

12 What if you could listen to this PDF?

by Philippe Teuwen

To honor the tradition of polyglot releases, this PDF is also an audio file featuring a 24-bit studio recording of fbz' *Root Rights are a Grrl's Best Friend*, which you can enjoy with MPlayer or VLC media player.

There are some official ways to embed an audio file in a PDF, such as L^AT_EX's `media9` package. Unfortunately, that would only work in Adobe Acrobat Reader, provided that you also install Adobe Flash—quite a reckless prerequisite nowadays. We are not such bad neighbors, so we looked for alternatives.

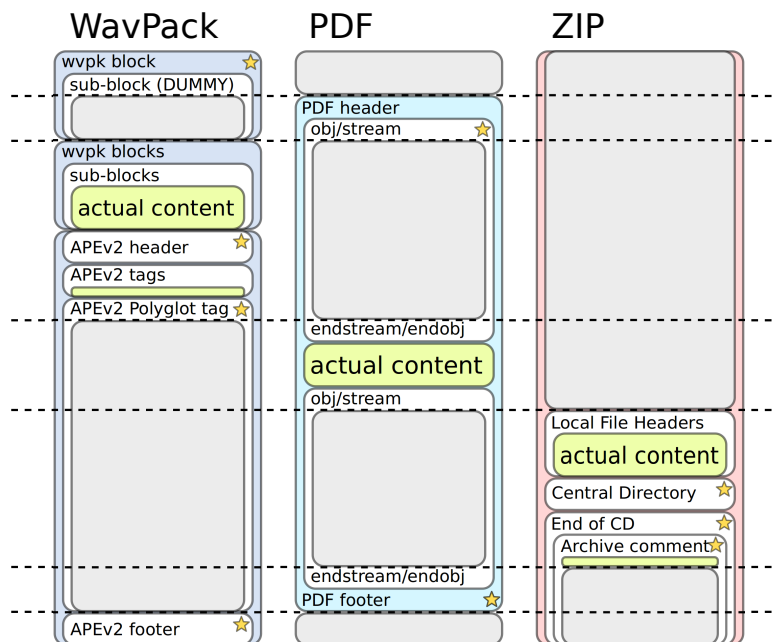
Adobe, once again, is out to search-and-destroy polyglots, so all common audio file types such as WAV, MP3, M4A, 3GP, AAC, FLAC, are prohibited. Still, some less popular formats remain undetected, up until now! Among the free lossless formats these are True Audio (`.tta`) and WavPack (`.wv`).

TTA frame structure³⁰ is unfortunately too rigid and doesn't allow much trickery to inject the start of the PDF within the first kilobyte. It supports standard tagging by ID3v1/v2 and APEv2, but prepending ID3 info is banned by Acrobat. The APEv2 specification,³¹ on the other hand, *strongly recommends* against using it at the beginning of a file. In practice, audio readers don't support files starting with APEv2.

The WavPack file format³² is quite unusual, but far more friendly to us: it doesn't have a file header, but every block starts with the same magic `wvpk`. We can add new metadata blocks at the beginning of the file, and they support `DUMMY` sub-blocks, meant for padding. So we can inject the beginning of a PDF, but can we use those sub-blocks to inject the full PDF in our WavPack? For each sub-block the theoretical size is 16 Mb, but in practice MPlayer accepts a maximum of 1,047,548 bytes and VLC 1,048,548 bytes and only one such sub-block per block. So it's possible, but it would be quite impractical to slice the PDF in 1Mb chunks. WavPack also supports ID3v1 and APEv2. ID3v1 is too limited (only ID3v2 allows `PRIV` frames), so we have to rely on APEv2 to inject the bulk of the PDF (and ZIP, as usual) in a large metadata frame.

We now have the ingredients to build a PDF/ZIP/WavPack polyglot file. The final file structure, from the three perspectives, is depicted on the right.

All starred items contain a size or an offset that depends on another part of the polyglot, so the file is built in two passes. The first pass puts the elements together, and then the second pass adjusts those fields in the WavPack and ZIP.



By the way, the artwork on page 60 is by Ange and myself, derived from Vectorportal's artwork³³ licensed under a Creative Commons Attribution 3.0 Unported License.

³⁰http://en.true-audio.com/TTA_Lossless_Audio_Codec_-_Format_Description

³¹http://wiki.hydrogenaud.io/index.php?title=APEv2_specification

³²http://www.wavpack.com/file_format.txt

³³<http://www.vecteezy.com/people/23511-marilyn-monroe-vector>