

The Infrastructure Behind TDF

- ▼ Alexander Werner and Florian Effenberger
- ▼ LibreOffice Conference, Paris, October 15th 2011

Who's disturbing my sleep?

▼ **Alexander Werner**

- ▼ member of the admin team at TDF
- ▼ been involved with TDF from the very beginning

▼ **Florian Effenberger**

- ▼ been with OpenOffice.org for seven years, eventually as Marketing Lead and Distribution Lead
- ▼ member of the Steering Committee at TDF and member of the admin team
- ▼ responsible for marketing and infrastructure

Infrastructure – the unknown being?

- ▼ you'll never see infrastructure until it doesn't work
- ▼ technically enables the community to do all sorts of fun things
- ▼ it provides an important framework and platform for them
- ▼ has to be dynamic, scaleable, stable
- ▼ “infrastructure provides the playground where things happen”

We didn't start the fire...

- ▼ existing OpenOffice.org infrastructure at CollabNet with migration to Kenai planned
 - ▼ mainly tied to one IT service provider
 - ▼ many integrated services, a few external ones
 - ▼ “somehow works” and contains 10 years worth of data
 - ▼ one of the high-profile sites on the internet
- ▼ back then: pending Oracle decision about TDF also decides about future infrastructure

Status quo: three-way approach

- ▼ TDF-hosted services
 - ▼ mailing lists plus archives
 - ▼ websites and wiki
 - ▼ download management system (MirrorBrain) and rsync
 - ▼ translation services, DMS
- ▼ FreeDesktop.org-hosted services
 - ▼ development mailing lists
 - ▼ git repository, bug tracker (BugZilla)
- ▼ external services
 - ▼ worldwide mirrors
 - ▼ WordPress Blog

Numb3rs

- ▼ team of **14** administrators from **5** countries
- ▼ **7** TDF servers, plus FreeDesktop.org, plus external
- ▼ TDF alone handling up to **10 TB** of monthly traffic w/o DLs
- ▼ **82** mirrors worldwide handled more than **6.000.000** DLs
- ▼ **110** mailing lists in 35 languages, over **600 MB** in total
- ▼ **15,500** subscribers, half of them for announcements
- ▼ in total, over **200 GB** of web-hosted data at TDF
- ▼ we handled **350.000** DLs the first week

I want to break free

- ▼ TDF not only managed to attract developers, maintain a product, create, shape and strengthen a brand, do marketing
- ▼ but also to take care of all technical needs by itself
- ▼ all based on free software
- ▼ all knowledge is within the community
- ▼ we are not tied to one specific service provider anymore
- ▼ we contribute back to other projects (mlmmj and PyMIME)
- ▼ in a nutshell:
Our infrastructure development reflects the freedom, openness and transparency that TDF stands for.

Some admin porn...

- ▼ Ubuntu 10.04 LTS as base system, moving to 11.10
- ▼ rented dedicated servers at a German datacenter
- ▼ well-known components used
 - ▼ Apache with mod_php or PHP-FastCGI, xcache
 - ▼ MySQL
 - ▼ Postfix, SpamAssassin, ClamAV, amavisd-new, Postgrey, PostFWD, policyd-weight
 - ▼ Dovecot
 - ▼ vsftpd
 - ▼ Mailman (for internal lists), mlmmj (for public lists)
 - ▼ rsnapshot for backup strategy

MirrorBrain - our mastermind

- ▼ introduced at the OpenOffice.org infrastructure as replacement for “Bouncer”, heavily refined and enhanced now
- ▼ Apache module with rsync
- ▼ database of mirrors with their respective bandwidth/priority
- ▼ selection of best download server based on GeoIP
- ▼ files are automatically searched, indexed and hashed
- ▼ GPG signatures are verified automatically and reported
- ▼ download statistics are generated
- ▼ three-tier mirroring system
 - ▼ push mirrors
 - ▼ staging mirrors
 - ▼ regular mirrors

Puppet on a string

- ▼ open source configuration management tool
- ▼ written in Ruby
- ▼ configuration is described using DSL (Domain Specific Language)
- ▼ every server can be configured differently
- ▼ configuration changes are automatically distributed amongst servers
- ▼ adding a new server is deadly easy

Puppet provides flexibility and expandability needed for the future.

Gimme gimme gimme...

- ▼ we contribute back to other projects
- ▼ mmmj is our mailing list system
 - ▼ still rather unknown, but very flexible
 - ▼ moderators can work completely via e-mail
 - ▼ one file per message for easier backups
- ▼ we filed bug reports and lots of enhancements suggestions
- ▼ currently the help texts are localized in various languages
- ▼ TDF's own development: PyMIME under GPL license
 - ▼ converts charsets
 - ▼ adds footers
 - ▼ strips attachments and HTML messages

Thank you for listening!

- ▼ Questions?
floeff@documentfoundation.org
alex@documentfoundation.org
- ▼ Twitter/Identicia: [@floeff](https://twitter.com/floeff)



All text and image content in this document is licensed under the [Creative Commons Attribution-Share Alike 3.0 License](https://creativecommons.org/licenses/by-sa/3.0/) (unless otherwise specified). "LibreOffice" and "The Document Foundation" are registered trademarks. Their respective logos and icons are subject to international copyright laws. The use of these therefore is subject to the [trademark policy](#).