# Projects on the Move



**The free software market is a very special little biotope, and life within it can be pretty treacherous at times, as the Tortoise SVN developer learned thanks to PayPal. Voice controls for the desktop are a far more harmonious affair.** *By Carsten Schnober*

One thing that makes the open source world likeable is the number of small projects whose maintainers are motivated by the fun of programming rather than any commercial or strategic concerns, giving the users of free software a choice of a more or less infinite variety of specialist applications that come free of charge and that they can develop and improve themselves if the need arises.

## Fun Not Cash

In many cases, this model has worked for years without any trouble, mainly because money wasn't an issue. As a free software distributor, you don't even need to worry about hosting costs because sites such as SourceForge [1] offer the infrastructure necessary for communications and development free of charge. In an ideal world, instead of economic pressure and struggles, free software will only experience transparency, openness, and self-determination.

But even in the world of free software, it is no contradiction for developers to sometimes ask for financial support for their hard work – preferably in the form of donations from satisfied users. These donations are typically small sums of money that rarely are enough to let developers make the project their mainstay, but even a symbolic amount can provide motivation as a sign of recognition. In many cases, donations do at least cover the costs of running a server.

But, this is where a programmer starts to depart from the realm of free software. Transferring money across the Internet means having an infrastructure that automatically includes a rat's nest of technical, legal, and compliance issues. SourceForge relies on the quasi-monopolist, PayPal, for this kind of thing.

## Money Not Fun

Swiss developer Stefan Küng recently discovered that the world of financial services has very little in common with the typically friendly intercourse between members of the open source community. Like many other developers, Küng added the prebuilt

PayPal button to his project homepage with the intent of giving a financial boost to his work on Tortoise SVN [2], which is a graphical Subversion client for Windows.

In September of this year, Küng received a message from PayPal telling him that his application for a donation account had been refused – which came as some surprise because he had not explicitly made an application. The unpleasant thing about this situation was that the existing account was blocked immediately so that Küng was unable to access the money that was already deposited there.

Küng's initial reaction was to question this decision by email. From the standard response, he discovered that, as a Swiss citizen without a registered charity, he was not permitted to maintain a donation account and that the money in his blocked PayPal account would be inaccessible for 180 days.

Stefan Küng and PayPal exchanged a number of mail messages, which Küng published along with his view of things on his blog page [3]. The dispute ended with a promise by Küng not to use his PayPal account to collect donations in the future, but this was not a voluntary act, because PayPal had threatened to refuse Küng access to his money for 180 days if he failed to acquiesce. Küng also referred the case to SourceForge, whose lawyers took up the cause.

## Bug in the System

The legal investigation had not been completed when this issue went to press. But SourceForge has now at least posted a warning to users of SourceForge and PayPal to read the general terms and conditions in detail before opening an account. Although SourceForge facilitates setting up PayPal donation accounts for free projects, the contract is exclusively between the account holder and PayPal [4].

Reading the terms sounds like something you would do automatically, but if you try reading the 25-page usage agreement, you will quickly see that it is hard going and not particularly intelligible without expert help.

At the end of the day, PayPal users either need a lawyer or need to be willing to expose themselves to boundless, legal machinations and thus finally to the whim of the PayPal providers.

SourceForge can't hide the fact that it is perplexed; the powers that be at SourceForge don't currently see a practicable alternative to PayPal. As an alternative, they suggest donations in kind via an Amazon wish list.

Stefan Küng advises all PayPal account holders (and he is understandably frustrated here) to withdraw their money before the online payment system's operators decide to block their accounts without any notice.

He compares PayPal's behavior with that of a car dealership that decides not to sell anything to a customer in the future, and instead locks away the customer's car for six months. You can follow the developments in this story via the Bug Report [5] on the SourceForge site (Figure 1).

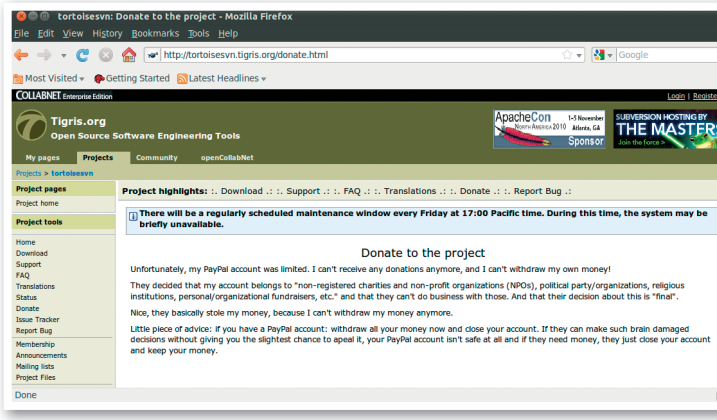


**Figure 1:** The Tortoise SVN project is fighting a very special bug. The project's donations account was blocked by the account provider.

## On Screen, Computer!

At least in science fiction series, this is how people talk to their computers. The computer not only understands what the user wants but also answers fluently and intelligibly in a human (or alien) language.

Even a simple scenario like this can cause enormous difficulty, however. First of all, the computer needs to identify the sounds made by the users as language, and then it needs to interpret them correctly. Once the speech recognition obstacle has been overcome, the application needs to plumb the depths of artificial intelligence: The context decides on how the command is interpreted. Giving a response seems almost trivial in comparison. The computer only needs to read something to the user.

In real life, as many readers will be aware, the first two sub-problems, speech recognition and interpretation, are currently only solvable in an ideal environment and with a fairly substantial error rate. Voice synthesis is still immature and often produces incorrect or unintelligible pronunciation.

OpenTTS [6] sees a new free project venture into the text-to-speech (TTS)

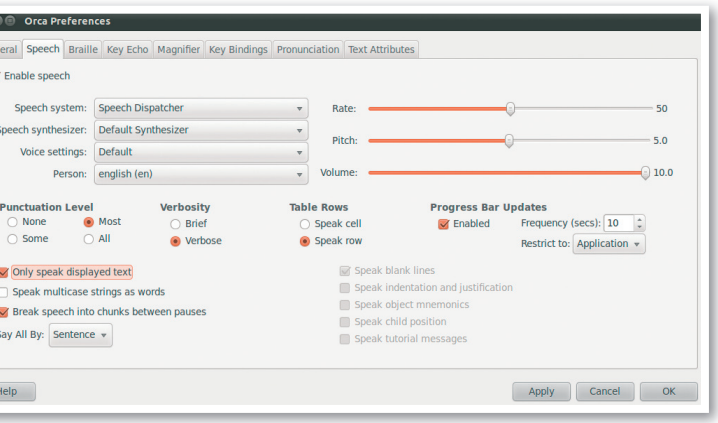


**Figure 2:** Gnome Speech will read on-screen text to help users with vision impairments use the desktop.

field to tackle the current issues. The project team only published the developer version 0.1 in June; however, the latest version of the Gnome desktop already integrates the tool and uses it for voice output.

The OpenTTS team didn't have to start from scratch; it forked the project from Speech Dispatcher [7]. Both projects provide interfaces for various front ends. Besides the Java interface, OpenTTS additionally provides a server application that can send texts for local and remote clients to read.

Whereas the OpenTTS project is looking to revive the creative process, which has been fairly lame since 2008, the developers behind the Speech Dispatcher project have criticized the fork and fear that parallel work will mean even less progress than in the past.

## Front and Back

Gnome Speech [8] version 2.32 and later has the ability to access OpenTTS, although it optionally uses other sound severs. As part of the accessibility package, Gnome Speech targets people with special needs to facilitate the use of the desktop by Linux users with sight impairments. Orca (see Figure 2) provides the graphical front end to this. You just need to tell the screen reader which text synthesis program to use and which areas of the desktop the system should read aloud.

Alternatively, Orca will work with the eSpeak [9] and Festival [10] text synthesizers. The screen reader output quality, however, depends less on the back end than on the voices available.

Modern TTS programs use a machine-learning approach to correct pronunciation. A human needs to read as much text as possible; phonetic markup is then applied to the text. Experts refer to the results as a voice.

When the TTS application sees a new written text, it cuts out parts of the existing data and composes a spoken message from them. The available systems mainly differ with respect to size and characteristics of the individual snippets of text, as well as to the statistical methods they apply to discover the best possible units.

This probably explains why Gnome Speech offers commercial voices for purchase in addition to the free voices provided with the open source back ends. The manufacturer, Cepstral [11], provides compatible records for German, Italian, Spanish, French, and various types of English, as well as in male and female variants. Another commercial alternative is Fonix [12]; the manufacturer sells a complete back-end system for Gnome Speech for around $US 40.

## Commander

The developers are also working hard on automated speech recognition. The CMU Sphinx [13] project from Carnegie Mellon University (CMU) in Pittsburgh will be publishing version 1.0 shortly, but as is often the case with academic projects, the work is mainly focused on the technical background, which means end users don't actually get to see much right now.

In the meantime, the front ends are making visible, or should I say audible, progress. Several projects are currently developing front ends for Sphinx. Gnome users, for example, can try Gnome Voice Control [14] as a desktop applet. And, even if the development of the front end itself has been dormant for three years, users will still benefit from the technical advances made by recent versions of Sphinx.

A speech recognition system does the opposite of what a text synthesizer does. Developers still rely on human speech for training purposes, however, with the user acting as the coach. When the program accepts a spoken command, the program then compares it with the text read by the user and, on the basis of various parameters, retrieves the statistically most probable counterpart from the database.

This summary of speech recognition and text synthesis is both positive and negative. It is a good thing that free software plays some kind of role in an area that is fraught with commercial solutions. But unfortunately, it will take quite a while before your Linux computer understands commands like "Earl Grey, piping hot!" ■■■

## INFO

[1] SourceForge: *http://www.sourceforge.net*
[2] Tortoise SVN: *http://tortoisesvn.net*
[3] Stefan Küng: "How PayPal screws open source projects", *http://tortoisesvn.net/howpaypalscrewsopensourceprojects*
[4] SourceForge: "Know Your Rights with PayPal", *http://sourceforge.net/blog/warning-to-open-source-projects-know-your-rights-with-paypal*
[5] Bug "Donations not legally sound": *http://sourceforge.net/apps/trac/sourceforge/ticket/13993*
[6] OpenTTS: *http://www.opentts.org*
[7] Speech Dispatcher: *http://www.freebsoft.org/speechd*
[8] Gnome Speech: *http://live.gnome.org/Orca/GnomeSpeech*
[9] eSpeak: *http://espeak.sourceforge.net*
[10] Festival: *http://www.cstr.ed.ac.uk/projects/festival*
[11] Cepstral: *http://www.cepstral.com*
[12] Fonix: *http://www.fonixspeech.com*
[13] CMU Sphinx: *http://cmusphinx.sourceforge.net*
[14] Gnome Voice Control: *http://live.gnome.org/GnomeVoiceControl*