

Mumble protocol 1.2.X reference (WIP)

Stefan Hacker, Mikko Rantanen

June 10, 2013

DISCLAIMER

THIS DOCUMENTATION IS PROVIDED BY THE MUMBLE PROJECT "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE MUMBLE PROJECT BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS DOCUMENTATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Contents

1	Introduction	4
2	Overview	4
3	Protocol stack (TCP)	5
4	Establishing a connection	6
4.1	Connect	6
4.2	Version exchange	7
4.3	Authenticate	8
4.4	Crypt setup	8
4.5	Channel states	9
4.6	User states	9
4.7	Server sync	10
4.8	Ping	10
5	Voice data	10
5.1	Enabling the UDP channel	10
5.2	Data	11
5.3	Codecs	12
5.3.1	Whispering	12
5.4	varint, 64-bit integer encoding	13
5.5	TCP tunnel	13
5.6	Encryption	13
5.7	Implementation notes	14
6	Messages	14
6.1	ACL	14
6.1.1	ACL.ChanACL	14
6.1.2	ACL.ChanGroup	15
6.2	Authenticate	15
6.3	BanList	15
6.3.1	BanList.BanEntry	16
6.4	ChannelRemove	16

6.5	ChannelState	16
6.6	CodecVersion	17
6.7	ContextAction	17
6.8	ContextActionAdd	17
6.8.1	Enumeration: ContextActionAdd_Context	17
6.9	CryptSetup	17
6.10	PermissionDenied	18
6.10.1	Enumeration: PermissionDenied_DenyType	18
6.11	PermissionQuery	18
6.12	Ping	19
6.13	QueryUsers	19
6.14	Reject	19
6.14.1	Enumeration: Reject_RejectType	20
6.15	RequestBlob	20
6.16	ServerConfig	20
6.17	ServerSync	21
6.18	SuggestConfig	21
6.19	TextMessage	21
6.20	UDPTunnel	22
6.21	UserList	22
6.21.1	UserList_User	22
6.22	UserRemove	22
6.23	UserState	23
6.24	UserStats	24
6.24.1	UserStats_Stats	24
6.25	Version	25
6.26	VoiceTarget	25
6.26.1	VoiceTarget_Target	25
A	Appendix	26
A.1	Mumble.proto	26

1 Introduction

This document is meant to be a reference for the Mumble VoIP 1.2.X server-client communication protocol. It reflects the state of the protocol implemented in the Mumble 1.2.4 client and might be outdated by the time you are reading this. Be sure to check for newer revisions of this document on our website <http://www.mumble.info>. At the moment this document is work in progress.

2 Overview

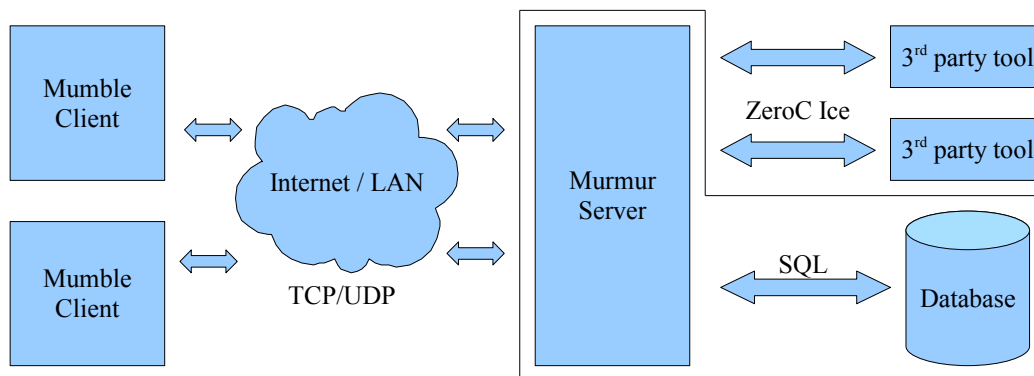


Figure 1: Mumble system overview

Mumble is based on a standard server-client communication model. It utilizes two channels of communication, the first one is a TCP connection which is used to reliably transfer control data between the client and the server. The second one is a UDP connection which is used for unreliable, low latency transfer of voice data.




Figure 2: Mumble crypto types

Both are protected by strong cryptography, this encryption is mandatory and cannot be disabled. The TCP control channel uses TLSv1 AES256-SHA¹ while the voice channel is encrypted with OCB-AES128².

While the TCP connection is mandatory the UDP connection can be compensated by tunnelling the UDP packets through the TCP connection as described in the protocol description later.

¹http://en.wikipedia.org/wiki/Transport_Layer_Security

²<http://www.cs.ucdavis.edu/~rogaway/ocb/ocb-back.htm>

3 Protocol stack (TCP)




Figure 3: Mumble protocol stack

Mumble has a shallow and easy to understand stack. Basically it uses Google's Protocol Buffers³ with simple prefixing to distinguish the different kinds of packets sent through an TLSv1 encrypted connection. This makes the protocol very easily expandable.




Figure 4: Mumble packet

The prefix consists out of the two bytes defining the type of the packet in the payload and 4 bytes stating the length of the payload in bytes followed by the payload itself. The following packet types are available in the current protocol and all but UDPTunnel are simple protobuf messages. If not mentioned otherwise all fields outside the protobuf encoding are big-endian.

³<http://code.google.com/p/protobuf/>

Table 1: Packet types

Type	Payload
0	Version
1	UDPTunnel
2	Authenticate
3	Ping
4	Reject
5	ServerSync
6	ChannelRemove
7	ChannelState
8	UserRemove
9	UserState
10	BanList
11	TextMessage
12	PermissionDenied
13	ACL
14	QueryUsers
15	CryptSetup
16	ContextActionModify
17	ContextAction
18	UserList
19	VoiceTarget
20	PermissionQuery
21	CodecVersion
22	UserStats
23	RequestBlob
24	ServerConfig
25	SuggestConfig

For raw representation of each packet type see the attached Mumble.proto file.

4 Establishing a connection

This section describes the communication between the server and the client during connection establishing, note that only the TCP connection needs to be established for the client to be connected. After this the client will be visible to the other clients on the server and able to send other types of messages.

4.1 Connect

As the basis for the synchronization procedure the client has to first establish the TCP connection to the server and do a common TLSv1 handshake. To be able to use the complete feature set of the Mumble protocol it is recommended that the client provides a strong certificate to the server. This however is not mandatory as you can connect to the server without providing a certificate. However the server must provide the client with its certificate and it is recommended that the client checks this.

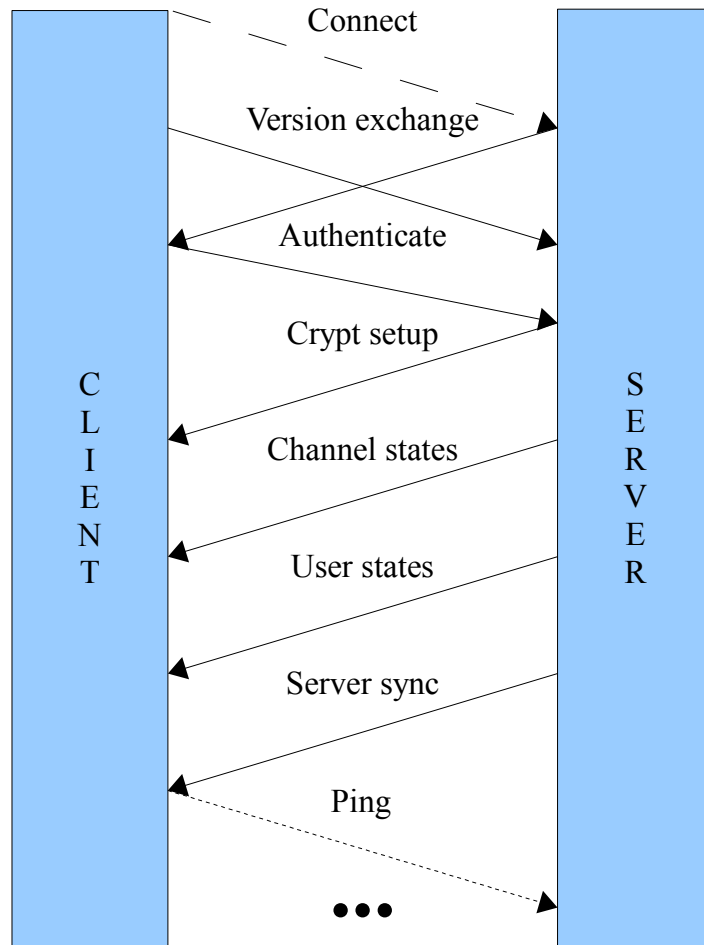


Figure 5: Mumble connection setup

4.2 Version exchange

Once the TLS handshake is completed both sides should transmit their version information using the Version message. The message structure is described below.

Version	
version	uint32
release	string
os	string
os_version	string

Figure 6: Version message

The version field is a combination of major, minor and patch version numbers (e.g. 1.2.0) so that major number takes two bytes and minor and patch numbers take one byte each. The structure is shown in figure 7. The release, os and os_version fields are common strings containing additional information.

Major	Minor	Patch
2B	1B	1B

Figure 7: `version` field structure

The version information may be used as part of the `SuggestConfig` checks, which usually refer to the standard client versions. The major changes between these versions are listed in table 2. The `release`, `os` and `os_version` information is not interpreted in any way at the moment.

Table 2: Mumble version differences

Version	Major changes
1.2.0	CELT 0.7.0 codec support
1.2.2	CELT 0.7.1 codec support
1.2.3	CELT 0.11.0 codec support, priority speakers
1.2.4	OPUS codec support, <code>SuggestConfig</code> message

4.3 Authenticate

Once the client has sent the version it should follow this with the `Authenticate` message. The message structure is described below in figure 8. This message may be sent immediately after sending the version message. The client does not need to wait for the server version message.

Authenticate	
username	string
password	string
tokens	repeated string

Figure 8: `Authenticate` message

The username and password are UTF-8 encoded strings. While the client is free to accept any username from the user the server is allowed to impose further restrictions. Furthermore if the client certificate has been registered with the server the client is primarily known with the username they had when the certificate was registered. For more information see the server documentation.

The password must only be provided if the server is passworded, the client provided no certificate but wants to authenticate to an account which has a password set, or to access the `SuperUser` account.

The third field contains a list of zero or more token strings which act as passwords that may give the client access to certain ACL groups without actually being a registered member in them, again see the server documentation for more information.

4.4 Crypt setup

Once the `Version` packets are exchanged the server will send a `CryptSetup` packet to the client. It contains the necessary cryptographic information for the `OCB-AES128`

encryption used in the UDP Voice channel. The packet is described in figure 9. The encryption itself is described later in section 5.6.

CryptSetup	
key	bytes
client_nonce	bytes
server_nonce	bytes

Figure 9: CryptSetup message

4.5 Channel states

After the client has successfully authenticated the server starts listing the channels by transmitting partial ChannelState message for every channel on this server. These messages lack the channel link information as the client does not yet have full picture of all the channels. Once the initial ChannelState has been transmitted for all channels the server updates the linked channels by sending new packets for these. The full structure of these ChannelState messages is shown in 10.

ChannelState	
channel_id	uint32
parent	uint32
name	string
links	uint32, repeated
description	string
links_add	uint32, repeated
links_remove	uint32, repeated
temporary	bool, optional
position	int32, optional

Figure 10: ChannelState message

The server must send a ChannelState for the root channel identified with ID 0.

4.6 User states

After the channels have been synchronized the server continues by listing the connected users. This is done by sending a UserState message for each user currently on the server, including the user that is currently connecting. The message structure is shown in figure 11.

ChannelState	
session	uint32
actor	uint32
name	string
user_id	uint32
channel_id	uint32
mute	bool
deaf	bool
suppress	bool
self_mute	bool
self_deaf	bool
texture	bytes
plugin_context	bytes
plugin_identity	string
comment	string
hash	string
comment_hash	bytes
texture_hash	bytes
priority_speaker	bool
recording	bool

Figure 11: UserState message

4.7 Server sync

The client has now received a copy of the parts of the server state he needs to know about. To complete the synchronization the server transmits a ServerSync message containing the session id of the clients session, the maximum bandwidth allowed on this server, the servers welcome text as well as the permissions the client has in the channel he ended up.

For more information please refer to Mumble.proto in the appendix.

4.8 Ping

If the client wishes to maintain the connection to the server it is required to ping the server. If the server does not receive a ping for 30 seconds it will disconnect the client.

5 Voice data

5.1 Enabling the UDP channel

Before the UDP channel can reliably be used both sides should be certain that the connection works. Before the server may use the UDP connection to the client the client must first open a UDP socket and communicate its address to the server by sending a packet over UDP. Once the server has received an UDP transmission the server should start using the UDP channel for the voice packets. Respectively the client should not

use the UDP channel for voice data until it is certain that the packets go through to the server.

In practice these requirements are filled with UDP ping. When the server receives a UDP ping packet (See figure 12) from the client it echoes the packet back. When the client receives this packet it can ascertain that the UDP channel works for two-way communication.

byte	:	type/flags	0010 0000 for Ping
varint	:	timestamp	Timestamp for the client.

Figure 12: UDP Ping packet

If the client stops receiving replies to the UDP packets at some point or never receives the first one it should immediately start tunneling the voice communication through TCP as described in section 5.5. When the server receives a tunneled packet over the TCP connection it must also stop using the UDP for communication. The client may continue sending UDP ping packets over the UDP channel and the server must echo these if it receives them. If the client later receives these echoes it may switch back to the UDP channel for voice communication. When the server receives an UDP voice communication packet from the client it should stop tunneling the packets as well.

5.2 Data

The voice data is transmitted in variable length packets that consist of header portion, followed by repeated data segments and an optional position part. The full packet structure is shown in figure 13. The decrypted data should never be longer than 1020 bytes, this allows the use of 1024 byte UDP buffer even after the 4-byte encryption header is added to the packet during the encryption. The protocol transfers 64-bit integers using variable length encoding. This encoding is specified in section 5.4.

Header	byte	:	type/target	Bit 1-3: Type, Bit 4-8: Target
	varint	:	session	The session number of the source user (only from server)
	varint	:	sequence	
Audio Repeated	byte	:	header	Bit 1: Terminator, Bit 2-8: Data length
	byte[]	:	data	Encoded voice frames
Position Optional	float	:	Pos 1	Positional audio positions
	float	:	Pos 2	Uses PacketDataStream encoding
	float	:	Pos 3	

Figure 13: UDP Voice packet

The first byte of the header contains the packet type and additional target specifier. The format of this byte is described below. If the voice packet comes from the server, the type is followed by a `varint` encoded value that specifies the session this voice packet originated from – this information is added by the server and the client omits this field. The last segment in the header is a sequence number for the first audio frame of the

packet. If there are for example two frames in the packet, the sequence field of the next packet should be incremented by two.

The type is stored in the first three bits and specifies the type and encoding of the packet. Current types are listed in table 3. The remaining 5 bits specify additional packet-wide options. For voice packets the values specify the voice target as listed in table 4.

Table 3: UDP Types

Type	Description
0	CELT Alpha encoded voice data
1	Ping packet (See section 5.1)
2	Speex encoded voice data
3	CELT Beta encoded voice data
4-7	Unused

Table 4: UDP targets

Target	Description
0	Normal talking
1	Whisper to channel
2-30	Direct whisper Always 2 for incoming whisper.
31	Server loopback

The audio frames consist of one byte long header and up to 127 bytes long data portion. The first bit in the header is the **Terminator bit** which informs the receiver whether there are more audio frames after this one. This bit is turned on (value 1) for all but the last frame in the current UDP packet. Rest of the seven bits in the header specify the length of the data portion. The data portion is encoded using one of the supported codecs. The exact codec is specified in the type portion of the whole packet (See table 3). *The data in each frame is encoded separately.*

5.3 Codecs

Mumble supports two distinct codecs; Low bit rate audio uses Speex and higher quality audio is encoded with CELT. Both of these codecs must be supported for full support of the Mumble protocol. Furthermore, as the CELT bitstream has not been frozen yet which places requirements for the exact CELT version: The clients must support CELT 0.7.1 bitstream. The protocol includes codec negotiation which allows clients to support other codec versions as well, in which case the server should attempt to negotiate a version that all clients support. The clients must respect the server resolution.

5.3.1 Whispering

Normal talking can be heard by the users of the current channel and all linked channels as long as the speaker has Talk permission on these channels. If the speaker wishes to broadcast the voice to specific users or channels, he may use whispering. This is achieved by registering a voice target using the VoiceTarget message (See 6.26) and specifying the target ID as the target in the first byte of the UDP packet.

5.4 varint, 64-bit integer encoding

The variable length integer encoding is used to encode long, 64-bit, integers so that short values do not need the full 8 bytes to be transferred. The basic idea behind the encoding is prefixing the value with a length prefix and then removing the leading zeroes from the value. The positive numbers are always right justified. That is to say that the least significant bit in the encoded presentation matches the least significant bit in the decoded presentation. Table 5 contains the definitions of the different length prefixes. The encoded `x` bits are part of the decoded number while the `_` signifies a unused bit. Encoding should be done by searching the first decoded description that fits the number that should be decoded, truncating it to the required bytes and combining it with the defined encoding prefix.

See the `quint64` shift operators in <https://github.com/mumble-voip/mumble/blob/master/src/PacketDataStream.h> for a reference implementation.

Table 5: varint prefixes

Encoded	Decoded
<code>0xxxxxxx</code>	1 byte with $7 \cdot 8 + 1$ leading zeroes
<code>10xxxxxx + 1 byte</code>	2 bytes with $6 \cdot 8 + 2$ leading zeroes
<code>110xxxxx + 2 bytes</code>	3 bytes with $5 \cdot 8 + 3$ leading zeroes
<code>1110xxxx + 3 bytes</code>	4 bytes with $4 \cdot 8 + 4$ leading zeroes
<code>111100__ + int (4 bytes)</code>	32-bit positive number
<code>111101__ + long (8 bytes)</code>	64-bit number
<code>111110__ + varint</code>	Negative varint
<code>111111xx</code>	Byte-inverted negative two byte number (<code>xx</code>)

5.5 TCP tunnel

If the UDP channel isn't available the voice packets must be transmitted through the TCP socket. These messages use the normal TCP prefixing shown in figure : 16-bit message type followed by 32-bit message length. However unlike other TCP messages, the UDP packets are not encoded as protocol buffer messages but instead the raw UDP packet described in chapter should be written to the TCP socket directly.

When the packets are received it is safe to parse the type and length fields normally. If the type matches that of the UDP tunnel the rest of the message should be processed as an UDP packet without attempting a protocol buffer decoding.

5.6 Encryption

All the packets are encrypted once during transfer. The actual encryption depends on the used transport layer. If the packets are tunneled through TCP they are encrypted using the TLS that encrypts the whole TCP connection and if they are sent directly using UDP they must be encrypted using the OCB-AES128 encryption. The OCB-AES128 encryption is described in section ??.

5.7 Implementation notes

When implementing the protocol it is easier to ignore the UDP transfer layer at first and just tunnel the UDP data through the TCP tunnel. The TCP layer must be implemented for authentication in any case. Making sure that the voice transmission works before implementing the UDP protocol simplifies debugging greatly. The UDP protocol is a required part of the specification though.

6 Messages

6.1 ACL

Field	Type	Rule	Description
<code>channel_id</code>	<code>uint32</code>	Req.	Channel ID of the channel this message affects
<code>inherit_acls</code>	<code>bool</code>	Opt.	True if the channel inherits its parent's ACLs, Default: true
<code>groups</code>	<code>ChanGroup</code>	Rep.	User group specifications
<code>acls</code>	<code>ChanACL</code>	Rep.	ACL specifications
<code>query</code>	<code>bool</code>	Opt.	True if the message is a query for ACLs instead of setting them, Default: false

6.1.1 ACL_ChanACL

Field	Type	Rule	Description
<code>apply_here</code>	<code>bool</code>	Opt.	True if this ACL applies to the current channel, Default: true
<code>apply_subs</code>	<code>bool</code>	Opt.	True if this ACL applies to the sub channels, Default: true
<code>inherited</code>	<code>bool</code>	Opt.	True if the ACL has been inherited from the parent, Default: true
<code>user_id</code>	<code>uint32</code>	Opt.	ID of the user that is affected by this ACL
<code>group</code>	<code>string</code>	Opt.	ID of the group that is affected by this group
<code>grant</code>	<code>uint32</code>	Opt.	Bit flag field of the permissions granted by this ACL
<code>deny</code>	<code>uint32</code>	Opt.	Bit flag field of the permissions denied by this ACL

6.1.2 ACL_ChanGroup

Field	Type	Rule	Description
<code>name</code>	<code>string</code>	Req.	Name of the channel group, UTF-8 encoded
<code>inherited</code>	<code>bool</code>	Opt.	True if the group has been inherited from the parent. Read only, Default: true
<code>inherit</code>	<code>bool</code>	Opt.	True if the group members are inherited, Default: true
<code>inheritable</code>	<code>bool</code>	Opt.	True if the group can be inherited by sub channels, Default: true
<code>add</code>	<code>uint32</code>	Rep.	Users explicitly included in this group, identified by <code>user_id</code>
<code>remove</code>	<code>uint32</code>	Rep.	Users explicitly removed from this group in this channel if the group has been inherited, identified by <code>user_id</code>
<code>inherited_members</code>	<code>uint32</code>	Rep.	Users inherited, identified by <code>user_id</code>

6.2 Authenticate

Used by the client to send the authentication credentials to the server.

Field	Type	Rule	Description
<code>username</code>	<code>string</code>	Opt.	UTF-8 encoded username
<code>password</code>	<code>string</code>	Opt.	Server or user password
<code>tokens</code>	<code>string</code>	Rep.	Additional access tokens for server ACL groups
<code>celt_versions</code>	<code>int32</code>	Rep.	A list of CELT bitstream version constants supported by the client.

6.3 BanList

Relays information on the bans. The client may send the BanList message to either modify the list of bans or query them from the server. The server sends this list only after a client queries for it.

Field	Type	Rule	Description
<code>bans</code>	<code>BanEntry</code>	Rep.	List of ban entries currently in place
<code>query</code>	<code>bool</code>	Opt.	True if the server should return the list, False if it should replace old ban list with this one, Default: false

6.3.1 BanList_BanEntry

Field	Type	Rule	Description
<code>address</code>	<code>bytes</code>	Req.	Banned IP address
<code>mask</code>	<code>uint32</code>	Req.	The length of the subnet mask for the ban
<code>name</code>	<code>string</code>	Opt.	User name for identification purposes, does not affect the ban
<code>hash</code>	<code>string</code>	Opt.	TODO ??
<code>reason</code>	<code>string</code>	Opt.	Reason for the ban, does not affect the ban
<code>start</code>	<code>string</code>	Opt.	Ban start time
<code>duration</code>	<code>uint32</code>	Opt.	Ban duration in seconds

6.4 ChannelRemove

Sent by the client when it wants a channel removed. Sent by the server when a channel has been removed and clients should be notified.

Field	Type	Rule	Description
<code>channel_id</code>	<code>uint32</code>	Req.	The <code>channel_id</code> of the channel to be removed

6.5 ChannelState

Used to communicate channel properties between the client and the server. Sent by the server during the login process (See 4.5) or when channel properties are updated. Client may use this message to update said channel properties.

Field	Type	Rule	Description
<code>channel_id</code>	<code>uint32</code>	Opt.	Unique ID for the channel within the server.
<code>parent</code>	<code>uint32</code>	Opt.	<code>channel_id</code> of the parent channel.
<code>name</code>	<code>string</code>	Opt.	Channel name, UTF-8 encoded.
<code>links</code>	<code>uint32</code>	Rep.	A collection of <code>channel_id</code> values of the linked channels. Absent during the first channel listing (See 4.5).
<code>description</code>	<code>string</code>	Opt.	Channel description, UTF-8 encoded. Only if the description is less than 128 bytes
<code>links_add</code>	<code>uint32</code>	Rep.	A collection of <code>channel_id</code> values that should be added to <code>links</code> .
<code>links_remove</code>	<code>uint32</code>	Rep.	A collection of <code>channel_id</code> values that should be removed from <code>links</code> .
<code>temporary</code>	<code>bool</code>	Opt.	True if the channel is temporary. Default: false
<code>position</code>	<code>uint32</code>	Opt.	Position weight to tweak the channel position in the channel list. Default: 0
<code>description_hash</code>	<code>bytes</code>	Opt.	SHA1 hash of the description if the description is 128 bytes or more. See 6.15

6.6 CodecVersion

Sent by the server to notify the users of the version of the CELT codec they should use. This may change during the connection when new users join.

Field	Type	Rule	Description
<code>alpha</code>	<code>int32</code>	Req.	The version of the CELT Alpha codec
<code>beta</code>	<code>int32</code>	Req.	The version of the CELT Beta codec
<code>prefer_alpha</code>	<code>bool</code>	Req.	True if the user should prefer Alpha over Beta, Default: true

6.7 ContextAction

Sent by the client when it wants to initiate a Context action. Refer to Mumble documentation (TODO Context action source) for more information.

Field	Type	Rule	Description
<code>session</code>	<code>uint32</code>	Opt.	The target User for the action, identified by <code>session</code>
<code>channel_id</code>	<code>uint32</code>	Opt.	The target Channel for the action, identified by <code>channel_id</code>
<code>action</code>	<code>string</code>	Req.	The action that should be executed

6.8 ContextActionAdd

Sent by the server to inform the client of available context actions.

Field	Type	Rule	Description
<code>action</code>	<code>string</code>	Req.	The action name
<code>text</code>	<code>string</code>	Req.	The display name of the action
<code>context</code>	<code>uint32</code>	Opt.	Context bit flags defining where the action should be displayed, see 6.8.1

6.8.1 Enumeration: ContextActionAdd.Context

Name	Value	Description
<code>Server</code>	<code>0x01</code>	Action is applicable to the server
<code>Channel</code>	<code>0x02</code>	Action can target a Channel
<code>User</code>	<code>0x04</code>	Action can target a User

6.9 CryptSetup

Used to initialize and resync the UDP encryption. See section ?? for more information. Either side may request a resync by sending the message without any values filled. The resync is performed by sending the message with only the client or server nonce filled.

Field	Type	Rule	Description
<code>key</code>	<code>bytes</code>	Opt.	Encryption key
<code>client_nonce</code>	<code>bytes</code>	Opt.	Client nonce
<code>server_nonce</code>	<code>bytes</code>	Opt.	Server nonce

6.10 PermissionDenied

Field	Type	Rule	Description
<code>permission</code>	<code>uint32</code>	Opt.	The denied permission when type is <code>Permission</code>
<code>channel_id</code>	<code>uint32</code>	Opt.	<code>channel_id</code> for the channel where the permission was denied when type is <code>Permission</code>
<code>session</code>	<code>uint32</code>	Opt.	The user who was denied permissions, identified by <code>session</code>
<code>reason</code>	<code>string</code>	Opt.	Textual reason for the denial
<code>type</code>	<code>DenyType</code>	Opt.	Type of the denial
<code>name</code>	<code>string</code>	Opt.	The name that is invalid when type is <code>UserName</code>

6.10.1 Enumeration: PermissionDenied_DenyType

Name	Value	Description
<code>Text</code>	0	Operation denied for other reason, see reason field
<code>Permission</code>	1	Permissions were denied
<code>SuperUser</code>	2	Cannot modify SuperUser
<code>ChannelName</code>	3	Invalid channel name
<code>TextTooLong</code>	4	Text message too long
<code>H9K</code>	5	The flux capacitor was spelled wrong.
<code>TemporaryChannel</code>	6	Operation not permitted in temporary channel
<code>MissingCertificate</code>	7	Operation requires certificate
<code>UserName</code>	8	Invalid username
<code>ChannelFull</code>	9	Channel is full

6.11 PermissionQuery

Sent by the client when it wants permissions for a certain channel. Sent by the server when it replies to the query or wants the user to resync all channel permissions.

Field	Type	Rule	Description
<code>channel_id</code>	<code>uint32</code>	Opt.	<code>channel_id</code> of the channel for which the permissions are queried
<code>permissions</code>	<code>uint32</code>	Opt.	Channel permissions. TODO: Encoded how?
<code>flush</code>	<code>bool</code>	Opt.	True if the client should drop its current permission information for all channels, Default: false

6.12 Ping

Sent by the client to notify the server that the client is still alive. Server must reply to the packet with the same timestamp and its own good/late/lost/resync numbers. None of the fields is strictly required.

Field	Type	Rule	Description
<code>timestamp</code>	<code>uint64</code>	Opt.	Client timestamp. Server should not attempt to decode.
<code>good</code>	<code>uint32</code>	Opt.	The amount of good packets received
<code>late</code>	<code>uint32</code>	Opt.	The amount of late packets received
<code>lost</code>	<code>uint32</code>	Opt.	The amount of packets never received
<code>resync</code>	<code>uint32</code>	Opt.	The amount of nonce resyncs
<code>udp_packets</code>	<code>uint32</code>	Opt.	The total amount of UDP packets received
<code>tcp_packets</code>	<code>uint32</code>	Opt.	The total amount of TCP packets received
<code>udp_ping_avg</code>	<code>float</code>	Opt.	UDP ping average
<code>udp_ping_var</code>	<code>float</code>	Opt.	UDP ping variance
<code>tcp_ping_avg</code>	<code>float</code>	Opt.	TCP ping average
<code>tcp_ping_var</code>	<code>float</code>	Opt.	TCP ping variance

6.13 QueryUsers

Client may use this message to refresh its registered user information. The client should fill the IDs or Names of the users it wants to refresh. The server fills the missing parts and sends the message back.

Field	Type	Rule	Description
<code>ids</code>	<code>uint32</code>	Rep.	User IDs
<code>names</code>	<code>string</code>	Rep.	User names in the same order as <code>ids</code>

6.14 Reject

Sent by the server when it rejects the user connection.

Field	Type	Rule	Description
<code>type</code>	<code>RejectType</code>	Opt.	Rejection type
<code>reason</code>	<code>string</code>	Opt.	Human readable rejection reason

6.14.1 Enumeration: Reject_RejectType

Name	Value	Description
<code>None</code>	0	TODO ??
<code>WrongVersion</code>	1	The client attempted to connect with an incompatible version
<code>InvalidUsername</code>	2	The user name supplied by the client was invalid
<code>WrongUserPW</code>	3	The client attempted to authenticate as a user with a password but it was wrong
<code>WrongServerPW</code>	4	The client attempted to connect to a passworded server but the password was wrong
<code>UsernameInUse</code>	5	Supplied username is already in use
<code>ServerFull</code>	6	Server is currently full and cannot accept more users
<code>NoCertificate</code>	7	The user did not provide a certificate but one is required

6.15 RequestBlob

Used by the client to request binary data from the server. By default large comments or textures are not sent within standard messages but instead the hash is. If the client does not recognize the hash it may request the resource when it needs it. The client does so by sending a RequestBlob message with the correct fields filled with the hashes it wants to receive. The server replies to this by sending a new UserState/ChannelState message with the resources filled even if they would normally be transmitted as hashes.

Field	Type	Rule	Description
<code>session_texture</code>	<code>uint32</code>	Rep.	Hashes of the requested UserState textures
<code>session_comment</code>	<code>uint32</code>	Rep.	Hashes of the requested UserState comments
<code>channel_description</code>	<code>uint32</code>	Rep.	Hashes of the requested ChannelState descriptions

6.16 ServerConfig

Sent by the server when it informs the clients on server configuration details.

Field	Type	Rule	Description
<code>max_bandwidth</code>	<code>uint32</code>	Opt.	The maximum bandwidth the clients should use
<code>welcome_text</code>	<code>string</code>	Opt.	Server welcome text
<code>allow_html</code>	<code>bool</code>	Opt.	True if the server allows HTML
<code>message_length</code>	<code>uint32</code>	Opt.	Maximum text message length
<code>image_message_length</code>	<code>uint32</code>	Opt.	Maximum image message length

6.17 ServerSync

ServerSync message is sent by the server when it has authenticated the user and finished synchronizing the server state. See section 4.7 for more information on the initial connection exchange.

Field	Type	Rule	Description
<code>session</code>	<code>uint32</code>	Opt.	The <code>session</code> of the current user
<code>max_bandwidth</code>	<code>uint32</code>	Opt.	Maximum bandwidth that the user should use
<code>welcome_text</code>	<code>string</code>	Opt.	Server welcome text
<code>permissions</code>	<code>uint64</code>	Opt.	Current user permissions TODO: Confirm??

6.18 SuggestConfig

Sent by the server to inform the clients of suggested client configuration specified by the server administrator.

Field	Type	Rule	Description
<code>version</code>	<code>uint32</code>	Opt.	Suggested client version
<code>positional</code>	<code>bool</code>	Opt.	True if the administrator suggests positional audio to be used on this server
<code>push_to_talk</code>	<code>bool</code>	Opt.	True if the administrator suggests push to talk to be used on this server

6.19 TextMessage

Used to send and broadcast text messages.

Field	Type	Rule	Description
<code>actor</code>	<code>uint32</code>	Opt.	The message sender, identified by its <code>session</code>
<code>session</code>	<code>uint32</code>	Rep.	Target users for the message, identified by their <code>session</code>
<code>channel_id</code>	<code>uint32</code>	Rep.	The channels to which the message is sent, identified by their <code>channel_ids</code>
<code>tree_id</code>	<code>uint32</code>	Rep.	The root channels when sending message recursively to several channels, identified by their <code>channel_ids</code>
<code>message</code>	<code>string</code>	Req.	The UTF-8 encoded message. May be HTML if the server allows.

6.20 UDPTunnel

Not used. Not even for tunneling UDP through TCP. See section 5.5 for more information.

Field	Type	Rule	Description
<code>packet</code>	<code>bytes</code>	Req.	Not used

6.21 UserList

Lists the registered users

Field	Type	Rule	Description
<code>users</code>	<code>User</code>	Rep.	A list of registered users

6.21.1 UserList_User

Field	Type	Rule	Description
<code>user_id</code>	<code>uint32</code>	Req.	Registered user ID
<code>name</code>	<code>string</code>	Opt.	Registered user name

6.22 UserRemove

Used to communicate user leaving or being kicked. May be sent by the client when it attempts to kick a user. Sent by the server when it informs the clients that a user is not present anymore.

Field	Type	Rule	Description
<code>session</code>	<code>uint32</code>	Req.	The user who is being kicked, identified by their <code>session</code> , not present when no one is being kicked
<code>actor</code>	<code>uint32</code>	Opt.	The user who initiated the removal. Either the user who performs the kick or the user who is currently leaving
<code>reason</code>	<code>string</code>	Opt.	Reason for the kick, stored as the ban reason if the user is banned
<code>ban</code>	<code>bool</code>	Opt.	True if the kick should result in a ban

6.23 UserState

Sent by the server when it communicates new and changed users to client. First seen during login procedure (See 4.6). May be sent by the client when it wishes to alter its state.

Field	Type	Rule	Description
<code>session</code>	<code>uint32</code>	Opt.	Unique user session ID of the user whose state this is, may change on reconnect
<code>actor</code>	<code>uint32</code>	Opt.	The <code>session</code> of the user who is updating this user
<code>name</code>	<code>string</code>	Opt.	User name, UTF-8 encoded
<code>user_id</code>	<code>uint32</code>	Opt.	Registered user ID if the user is registered
<code>channel_id</code>	<code>uint32</code>	Opt.	Channel on which the user is
<code>mute</code>	<code>bool</code>	Opt.	True if the user is muted by admin
<code>deaf</code>	<code>bool</code>	Opt.	True if the user is deafened by admin
<code>suppress</code>	<code>bool</code>	Opt.	True if the user has been suppressed from talking by a reason other than being muted
<code>self_mute</code>	<code>bool</code>	Opt.	True if the user has muted self
<code>self_deaf</code>	<code>bool</code>	Opt.	True if the user has deafened self
<code>texture</code>	<code>bytes</code>	Opt.	User image if it is less than 128 bytes
<code>plugin_context</code>	<code>bytes</code>	Opt.	TODO ??
<code>plugin_identity</code>	<code>string</code>	Opt.	TODO ??
<code>comment</code>	<code>string</code>	Opt.	User comment if it is less than 128 bytes
<code>hash</code>	<code>string</code>	Opt.	The hash of the user certificate
<code>comment_hash</code>	<code>bytes</code>	Opt.	SHA1 hash of the user comment if it 128 bytes or more. See 6.15
<code>texture_hash</code>	<code>bytes</code>	Opt.	SHA1 hash of the user picture if it 128 bytes or more. See 6.15
<code>priority_speaker</code>	<code>bool</code>	Opt.	True if the user is a priority speaker
<code>recording</code>	<code>bool</code>	Opt.	True if the user is currently recording

6.24 UserStats

Used to communicate user stats between the server and clients.

Field	Type	Rule	Description
<code>session</code>	<code>uint32</code>	Opt.	User whose stats these are
<code>stats_only</code>	<code>bool</code>	Opt.	True if the message contains only mutable stats (packets, ping), Default: false
<code>certificates</code>	<code>bytes</code>	Rep.	Full user certificate chain of the user certificate in DER format
<code>from_client</code>	<code>Stats</code>	Opt.	Packet statistics for packets received from the client
<code>from_server</code>	<code>Stats</code>	Opt.	Packet statistics for packets sent by the server
<code>udp_packets</code>	<code>uint32</code>	Opt.	Amount of UDP packets sent
<code>tcp_packets</code>	<code>uint32</code>	Opt.	Amount of TCP packets sent
<code>udp_ping_avg</code>	<code>float</code>	Opt.	UDP ping average
<code>udp_ping_var</code>	<code>float</code>	Opt.	UDP ping variance
<code>tcp_ping_avg</code>	<code>float</code>	Opt.	TCP ping average
<code>tcp_ping_var</code>	<code>float</code>	Opt.	TCP ping variance
<code>version</code>	<code>Version</code>	Opt.	Client version, see 6.25
<code>celt_versions</code>	<code>int32</code>	Rep.	A list of CELT bitstream version constants supported by the client of this user.
<code>address</code>	<code>bytes</code>	Opt.	Client IP address
<code>bandwidth</code>	<code>uint32</code>	Opt.	Bandwidth used by this client
<code>onlinesecs</code>	<code>uint32</code>	Opt.	Connection duration
<code>idlesecs</code>	<code>uint32</code>	Opt.	Duration since last activity
<code>strong_certificate</code>	<code>bool</code>	Opt.	True if the user has a strong certificate, Default: false

6.24.1 UserStats.Stats

Field	Type	Rule	Description
<code>good</code>	<code>uint32</code>	Opt.	The amount of good packets received
<code>late</code>	<code>uint32</code>	Opt.	The amount of late packets received
<code>lost</code>	<code>uint32</code>	Opt.	The amount of packets never received
<code>resync</code>	<code>uint32</code>	Opt.	The amount of nonce resyncs

6.25 Version

Field	Type	Rule	Description
version	uint32	Opt.	2-byte Major, 1-byte Minor and 1-byte Patch version number
release	string	Opt.	Client release name
os	string	Opt.	Client OS name
os.version	string	Opt.	Client OS version

6.26 VoiceTarget

Sent by the client when it wants to register or clear whisper targets. See 5.3.1 for more information. **Note: The first available target ID is 1 as 0 is reserved for normal talking. Maximum target ID is 30**

Field	Type	Rule	Description
id	uint32	Opt.	Voice target ID
targets	Target	Rep.	The receivers that this voice target includes

6.26.1 VoiceTarget.Target

Field	Type	Rule	Description
session	uint32	Rep.	Users that are included as targets
channel_id	uint32	Opt.	Channels that are included as targets
group	string	Opt.	TODO ??
links	bool	Opt.	True if the voice should follow links from the specified channel, Default: false
children	bool	Opt.	True if the voice should also be sent to children of the specified channel, Default: false

This document is WIP

SORRY BUT THIS DOCUMENT IS WORK IN PROGRESS. AT THE MOMENT IT LACKS A LOT OF IMPORTANT INFORMATION BUT WE HOPE TO BE ABLE TO FINISH THIS DOCUMENT SOMEDAY :-)

We're getting there though! Currently the largest omission is the UDP channel encryption. Most other bits are there.

A Appendix

A.1 Mumble.proto

```
1 package MumbleProto;
2
3 option optimize_for = SPEED;
4
5 message Version {
6     optional uint32 version = 1;
7     optional string release = 2;
8     optional string os = 3;
9     optional string os_version = 4;
10 }
11
12 message UDPTunnel {
13     required bytes packet = 1;
14 }
15
16 message Authenticate {
17     optional string username = 1;
18     optional string password = 2;
19     repeated string tokens = 3;
20     repeated int32 celt_versions = 4;
21     optional bool opus = 5 [default = false];
22 }
23
24 message Ping {
25     optional uint64 timestamp = 1;
26     optional uint32 good = 2;
27     optional uint32 late = 3;
28     optional uint32 lost = 4;
29     optional uint32 resync = 5;
30     optional uint32 udp_packets = 6;
31     optional uint32 tcp_packets = 7;
32     optional float udp_ping_avg = 8;
33     optional float udp_ping_var = 9;
34     optional float tcp_ping_avg = 10;
35     optional float tcp_ping_var = 11;
36 }
37
38 message Reject {
39     enum RejectType {
40         None = 0;
41         WrongVersion = 1;
42         InvalidUsername = 2;
43         WrongUserPW = 3;
44         WrongServerPW = 4;
45         UsernameInUse = 5;
46         ServerFull = 6;
47         NoCertificate = 7;
```

```

48             AuthenticatorFail = 8;
49         }
50         optional RejectType type = 1;
51         optional string reason = 2;
52     }
53
54     message ServerConfig {
55         optional uint32 max_bandwidth = 1;
56         optional string welcome_text = 2;
57         optional bool allow_html = 3;
58         optional uint32 message_length = 4;
59         optional uint32 image_message_length = 5;
60     }
61
62     message ServerSync {
63         optional uint32 session = 1;
64         optional uint32 max_bandwidth = 2;
65         optional string welcome_text = 3;
66         optional uint64 permissions = 4;
67     }
68
69     message ChannelRemove {
70         required uint32 channel_id = 1;
71     }
72
73     message ChannelState {
74         optional uint32 channel_id = 1;
75         optional uint32 parent = 2;
76         optional string name = 3;
77         repeated uint32 links = 4;
78         optional string description = 5;
79         repeated uint32 links_add = 6;
80         repeated uint32 links_remove = 7;
81         optional bool temporary = 8 [default = false];
82         optional int32 position = 9 [default = 0];
83         optional bytes description_hash = 10;
84     }
85
86     message UserRemove {
87         required uint32 session = 1;
88         optional uint32 actor = 2;
89         optional string reason = 3;
90         optional bool ban = 4;
91     }
92
93     message UserState {
94         optional uint32 session = 1;
95         optional uint32 actor = 2;
96         optional string name = 3;
97         optional uint32 user_id = 4;
98         optional uint32 channel_id = 5;

```

```

99         optional bool mute = 6;
100        optional bool deaf = 7;
101        optional bool suppress = 8;
102        optional bool self_mute = 9;
103        optional bool self_deaf = 10;
104        optional bytes texture = 11;
105        optional bytes plugin_context = 12;
106        optional string plugin_identity = 13;
107        optional string comment = 14;
108        optional string hash = 15;
109        optional bytes comment_hash = 16;
110        optional bytes texture_hash = 17;
111        optional bool priority_speaker = 18;
112        optional bool recording = 19;
113    }
114
115    message BanList {
116        message BanEntry {
117            required bytes address = 1;
118            required uint32 mask = 2;
119            optional string name = 3;
120            optional string hash = 4;
121            optional string reason = 5;
122            optional string start = 6;
123            optional uint32 duration = 7;
124        }
125        repeated BanEntry bans = 1;
126        optional bool query = 2 [default = false];
127    }
128
129    message TextMessage {
130        optional uint32 actor = 1;
131        repeated uint32 session = 2;
132        repeated uint32 channel_id = 3;
133        repeated uint32 tree_id = 4;
134        required string message = 5;
135    }
136
137    message PermissionDenied {
138        enum DenyType {
139            Text = 0;
140            Permission = 1;
141            SuperUser = 2;
142            ChannelName = 3;
143            TextTooLong = 4;
144            H9K = 5;
145            TemporaryChannel = 6;
146            MissingCertificate = 7;
147            UserName = 8;
148            ChannelFull = 9;
149            NestingLimit = 10;

```

```

150     }
151     optional uint32 permission = 1;
152     optional uint32 channel_id = 2;
153     optional uint32 session = 3;
154     optional string reason = 4;
155     optional DenyType type = 5;
156     optional string name = 6;
157 }
158
159 message ACL {
160     message ChanGroup {
161         required string name = 1;
162         optional bool inherited = 2 [default = true];
163         optional bool inherit = 3 [default = true];
164         optional bool inheritable = 4 [default = true];
165         repeated uint32 add = 5;
166         repeated uint32 remove = 6;
167         repeated uint32 inherited_members = 7;
168     }
169     message ChanACL {
170         optional bool apply_here = 1 [default = true];
171         optional bool apply_subs = 2 [default = true];
172         optional bool inherited = 3 [default = true];
173         optional uint32 user_id = 4;
174         optional string group = 5;
175         optional uint32 grant = 6;
176         optional uint32 deny = 7;
177     }
178     required uint32 channel_id = 1;
179     optional bool inherit_acls = 2 [default = true];
180     repeated ChanGroup groups = 3;
181     repeated ChanACL acls = 4;
182     optional bool query = 5 [default = false];
183 }
184
185 message QueryUsers {
186     repeated uint32 ids = 1;
187     repeated string names = 2;
188 }
189
190 message CryptSetup {
191     optional bytes key = 1;
192     optional bytes client_nonce = 2;
193     optional bytes server_nonce = 3;
194 }
195
196 message ContextActionModify {
197     enum Context {
198         Server = 0x01;
199         Channel = 0x02;
200         User = 0x04;

```

```

201     }
202     enum Operation {
203         Add = 0;
204         Remove = 1;
205     }
206     required string action = 1;
207     optional string text = 2;
208     optional uint32 context = 3;
209     optional Operation operation = 4;
210 }
211
212 message ContextAction {
213     optional uint32 session = 1;
214     optional uint32 channel_id = 2;
215     required string action = 3;
216 }
217
218 message UserList {
219     message User {
220         required uint32 user_id = 1;
221         optional string name = 2;
222     }
223     repeated User users = 1;
224 }
225
226 message VoiceTarget {
227     message Target {
228         repeated uint32 session = 1;
229         optional uint32 channel_id = 2;
230         optional string group = 3;
231         optional bool links = 4 [default = false];
232         optional bool children = 5 [default = false];
233     }
234     optional uint32 id = 1;
235     repeated Target targets = 2;
236 }
237
238 message PermissionQuery {
239     optional uint32 channel_id = 1;
240     optional uint32 permissions = 2;
241     optional bool flush = 3 [default = false];
242 }
243
244 message CodecVersion {
245     required int32 alpha = 1;
246     required int32 beta = 2;
247     required bool prefer_alpha = 3 [default = true];
248     optional bool opus = 4 [default = false];
249 }
250
251 message UserStats {

```

```

252     message Stats {
253         optional uint32 good = 1;
254         optional uint32 late = 2;
255         optional uint32 lost = 3;
256         optional uint32 resync = 4;
257     }
258
259     optional uint32 session = 1;
260     optional bool stats_only = 2 [default = false];
261     repeated bytes certificates = 3;
262     optional Stats from_client = 4;
263     optional Stats from_server = 5;
264
265     optional uint32 udp_packets = 6;
266     optional uint32 tcp_packets = 7;
267     optional float udp_ping_avg = 8;
268     optional float udp_ping_var = 9;
269     optional float tcp_ping_avg = 10;
270     optional float tcp_ping_var = 11;
271
272     optional Version version = 12;
273     repeated int32 celt_versions = 13;
274     optional bytes address = 14;
275     optional uint32 bandwidth = 15;
276     optional uint32 onlinesecs = 16;
277     optional uint32 idlesecs = 17;
278     optional bool strong_certificate = 18 [default = false];
279     optional bool opus = 19 [default = false];
280 }
281
282 message SuggestConfig {
283     optional uint32 version = 1;
284     optional bool positional = 2;
285     optional bool push_to_talk = 3;
286 }
287
288 message RequestBlob {
289     repeated uint32 session_texture = 1;
290     repeated uint32 session_comment = 2;
291     repeated uint32 channel_description = 3;
292 }

```