

CESNET Technical Report 18/2008

**Broadcasting and Transmitting Audio Signals
Using the Lossless FLAC Compression
Encoding**

MILOŠ WIMMER

Received 5.12.2008

Abstract

This technical report describes the cooperation of CESNET and Czech Radio on streaming high audio quality broadcasts over high speed networks. The report describes the FLAC codec used for this purpose and gives a detailed overview of the streaming server. Available clients are also discussed.

Keywords: audio streaming, FLAC, high-speed network

1 Introduction

CESNET has been maintaining a long-term partnership with the Czech Radio, assisting the live broadcasting of its radio stations over the Internet. We focus mostly on high-quality broadcasts, relying on a purpose-built streaming system based on the lossy Ogg Vorbis compression format and bitrates of either 128 kbps or 224 kbps [1], [2], [3].

Growing availability of high-speed Internet connections increases the importance of providing a lossless streaming solution that could transport audio signals from producers to listeners' client equipment unmodified, in their original form. Our technical report aimed at designing and implementing an experimental streaming system relying on the FLAC¹ lossless compression format.

2 System and Environment Description

Every streaming system can be split into four basic components:

1. one that acquires and processes the inputs and provides the input audio signal in a digital form,
2. a producer encoding the audio signal and transmitting the resulting compressed audio to a streaming server,
3. a streaming server accepting streams from producers and redistributing them to end user clients,

¹ <http://flac.sourceforge.net/>

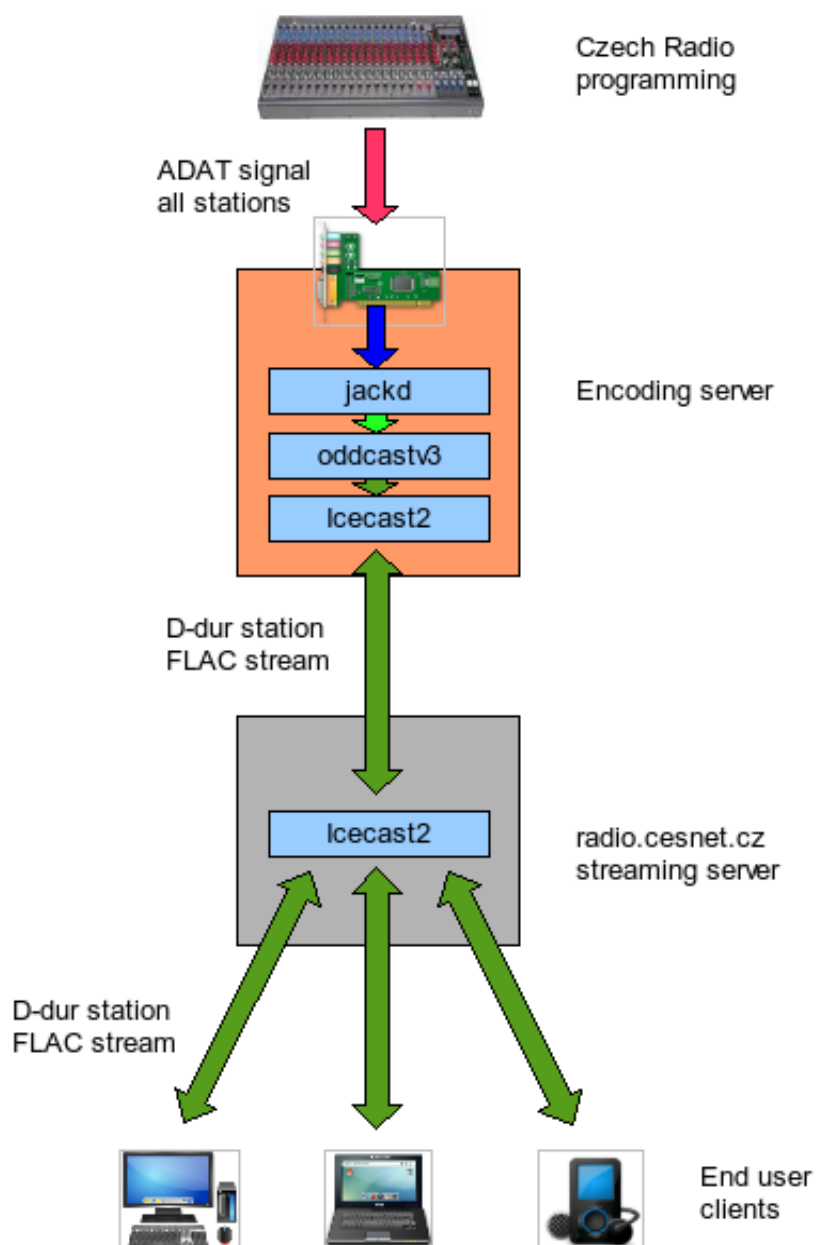


Figure 1. System topology diagram

4. clients used to play the streams on the end users' (listeners') equipment.

Figure 1 illustrates the topology of the system.

The experimental streaming system receives the digital input audio signal directly from the Czech Radio center at Vinohrady, Prague. The signal is being processed by an encoding server transforming it into a compression format stored in an Ogg transport container, and the resulting stream is then transmitted to an application streaming server running on the same machine. From here, the stream is being received by a production streaming server located in CESNET's core network, and possibly also by additional clients located in the Czech Radio's private network. The production streaming server offers the stream to end-user clients.

Similar to the Ogg Vorbis-based production system, the experimental design is also based solely on free software. All servers within the system run Debian GNU/Linux², the producer uses jackd³ and oddcastv³⁴ applications, and the streaming server runs Icecast²⁵. All these software products, as well as the FLAC codec and the Ogg container format, are being developed under the GNU license, and can be used freely.

2.1 The FLAC Lossless Codec

Compression codecs may be used to decrease the bandwidth required by audio transmissions. Compression is either lossless (full audio information is preserved), or lossy (part of the audio data is lost but the resulting signal sounds almost like the original). All lossy codecs implement so called psychoacoustic models, mathematic representations of human hearing characteristics. These allow the codecs to decide, which part of the information can be most easily recognized by humans. Parts that cannot be heard or do not fit into the required bitrate are stripped. Generally, the lower the compression, the higher the quality and bandwidth, and vice versa.

Standard audio CDs store uncompressed digitized audio at 44.1 kHz and 16 bits in two channels. The resulting bandwidth is 1,411 kbps.

Employing the FLAC lossless codec can decrease the bitrate approximately to two thirds, resulting in a bitrate of about 800 kbps. Lossy codecs are capable of reducing the volume of data even more, the results depending on the level of compression and the amount of information lost. Obviously, the quality of the codec is also a factor. Typically, high-quality music is stored at 320 kbps, or broadcast over the Internet at 240 kbps.

The choice of the FLAC lossless codec has been easy as it is considered the best available option. It is a professional open codec and is not subject to any licensing conditions. More and more players support FLAC and it is the most widely used format for lossless audio compression.

It is, however, necessary to point out that while static FLAC encoding (i. e. converting an original sound file into a FLAC-encoded sound file) and playback are simple to achieve, technologies for streaming live inputs are experimental and not widely supported. As of now, we do not know of any radio station broadcasting in FLAC, however we believe that the use of this technology will spread over time.

2.2 Input Source

Most programs broadcast by Czech Radio stations are stored in an internal DALET production system, converted into the lossy MPEG2 format. As digital technologies grow, however, the Czech Radio starts to process and store some of its programs in lossless formats, mostly wav. Majority of programs broadcast by the new Czech Radio's D-dur station are stored in that format, which is why we have chosen the D-

² <http://www.debian.org/>

³ <http://jackaudio.org/>

⁴ <http://www.oddsock.org/>

⁵ <http://www.icecast.org/>

dur station for our lossless broadcasting. Listeners will be able to receive unaltered audio signal identical with the original.

Digital audio signal produced by the D-dur station is encoded using the ADAT format. Therefore the encoding server needs to be fitted with an ADAT compatible sound card and software capable of capturing and decoding the audio. The ADAT data stream may contain multiple channels – a technology also used by the Czech Radio. The D-dur Station signal is ADAT-encoded at 48 kHz and 16 bits in two channels.

2.3 Producer

The encoding server is located in the Czech Radio's server room. It is a Dell PowerEdge 1800 server fitted with two Intel P4 Xeon/3.4GHz processors and 1 GB of RAM. Incoming ADAT signal is being processed by an RME Hammerfall HDSP 9632 sound card. The server runs 64-bit Debian GNU/Linux version 4.0 (amd64 architecture); kernel version is 2.6.27.5.

A *jackd* low-latency sound server is used to separate the left and right channel of the D-dur station from other channels belonging to other Czech Radio stations distributed via ADAT. Then, both channels are taken over by the *oddcastv3* application, which is capable of receiving the sound input from *jackd*, FLAC-encode it, store in the Ogg transport format, and forward it to the *Icecast2* streaming server.

To preserve the original audio signal, our FLAC encoding also uses the sampling frequency of 48 kHz. The variable bitrate (VBR) of the output stream averages at approx. 800 kbps. As far as processing power is concerned, FLAC encoding is less demanding than encoding in Ogg Vorbis. The *oddcastv3* encoder is the largest consumer of server resources here, yet it utilizes no more than 1 % of CPU time.

Unlike the *ices* server used in our production environment to generate Ogg Vorbis streams, *oddcastv3* cannot handle interrupted output connections to streaming servers, and cannot reestablish them automatically. We regard this as a flaw, as it is not possible to ensure uninterrupted connectivity or streaming server availability. We didn't fix this flaw, however, as we have decided to redesign the solution using an instance of the *Icecast2* streaming server locally on the encoding server, and *oddcastv3* is configured to send its output to that instance. The remote production streaming server downloads the stream from the local one, and renews the connection automatically in case of any interruptions. The advantage of this design is that allows other clients located within the Czech Radio network to receive the stream and use it for further processing.

2.4 Streaming Server

The *radio.cesnet.cz* production streaming server is located on CESNET premises. The server (Dell PowerEdge 860, Intel 3050/2.13GHz CPU, 2 GB of RAM) also serves as a main streaming server for Ogg Vorbis streams broadcast by other Czech Radio stations.

The streaming server runs an *Icecast2* application server receiving output from Ogg Vorbis stream producers. To transfer the FLAC stream, the application server is set up as a Single Broadcast Relay, actively downloading the stream from the Mas-

ter *Icecast2* server running on the encoding server located in Czech Radio's internal network.

All streams are then made available to internet clients via standard http. Clients may access the streams at addresses matching the following format:

http://radio.cesnet.cz:8000/required_stream_mount_point,

for example *http://radio.cesnet.cz:8000/cro-d-dur.ogg.*

The Czech Radio D-dur station is available at:

http://radio.cesnet.cz:8000/cro-d-dur.flac

All individual streams are published as links at home page⁶ of the streaming servers, and some of them are additionally available through Czech Radio's live broadcasting web pages⁷ and internet radio index servers.

A full list of streams broadcast by our streaming server is available at *http://radio.cesnet.cz/.*

2.5 Clients

Clients are computer programs run by end users (listeners) on their computers to receive and process streams from the streaming server. Most usually, processing consists of decoding into PCM signal and playing on a sound card.

Static FLAC file playback is supported by many contemporary clients. However, there are only a handful of clients supporting FLAC streams. We have positive experience with foobar2000⁸ and vlc⁹ but some advanced table-top receivers also support FLAC streams.

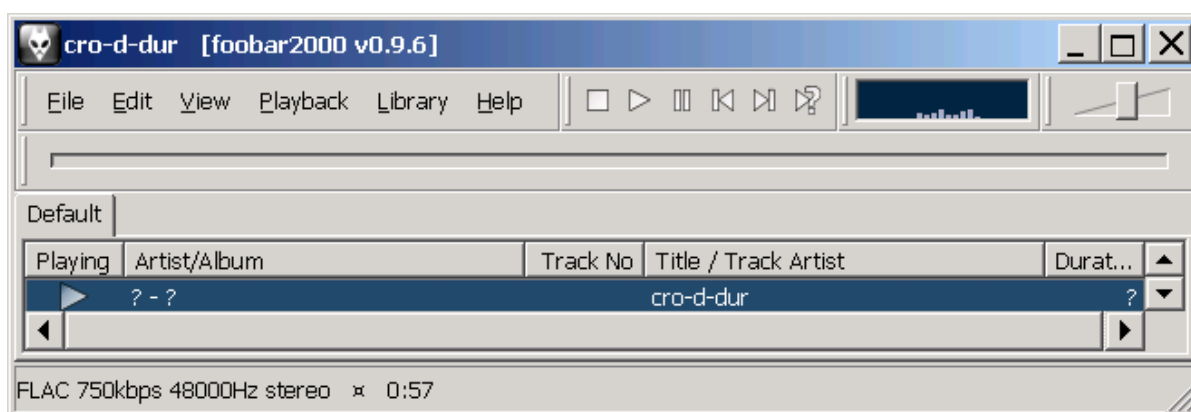


Figure 2. Foobar2000 client

⁶ <http://radio.cesnet.cz/>

⁷ <http://www.rozhlas.cz/>

⁸ <http://foobar2000.org/>

⁹ <http://www.videolan.org/>

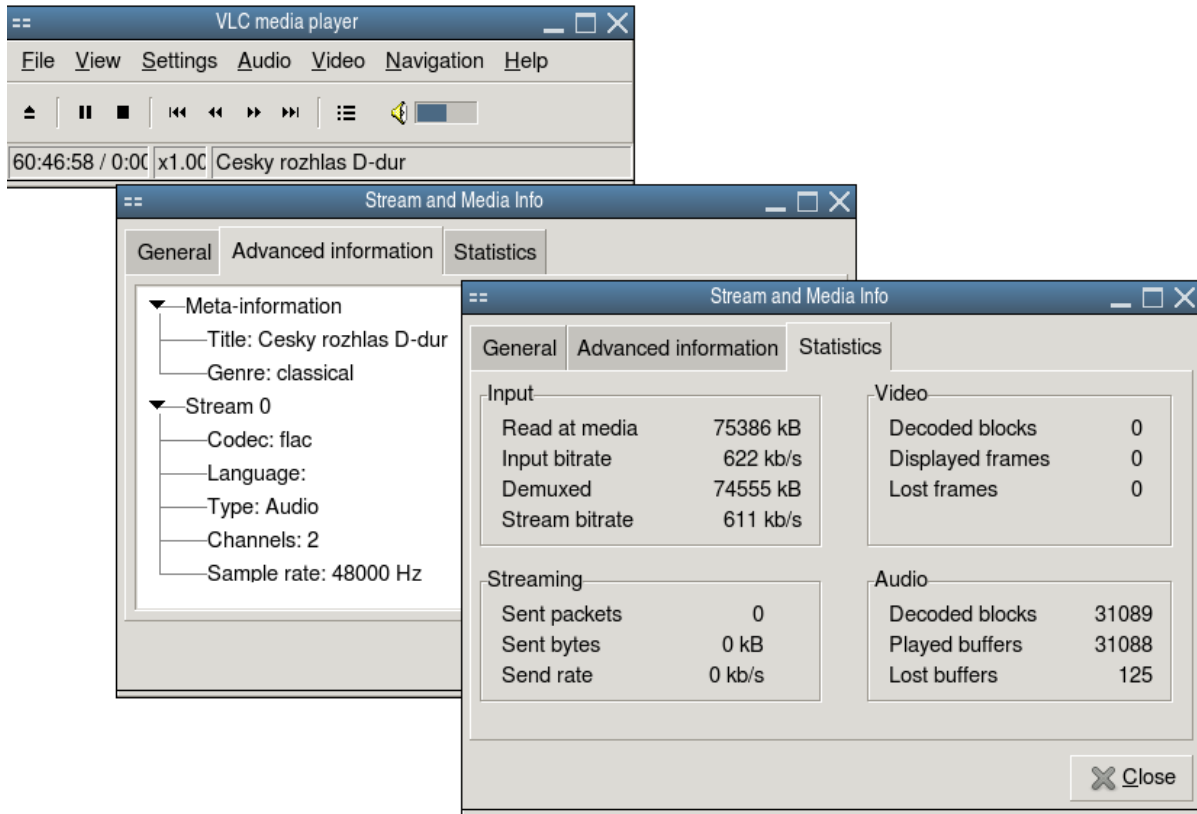


Figure 3. VLC client

3 Appendix – Technical Details

The previous section gave an overview of the system. This Appendix focuses on technical details and discusses problems encountered during implementation.

The choice of producers capable of generating a FLAC stream is limited. The *ices* producer used in our environment to produce Ogg Vorbis streams does not support FLAC. We have identified three candidates: *oddcastv3-jack*, *edcast-jack*, and *vlc*. Since the input signal is being produced by a *jackd* daemon, the producer needs to support that kind of input. Although *vlc* versions 0.9.x already support that, it does not work. The *edcast-jack* program is the direct descendant of *oddcastv3-jack*, it has preliminary support for FLAC encoding in its

source code, but it is not working yet as of version 3.1.7. The only option left is *oddcastv3-jack*, which – built with *libFLAC* and *libjack* libraries and slightly modified as shown below – works correctly.

3.1 jackd server (encoding server)

As we are extracting the Czech Radio D-dur Station signal from the ADAT stream containing other audio channels as well, we need to employ the *jackd* low-latency sound application server connected to an ADAT sound card to separate individual channels.

There were problems with *jackd* server version 0.101.1 available from the operating system distribution, which is why we had to build the latest version 0.109.2

including libraries. We use the following command to start the *jackd* daemon:

```
/usr/local/bin/jackd --driver alsa --device hw:1 \  
--rate 48000 --hwmon --hwmeter --softmode
```

3.2 oddcastv3-jack (encoding server)

Our producer is *oddcastv3-jack-3.1.6* with its source code extended with a function used to set the FLAC encoder sampling frequency:

```
diff ./src/liboddcast.cpp.orig ./src/liboddcast.cpp  
2315a2316  
> OggFLAC__stream_encoder_set_sample_rate(g->flacEncoder, g-currentSamplerate);
```

Real-life use, however, revealed occasional faults occurring randomly, sometimes after as much as a week of operation. These resulted in the application stopping and interrupting the stream. We were able to track the error to the system library *libOggFLAC.so.3*, relied upon by *oddcastv3*. That is why we have built the *flac-1.1.2* package from source, including the *libOggFLAC.so.3* and *libFLAC.so.7* libraries, and installed the resulting binaries in */usr/local/*. Then, we went on to build a new version of *oddcastv3* using those libraries.

We have also used the *external* directory found within the *oddcastv3* SVN tree, and copied its *include* subdirectory with its header files into the root of the unpacked *oddcastv3-jack-3.1.6* source package. Then, we have used the following commands to build the product:

```
$ ./configure --prefix=/usr/opt/oddcast --without-lame \  
--with-libFLAC-includes=include/FLAC --disable-libFLACtest  
$ make  
# make install
```

The resulting *oddcastv3* binary is linked with the */usr/local/lib/libjack.so.0* library originating from the *jack-audio-connection-kit_0.109.2* we have also built as explained above.

This gave us a stable producer that meets our expectations. The following command is used to start the producer:

```
/usr/opt/oddcast/bin/oddcastv3 \  
-c /usr/opt/oddcast/etc/cro-d-dur.cfg \  
system:capture_3 system:capture_4
```

Parameters *system:capture_3* and *system:capture_4* denote the left and right channel of the Czech Radio's D-dur station as delivered to the *jackd* server interface.

Important arguments specified in the *cro-d-dur.cfg* configuration file are set as follows:

```
# x.x.x.x is the encoding server's IP address,  
# the producer sends its output stream directly to the Master  
# Icecast2 server  
Server=x.x.x.x  
Port=8000
```

```

ServerPassword=xxxxx
ServerMountpoint=/cro-d-dur.flac
ServerPublic=0
AutomaticReconnectSecs=10
Encode=Ogg FLAC
NumberChannels=2
Samplerate=48000
ServerType=Icecast2
ServerStreamURL=http://www.rozhlas.cz
ServerName=Cesky rozhlas D-dur
ServerDescription=European-style cultural station
ServerGenre=classical
LogLevel=1
LogFile=/usr/opt/oddcast/logs/cro-d-dur.log
SaveAsWAV=0
OutputControl=SERVER,GENERAL,OUTPUT

```

3.3 Icecast2 server (streaming server)

For reasons explained above, we made the *oddcastv3* producer transmit its output FLAC stream to an *Icecast2* streaming server running locally on the same machine in Master mode. The *radio.cesnet.cz* production streaming server is configured as a Single Broadcast Relay, automatically receiving the stream from the Master server. This is achieved by including the following directive in the *icecast.xml* configuration file:

```

<!-- # radio.cesnet.cz configuration file -->
<relay>
  <server>x.x.x.x</server>
  <port>8000</port>
  <mount>/cro-d-dur.flac</mount>
  <local-mount>/cro-d-dur.flac</local-mount>
</relay>

```

4 Conclusion

The technical report aimed at designing and implementing a streaming system capable of distributing the live broadcasting of the Czech Radio's D-dur station with lossless encoding.

The implementation was successful. It is based solely on free technologies, and relies on the FLAC compression format, the Ogg transport container, the *oddcastv3* producer, and the *Icecast2* streaming server.

The lossless stream of the D-dur station can be received from CESNET's streaming server at *radio.cesnet.cz*. It gives the listeners a unique chance to experience top quality audio signal in a form identical with that produced by Czech Radio's own

equipment. It brings even higher quality than DVB-S or DVB-T digital broadcasting, as it does not include any lossy compression at any stage.

References

- [1] WIMMER, M. *Vysílání a přenos audio signálu ve velmi vysoké kvalitě do sítě Internet – rozvoj projektu* (In Czech: Sending and transmission of very high quality audio signal to the Internet – project development). Technical report 16/2004¹⁰, Praha: CESNET, 2004.
- [2] WIMMER, M. *Vysílání a přenos audio signálu ve velmi vysoké kvalitě*. (In Czech: Sending and transmission of very high quality audio signal). Technical report 28/2005¹¹, Praha: CESNET, 2005.
- [3] WIMMER, M. *Improving Reliability of a Streaming System*. Technical report 24/2006¹², Praha: CESNET, 2006.

¹⁰ <http://www.cesnet.cz/doc/techzpravy/2004/audio/>

¹¹ <http://www.cesnet.cz/doc/techzpravy/2005/audio/>

¹² <http://www.cesnet.cz/doc/techzpravy/2006/vps/>