

Mumble protocol 1.2.X reference (WIP)

Stefan Hacker

January 24, 2010

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	3
2	Overview	3
3	Protocol stack (TCP)	4
4	Establishing a connection	5
4.1	Connect	5
4.2	Version exchange	6
4.3	Crypt setup	6
4.4	Authenticate	6
4.5	Channel states	7
4.6	User states	7
4.7	Server sync	7
4.8	PacketDataStream	7
5	This document is WIP	7
A	Appendix	8
A.1	Mumble.proto	8

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 It reflects the state of the protocol implemented in the Mumble 1.2.2 client and might be outdated by the time you are reading this.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

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

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.

3 Protocol stack (TCP)




Figure 3: Mumble protocol stack

Mumble has a shallow and easy to understand stack. Basically it uses Googles 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 are little-endian encoded.

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

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

³<http://code.google.com/p/protobuf/>

4 Establishing a connection

The following section is going to describe the communication between the server and the client during connection establishing, note that the first part of this section only contains the procedures for the TCP connection. After this the client will be visible to the other clients on the server and able to send other types of messages.

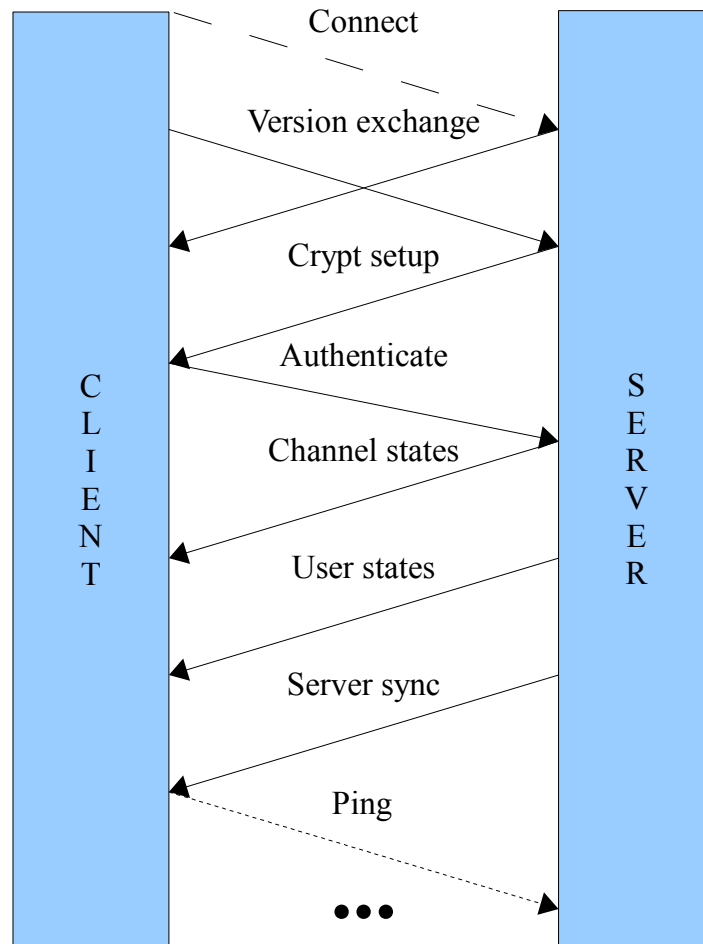


Figure 5: Mumble connection setup

4.1 Connect

As a basis for the synchronization procedure you first have to 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 your client provides a

strong certificate to the server. This however is not mandatory, you can connect to the server without providing a certificate, we do recommend to check the servers certificate though.

4.2 Version exchange

Once the TLS handshake is completed the server will send a Version packet to the client containing following information:

The client is supposed to send a Version packet with his information before any other messages. The version field of the packet contains the (major, minor, patch) tuple (e.g. 1.2.0) encoded like described in the following figure.

The release, os and os_version fields are common strings containing additional information about the client. This information is not interpreted in any way at the moment.

4.3 Crypt setup

Once the Version packets are exchanged the server will send a CryptSetup packet to the client. It contains the necessary cryptographic information to establish the OCB-AES128 encrypted UDP Voice channel. This will be described later in the section.

4.4 Authenticate

Before the client can be synchronized with the server state it has to authenticate itself to the server. This is done by sending an Authenticate packet.

The username and password are encoded as simple strings. Be aware that the server can impose restrictions on the username, also once the client registered a certificate with the server this field is only displayed in brackets behind the name the client possessed when he 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 called tokens contains a list of zero or more strings called tokens which are basically password which can give you access to a certain ACL group without actually being a registered member in them, again see the server documentation for more information.

4.5 Channel states

After the client is successfully authenticated the server starts synchronizing the state by transmitting a ChannelState message for every channel on this server. Note that these do not yet contain channel links. These are transmitted as updated directly after every channel has been transmitted. It contains the following information:

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

4.6 User states

When the channels are synchronized the server send a UserState message for every user connected to the client containing the following data:

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

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 PacketDataStream

The PacketDataStream class is used to serialize/deserialize the data packets received on the UDP connection or via the TCP-Tunneling. As the name implies it provides a stream based access to the data it contains. To pull data from it the user has to know what is located on the current position in the stream (e.g. a uint32, utf8 string and so on), the class itself is not aware of it's contents.

5 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 :-)

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 }
22
23 message Ping {
24     optional uint64 timestamp = 1;
25     optional uint32 good = 2;
26     optional uint32 late = 3;
27     optional uint32 lost = 4;
28     optional uint32 resync = 5;
29     optional uint32 udp_packets = 6;
30     optional uint32 tcp_packets = 7;
31     optional float udp_ping_avg = 8;
32     optional float udp_ping_var = 9;
33     optional float tcp_ping_avg = 10;
34     optional float tcp_ping_var = 11;
35 }
36
37 message Reject {
38     enum RejectType {
39         None = 0;
```



```

40         WrongVersion = 1;
41         InvalidUsername = 2;
42         WrongUserPW = 3;
43         WrongServerPW = 4;
44         UsernameInUse = 5;
45         ServerFull = 6;
46         NoCertificate = 7;
47     }
48     optional RejectType type = 1;
49     optional string reason = 2;
50 }
51
52 message ServerSync {
53     optional uint32 session = 1;
54     optional uint32 max_bandwidth = 2;
55     optional string welcome_text = 3;
56     optional uint64 permissions = 4;
57     optional bool allow_html = 5 [default = true];
58 }
59
60 message ChannelRemove {
61     required uint32 channel_id = 1;
62 }
63
64 message ChannelState {
65     optional uint32 channel_id = 1;
66     optional uint32 parent = 2;
67     optional string name = 3;
68     repeated uint32 links = 4;
69     optional string description = 5;
70     repeated uint32 links_add = 6;
71     repeated uint32 links_remove = 7;
72     optional bool temporary = 8 [default = false];
73     optional int32 position = 9 [default = 0];
74     optional bytes description_hash = 10;
75 }
76
77 message UserRemove {
78     required uint32 session = 1;
79     optional uint32 actor = 2;
80     optional string reason = 3;
81     optional bool ban = 4;
82 }

```

```

83
84 message UserState {
85     optional uint32 session = 1;
86     optional uint32 actor = 2;
87     optional string name = 3;
88     optional uint32 user_id = 4;
89     optional uint32 channel_id = 5;
90     optional bool mute = 6;
91     optional bool deaf = 7;
92     optional bool suppress = 8;
93     optional bool self_mute = 9;
94     optional bool self_deaf = 10;
95     optional bytes texture = 11;
96     optional bytes plugin_context = 12;
97     optional string plugin_identity = 13;
98     optional string comment = 14;
99     optional string hash = 15;
100    optional bytes comment_hash = 16;
101    optional bytes texture_hash = 17;
102 }
103
104 message BanList {
105     message BanEntry {
106         required bytes address = 1;
107         required uint32 mask = 2;
108         optional string name = 3;
109         optional string hash = 4;
110         optional string reason = 5;
111         optional string start = 6;
112         optional uint32 duration = 7;
113     }
114     repeated BanEntry bans = 1;
115     optional bool query = 2 [default = false];
116 }
117
118 message TextMessage {
119     optional uint32 actor = 1;
120     repeated uint32 session = 2;
121     repeated uint32 channel_id = 3;
122     repeated uint32 tree_id = 4;
123     required string message = 5;
124 }
125

```

```

126 message PermissionDenied {
127     enum DenyType {
128         Text = 0;
129         Permission = 1;
130         SuperUser = 2;
131         ChannelName = 3;
132         TextTooLong = 4;
133         H9K = 5;
134         TemporaryChannel = 6;
135         MissingCertificate = 7;
136         UserName = 8;
137         ChannelFull = 9;
138     }
139     optional uint32 permission = 1;
140     optional uint32 channel_id = 2;
141     optional uint32 session = 3;
142     optional string reason = 4;
143     optional DenyType type = 5;
144     optional string name = 6;
145 }
146
147 message ACL {
148     message ChanGroup {
149         required string name = 1;
150         optional bool inherited = 2 [default = true];
151         optional bool inherit = 3 [default = true];
152         optional bool inheritable = 4 [default = true];
153         repeated uint32 add = 5;
154         repeated uint32 remove = 6;
155         repeated uint32 inherited_members = 7;
156     }
157     message ChanACL {
158         optional bool apply_here = 1 [default = true];
159         optional bool apply_subs = 2 [default = true];
160         optional bool inherited = 3 [default = true];
161         optional uint32 user_id = 4;
162         optional string group = 5;
163         optional uint32 grant = 6;
164         optional uint32 deny = 7;
165     }
166     required uint32 channel_id = 1;
167     optional bool inherit_acls = 2 [default = true];
168     repeated ChanGroup groups = 3;

```

```

169         repeated ChanACL acls = 4;
170         optional bool query = 5 [default = false];
171     }
172
173     message QueryUsers {
174         repeated uint32 ids = 1;
175         repeated string names = 2;
176     }
177
178     message CryptSetup {
179         optional bytes key = 1;
180         optional bytes client_nonce = 2;
181         optional bytes server_nonce = 3;
182     }
183
184     message ContextActionAdd {
185         enum Context {
186             Server = 0x01;
187             Channel = 0x02;
188             User = 0x04;
189         }
190         required string action = 1;
191         required string text = 2;
192         optional uint32 context = 3;
193     }
194
195     message ContextAction {
196         optional uint32 session = 1;
197         optional uint32 channel_id = 2;
198         required string action = 3;
199     }
200
201     message UserList {
202         message User {
203             required uint32 user_id = 1;
204             optional string name = 2;
205         }
206         repeated User users = 1;
207     }
208
209     message VoiceTarget {
210         message Target {
211             repeated uint32 session = 1;

```

```

212         optional uint32 channel_id = 2;
213         optional string group = 3;
214         optional bool links = 4 [default = false];
215         optional bool children = 5 [default = false];
216     }
217     optional uint32 id = 1;
218     repeated Target targets = 2;
219 }
220
221 message PermissionQuery {
222     optional uint32 channel_id = 1;
223     optional uint32 permissions = 2;
224     optional bool flush = 3 [default = false];
225 }
226
227 message CodecVersion {
228     required int32 alpha = 1;
229     required int32 beta = 2;
230     required bool prefer_alpha = 3 [default = true];
231 }
232
233 message UserStats {
234     message Stats {
235         optional uint32 good = 1;
236         optional uint32 late = 2;
237         optional uint32 lost = 3;
238         optional uint32 resync = 4;
239     }
240
241     optional uint32 session = 1;
242     optional bool stats_only = 2 [default = false];
243     repeated bytes certificates = 3;
244     optional Stats from_client = 4;
245     optional Stats from_server = 5;
246
247     optional uint32 udp_packets = 6;
248     optional uint32 tcp_packets = 7;
249     optional float udp_ping_avg = 8;
250     optional float udp_ping_var = 9;
251     optional float tcp_ping_avg = 10;
252     optional float tcp_ping_var = 11;
253
254     optional Version version = 12;

```

```
255     repeated int32 celt_versions = 13;
256     optional bytes address = 14;
257     optional uint32 bandwidth = 15;
258     optional uint32 onlinesecs = 16;
259     optional uint32 idlesecs = 17;
260     optional bool strong_certificate = 18 [default = false];
261 }
262
263 message RequestBlob {
264     repeated uint32 session_texture = 1;
265     repeated uint32 session_comment = 2;
266     repeated uint32 channel_description = 3;
267 }
```