Attack II: Silencing a Client

In this attack, the attacker provides a STUN client with a faked reflexive address. The reflexive address it provides is a transport address that routes to nowhere. As a result, the client won't receive any of the packets it expects to receive when it hands out the reflexive address. This exploitation is not very interesting for the attacker. It impacts a single client, which is frequently not the desired target. Moreover, any attacker that can mount the attack could also deny service to the client by other means, such as preventing the client from receiving any response from the STUN server, or even a DHCP server.

[15.2.3.](#) Attack III: Assuming the Identity of a Client

This attack is similar to attack III. However, the faked reflexive address points to the attacker itself. This allows the attacker to

receive traffic which was destined for the client.

[15.2.4.](#) Attack IV: Eavesdropping

In this attack, the attacker forces the client to use a reflexive address that routes to itself. It then forwards any packets it receives to the client. This attack would allow the attacker to observe all packets sent to the client. However, in order to launch the attack, the attacker must have already been able to observe packets from the client to the STUN server. In most cases (such as when the attack is launched from an access network), this means that the attacker could already observe packets sent to the client. This attack is, as a result, only useful for observing traffic by attackers on the path from the client to the STUN server, but not generally on the path of packets being routed towards the client.

[15.3.](#) Hash Agility Plan

This specification uses SHA-1 for computation of the message integrity. If, at a later time, SHA-1 is found to be compromised, the following is the remedy that will be applied.

We will define a STUN extension which introduces a new message integrity attribute, computed using a new hash. Clients would be required to include both the new and old message integrity attributes in their requests or indications. A new server will utilize the new message integrity attribute, and an old one, the old. After a transition period where mixed implementations are in deployment, the old message-integrity attribute will be deprecated by another specification, and clients will cease including it in requests.

[16.](#) IAB Considerations

The IAB has studied the problem of "Unilateral Self Address Fixing" (UNSAF), 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](#) [[RFC3424](#)]). STUN can be used to perform this function using a BindingRequest/BindingResponse transaction if one agent is behind a NAT and the other is on the public side of the NAT.

The IAB has mandated that protocols developed for this purpose document a specific set of considerations. Because some STUN usages provide UNSAF functions (such as ICE [[I-D.ietf-mmusic-ice](#)]), and others do not (such as SIP Outbound [[I-D.ietf-sip-outbound](#)]), answers

to these considerations need to be addressed by the usages themselves.

[17.](#) IANA Considerations

IANA is hereby requested to create three new registries: a STUN methods registry, a STUN Attributes registry, and a STUN Error Codes registry.

[17.1.](#) STUN Methods Registry

A STUN method is a hex number in the range 0x000 - 0x3FF. The encoding of STUN method into a STUN message is described in [Section 6](#).

The initial STUN methods are:

- 0x000: (Reserved)
- 0x001: Binding
- 0x002: (Reserved; was SharedSecret)

STUN methods in the range 0x000 - 0x1FF are assigned by IETF Consensus [[RFC2434](#)]. STUN methods in the range 0x200 - 0x3FF are assigned on a First Come First Served basis [[RFC2434](#)]

[17.2.](#) STUN Attribute Registry

A STUN Attribute type is a hex number in the range 0x0000 - 0xFFFF. STUN attribute types in the range 0x0000 - 0x7FFF are considered comprehension-required; STUN attribute types in the range 0x8000 - 0xFFFF are considered comprehension-optional. A STUN agent handles unknown comprehension-required and comprehension-optional attributes differently.

The initial STUN Attributes types are:

Comprehension-required range (0x0000-0x7FFF):

- 0x0000: (Reserved)
- 0x0001: MAPPED-ADDRESS
- 0x0006: USERNAME
- 0x0007: (Reserved; was PASSWORD)
- 0x0008: MESSAGE-INTEGRITY
- 0x0009: ERROR-CODE
- 0x000A: UNKNOWN-ATTRIBUTES
- 0x0014: REALM
- 0x0015: NONCE
- 0x0020: XOR-MAPPED-ADDRESS

Comprehension-optional range (0x8000-0xFFFF)

- 0x8022: SERVER
- 0x8023: ALTERNATE-SERVER
- 0x8028: FINGERPRINT

STUN Attribute types in the first half of the comprehension-required range (0x0000 - 0x3FFF) and in the first half of the comprehension-optional range (0x8000 - 0xBFFF) are assigned by IETF Consensus [[RFC2434](#)]. STUN Attribute types in the second half of the comprehension-required range (0x4000 - 0x7FFF) and in the second half of the comprehension-optional range (0xC000 - 0xFFFF) are assigned on a First Come First Served basis [[RFC2434](#)].

[17.3](#). STUN Error Code Registry

A STUN Error code is a number in the range 0 - 699. STUN error codes are accompanied by a textual reason phrase in UTF-8 which is intended only for human consumption and can be anything appropriate; this

document proposes only suggested values.

STUN error codes are consistent in codepoint assignments and semantics with SIP [[RFC3261](#)] and HTTP [[RFC2616](#)].

The initial values in this registry are given in [Section 14.6](#).

New STUN error codes are assigned on a Specification-Required basis [[RFC2434](#)]. The specification must carefully consider how clients that do not understand this error code will process it before granting the request. See the rules in [Section 7.3.4](#).

[18](#). Changes Since [RFC 3489](#)

This specification obsoletes [RFC3489](#) [[RFC3489](#)]. This specification

differs from [RFC3489](#) in the following ways:

- o Removed the notion that STUN is a complete NAT traversal solution. STUN is now a toolkit that can be used to produce a NAT traversal solution. As a consequence, changed the name of the protocol to Session Traversal Utilities for NAT.
- o Introduced the concept of STUN usages, and described what a usage of STUN must document.
- o Removed the usage of STUN for NAT type detection and binding lifetime discovery. These techniques have proven overly brittle due to wider variations in the types of NAT devices than described in this document. Removed the RESPONSE-ADDRESS, CHANGED-ADDRESS, CHANGE-REQUEST, SOURCE-ADDRESS, and REFLECTED-FROM attributes.
- o Added a fixed 32-bit magic cookie and reduced length of transaction ID by 32 bits. The magic cookie begins at the same offset as the original transaction ID.
- o Added the XOR-MAPPED-ADDRESS attribute, which is included in Binding Responses if the magic cookie is present in the request. Otherwise the [RFC3489](#) behavior is retained (that is, Binding Response includes MAPPED-ADDRESS). See discussion in XOR-MAPPED-ADDRESS regarding this change.

- o Introduced formal structure into the Message Type header field, with an explicit pair of bits for indication of request, response, error response or indication. Consequently, the message type field is split into the class (one of the previous four) and method.
- o Explicitly point out that the most significant two bits of STUN are 0b00, allowing easy differentiation with RTP packets when used with ICE.
- o Added the FINGERPRINT attribute to provide a method of definitely detecting the difference between STUN and another protocol when the two protocols are multiplexed together.
- o Added support for IPv6. Made it clear that an IPv4 client could get a v6 mapped address, and vice-a-versa.
- o Added long-term credential-based authentication.
- o Added the SERVER, REALM, NONCE, and ALTERNATE-SERVER attributes.

- o Removed the SharedSecret method, and thus the PASSWORD attribute. This method was almost never implemented and is not needed with current usages.
- o Removed recommendation to continue listening for STUN Responses for 10 seconds in an attempt to recognize an attack.
- o Changed transaction timers to be more TCP friendly.
- o Removed the STUN example that centered around the separation of the control and media planes. Instead, provided more information on using STUN with protocols.
- o Defined a generic padding mechanism that changes the interpretation of the length attribute. This would, in theory, break backwards compatibility. However, the mechanism in [RFC 3489](#) never worked for the few attributes that weren't aligned naturally on 32 bit boundaries.

- o REALM, USERNAME, SERVER, reason phrases and NONCE limited to 127 characters.

19. Acknowledgements

The authors would like to thank Cedric Aoun, Pete Cordell, Cullen Jennings, Bob Penfield, Xavier Marjou, Bruce Lowekamp and Chris Sullivan for their comments, and Baruch Sterman and Alan Hawrylyshen for initial implementations. Thanks for Leslie Daigle, Allison Mankin, Eric Rescorla, and Henning Schulzrinne for IESG and IAB input on this work.

20. References

20.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), February 2000.
- [RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), May 2000.

Rosenberg, et al. Expires January 9, 2008 [Page 42]

Internet-Draft STUN July 2007

- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", [RFC 2617](#), June 1999.
- [RFC2988] Paxson, V. and M. Allman, "Computing TCP's Retransmission Timer", [RFC 2988](#), November 2000.
- [ITU.V42.1994] International Telecommunications Union, "Error-correcting Procedures for DCEs Using Asynchronous-to-Synchronous

Conversion", ITU-T Recommendation V.42, 1994.

20.2. Informational References

- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [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.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [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-16](#) (work in progress), June 2007.
- [RFC3489] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", [RFC 3489](#), March 2003.
- [I-D.ietf-behave-turn]
Rosenberg, J., "Obtaining Relay Addresses from Simple Traversal Underneath NAT (STUN)", [draft-ietf-behave-turn-03](#) (work in progress), March 2007.
- [I-D.ietf-sip-outbound]
Jennings, C. and R. Mahy, "Managing Client Initiated Connections in the Session Initiation Protocol (SIP)", [draft-ietf-sip-outbound-09](#) (work in progress), June 2007.

Rosenberg, et al. Expires January 9, 2008 [Page 43]

Internet-Draft

STUN

July 2007

- [I-D.ietf-behave-nat-behavior-discovery]
MacDonald, D. and B. Lowekamp, "NAT Behavior Discovery Using STUN", [draft-ietf-behave-nat-behavior-discovery-00](#) (work in progress), February 2007.

- [I-D.ietf-mmusic-ice-tcp]
Rosenberg, J., "TCP Candidates with Interactive Connectivity Establishment (ICE",
[draft-ietf-mmusic-ice-tcp-03](#) (work in progress),
March 2007.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#),
June 2002.
- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", [RFC 3424](#), November 2002.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#),
October 1998.

[Appendix A](#). C Snippet to Determine STUN Message Types

Given an 16-bit STUN message type value in host byte order in msg_type parameter, below are C macros to determine the STUN message types:

```
#define IS_REQUEST(msg_type)      (((msg_type) & 0x0110) == 0x0000)
#define IS_INDICATION(msg_type)  (((msg_type) & 0x0110) == 0x0010)
#define IS_SUCCESS_RESP(msg_type) (((msg_type) & 0x0110) == 0x0100)
#define IS_ERR_RESP(msg_type)    (((msg_type) & 0x0110) == 0x0110)
```

Authors' Addresses

Jonathan Rosenberg
Cisco
Edison, NJ
US

Email: jdrosen@cisco.com
URI: <http://www.jdrosen.net>

Christian Huitema
Microsoft
One Microsoft Way
Redmond, WA 98052
US

Email: huitema@microsoft.com

Rohan Mahy
Plantronics
345 Encinal Street
Santa Cruz, CA 95060
US

Email: rohan@ekabal.com

Philip Matthews
Avaya
1135 Innovation Drive
Ottawa, Ontario K2K 3G7
Canada

Phone: +1 613 592 4343 x224
Fax:
Email: philip_matthews@magma.ca
URI:

Dan Wing
Cisco
771 Alder Drive
San Jose, CA 95035
US

Email: dwing@cisco.com

Internet-Draft

STUN

July 2007

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

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