

Mumble Conference With ALSA

About

by [Daniel O'Neill](#)

How to set up a Mumble conference with ALSA.

Prerequisites

You'll need to verify you have a few things installed before you can get started:

```
* Linux 2.6+
* ALSA Loopback Module (snd-aloop or snd_aloop)
* Xorg, Xvfb, or Xvnc
* Mumble (1.2.2 used here)
```

For the FreeSWITCH, you'll need [mod_portaudio](#) and [mod_conference](#).

Verify that you have these modules by looking in mod/ for their filenames (mod_portaudio.so and mod_conference.so)

Also, make sure they're enabled in autoload_configs/modules.conf.xml.

Conference Setup

Initially all you'll need to do is setup a conference, which is configured in autoload_configs/conference.conf.xml:

Mine looks like this:

```

<configuration name="conference.conf" description="Audio Conference">
  <advertise>
    <room name="\*94@${domain}" status="FreeSWITCH"/>
  </advertise>
  <caller-controls>
    <group name="default">
      <control action="mute" digits="0"/>
      <control action="deaf mute" digits="*/>
      <control action="energy up" digits="9"/>
      <control action="energy equ" digits="8"/>
      <control action="energy dn" digits="7"/>
      <control action="vol talk up" digits="3"/>
      <control action="vol talk zero" digits="2"/>
      <control action="vol talk dn" digits="1"/>
      <control action="vol listen up" digits="6"/>
      <control action="vol listen zero" digits="5"/>
      <control action="vol listen dn" digits="4"/>
      <control action="hangup" digits="#" />
    </group>
  </caller-controls>
  <profiles>
    <profile name="default">
      <param name="domain" value="${domain}"/>
      <param name="rate" value="48000"/>
      <param name="interval" value="20"/>
      <param name="energy-level" value="300"/>
      <param name="caller-controls" value="default"/>
      <param name="sound-prefix" value="${sounds_dir}/en/us/callie"/>
      <param name="muted-sound" value="conference/conf-muted.wav"/>
      <param name="unmuted-sound" value="conference/conf-unmuted.wav"/>
      <param name="alone-sound" value="conference/conf-alone.wav"/>
      <param name="enter-sound" value="tone_stream://%(200,0,500,600,700)"/>
      <param name="exit-sound" value="tone_stream://%(500,0,300,200,100,50,25)"/>
      <param name="kicked-sound" value="conference/conf-kicked.wav"/>
      <param name="locked-sound" value="conference/conf-locked.wav"/>
      <param name="is-locked-sound" value="conference/conf-is-locked.wav"/>
      <param name="is-unlocked-sound" value="conference/conf-is-unlocked.wav"/>
      <param name="pin-sound" value="conference/conf-pin.wav"/>
      <param name="bad-pin-sound" value="conference/conf-bad-pin.wav"/>
      <param name="caller-id-name" value="${outbound_caller_name}"/>
      <param name="caller-id-number" value="${outbound_caller_id}"/>
      <param name="comfort-noise" value="true"/>
    </profile>
  </profiles>
</configuration>

```

There's probably 'problems' with this, but it works perfectly for me, so pundits please update this example if you spot any errors.

As you can see, it's a basic conference, but with some key differences: the RATE is quite high (48khz), but this is because Mumble (on my system) uses this samplerate. Using the same samplerate not only avoids resampling, but also improves stability (no jitters, or crackles). Take note of this setting though for later.

Dialplan Setup

Assign an extension to this conference, and test it out. My dialplan entry looks like this:

```

<extension name="mumbleconference">
  <condition field="destination_number" expression="^\*94$">
    <action application="conference" data="hydway"/>
  </condition>
</extension>

```

In our case, the conference is configured as 'default', and there is no conference named 'hydway', so .. it works. But I think if you have more than one conference, you should point your dialplan at it properly.

PortAudio Endpoint

The Portaudio plugin must also be setup, and the configuration is located at `autoload_configs/portaudio.conf.xml`. Mine looks like this:

```
<configuration name="portaudio.conf" description="Soundcard Endpoint">
  <settings>
    <param name="context" value="local"/>
    <param name="indev" value="cloop"/>
    <param name="outdev" value="ploop"/>
    <param name="dialplan" value="XML"/>
    <param name="cid-name" value="${outbound_caller_name}"/>
    <param name="cid-num" value="${outbound_caller_id}"/>
    <param name="sample-rate" value="48000"/>
    <param name="codec-ms" value="20"/>
  </settings>
</configuration>
```

Again, take note of the 'sample-rate' parameter, we want it to match both our conference and our mumble client.

ALSA asoundrc or asound.conf

Notice the 'indev' and 'outdev' values as 'cloop' and 'ploop' in the portaudio setup. These devices will be created using the `snd-aloop` ALSA module by adding (or creating) the following to `/etc/asound.conf`:

```
pcm.amix {
  type dmix
  ipc_key 219345
  slave.pcm "hw:Loopback,0,0"
}
pcm.asnoop {
  type dsnoop
  ipc_key 219346
  slave.pcm "hw:Loopback,0,1"
}
# -----
pcm.ploop {
  type dmix
  ipc_key 219356
  slave.pcm "hw:Loopback,1,1"
}
pcm.cloop {
  type dsnoop
  ipc_key 219348
  slave.pcm "hw:Loopback,1,0"
}
```

Effectively, this will connect 'asnoop' with 'ploop', and 'amix' with 'cloop'. When audio is piped into asnoop, it'll come out of ploop. The same goes for the other two. (You can name these whatever you want; `mumbleAplayback/mumbleAcapture` and `mumbleBplayback/mumbleBcapture` for example may make more sense.)

In this configuration, everything runs through `dmix` and `dsnoop`. This isn't required if your samplerates match. To freeball it, change the 'type' field to 'plug', and remove the `ipc_key` lines.

Preparing FreeSWITCH

In case your freeswitch is currently running, connect using `fs_cli` and issue:

```
reloadxml
reload mod_portaudio
reload mod_conference
```

Now, while still in `fs_cli`, list your portaudio devices to verify it is using our new loopback devices:

```
pa devlist
```

If configured correctly, it should look something like this:

```
freeswitch@internal> pa devlist
0;VIA 8237: VIA 8237 (hw:0,0)(ALSA);2;2;
1;VIA 8237: VIA 8237 (hw:0,1)(ALSA);2;6;
2;Dummy: Dummy PCM (hw:1,0)(ALSA);2;2;
3;Dummy: Dummy PCM (hw:1,1)(ALSA);2;2;
4;Dummy: Dummy PCM (hw:1,2)(ALSA);2;2;
5;Dummy: Dummy PCM (hw:1,3)(ALSA);2;2;
6;Loopback: Loopback PCM (hw:2,0)(ALSA);32;32;
7;Loopback: Loopback PCM (hw:2,1)(ALSA);2;32;
8;front(ALSA);0;2;
9;surround40(ALSA);0;4;
10;surround41(ALSA);0;128;
11;surround50(ALSA);0;128;
12;surround51(ALSA);0;128;
13;iec958(ALSA);0;2;
14;spdif(ALSA);0;2;
15;ploop(ALSA);0;2;r,o
16;cloop(ALSA);2;0;i
17;default(ALSA);128;128;
18;dmix(ALSA);0;2;
```

The entries to note are 'ploop' and 'cloop', entries 15 and 16. At the end of each line you'll see 'r,o' which means 'ringer, output', meaning this device is where your portaudio endpoint will both ring on and send output to. The next line has an 'i' at the end, meaning the endpoint will use this device as input.

The 'Loopback' entries (6 and 7) are possibly unusable. I haven't found any way to directly address the subdevices ('2,0,0', '2,0,1', etc.) Luckily, it picks up our handrolled PCMs, which is good enough.

Mumble Configuration

If you've made it this far, it gets much easier. Get/build/install Mumble. I'm not going to explain how, their website should help you out.

Once you have Mumble built and running on your server (using Xorg or Xvnc), open the settings via Configure->Settings, and check the Advanced box in the bottom left.

Select Audio Input on the left-side menu.

Your devices (amix/asnoop) probably won't appear in the dropdown. If they do, select them. If they don't, don't worry, we'll fix that in a minute.

Under Transmission, set Transmit to Voice Activity. Adjust your Compression settings (I always suggest using Speex codec for compatibility with all Mumble clients, which may be achieved by selecting a Quality level below 45.5kbit. If it stays on CELT, your Mumble client probably wasn't compiled with Speex. I -highly- suggest recompiling with support for both, unless you have control over -every- Mumble client which this configuration will communicate with; otherwise you may not 'hear' what they say). Tune your Audio per packet as you like as well. I've found that when speaking to a murmur server running on my LAN, this can be set to the minimum (10ms) without any issues. It also improves responsiveness to do this, but higher settings may be required if you're connecting to a server off-site.

Switch to Audio Output on the left-side menu.

Default Jitter Buffer may be set quite low, I use 10ms. Output Delay also. Again, for longer range Mumble links, higher levels may be required on these two fields.

Switch to Messages on the left-side menu.

These are at your own discretion. Basically every message which has an audio-generating alert associated with it will be broadcast over the conference. This bothered me, so I disabled all message notifications. (Text-to-speech is sorta cool, but believe me, it gets old really fast on a conference!)

You may want to finish tuning your setup, but for great victory, click OK to save, and close Mumble.

Mumble manual device configuration

If your 'amix' and 'asnoop' PCM devices didn't appear in your Audio Input and Audio Output device selection fields in Mumble, some configuration hacking is required. Open `~/config/Mumble/Mumble.conf` in your favourite text editor, and locate a section which looks like this:

```
[alsa]
output="front:CARD=Headset,DEV=0"
input="front:CARD=Headset,DEV=0"
```

... now, change it to look like this:

```
[alsa]
input=asnoop
output=amix
```

Save, close, restart Mumble. Easy, right? Sure, but the next time you open Configure->Settings and click Apply or OK, you'll have to do this all over again. In my case, the section itself was removed from the config file entirely, I just re-added it. Other than this inconvenience, it works perfectly. And once running, you should have little or no reason to reopen your configuration anyway, right? :)

Connecting it together

If you've managed all of this, connect to any server you like in Mumble. If anybody speaks to you on there, you won't be able to hear it (if you have speakers or a headset connected to your FreeSWITCH server at all, you weirdos) because Mumble is routing directly into our Loopback device, which is tied, on the opposing end, to our `mod_portaudio` instance.

Go back to `fs_cli` and issue the following command:

```
pa call *94
```

...with your conference extension in place of '*94'. If all is well, you should see the 'lips' of your FreeSWITCH-controlled-Mumble-client turn red with fabulous activity-- briefly.

In the terminal where you executed Mumble, note the sample rates it is using for both input and output. Configure your conference and/or portaudio to match, unless you're satisfied by `dsnoop` and `dmix`'s resampling.

Testing

Now to test both ends:

Connect to the same mumble server from another computer, as usual. Pick up your SIP phone (or whatever you have connecting to FreeSWITCH, softphone, iax phone, whatever.) and dial into your conference. Talk into each, and listen with the respective other. If you have latency, adjust buffering accordingly.

Deploying

At this point, you may wish to set `snd-aloop` to automatically load on system startup by adding it to your `/etc/modules` file.

TODO List

I'd also like to request input on methods to do the following:

```
* Have mod_portaudio automatically connect to the conference instead of doing so manually with fs_cli.  
+ This is possible by calling something similar to:  
  /usr/local/freeswitch/bin/fs_cli -x "pa call *94"  
  ... near the end of a FreeSWITCH initscript, with your conference extension in place of '*94'
```

```
* Have Mumble automatically connect to a configured Murmur server on startup.  
+ This is possible with Mumble 1.2.2 and this fantastic patch:  
  https://sourceforge.net/tracker/?func=detail&aid=3044142&group_id=147372&atid=768007  
+ According to slicer at the Mumble project, the current (Aug 13th, 2010) GIT trunk of  
  Mumble is capable of connecting to the 'last server' automatically, which could also  
  work. This is probably the best method, since the feature will be actively maintained.
```

If there are better solutions, please comment. Otherwise, once/if the Autoconnect patch is accepted on Mumble upstream, instructions will be adjusted accordingly.

Credits

I stole the snd-aloop and asoundrc idea from this: http://alsa.opensrc.org/Jack_and_Loopback_device_as_Alsa-to-Jack_bridge