

# Linux, the Universe and Everything

A closer look on our open source work environment. Trends. Challenges.

Stephan Wiesand  
Technical Seminar  
Zeuthen, 2016-05-24



## About

- > A typical scientist's daily routine:
  - Log in to Linux desktop; start some local applications
  - ssh-Login to Linux WGS; develop, compile, test code
  - Run jobs on the farm or cluster; store results
  - Combine & analyze results; include in thesis or paper
- > This obviously requires a lot of software, including
  - "Linux" distribution(s), comprising much more than the kernel
  - Network file systems
- > With very few exceptions, we don't pay anything for this software
  - The price is certainly right, and we get a lot more than what we pay for
    - > Maybe even what we need
- > Can we take this for granted?



# Boundary Conditions

- > Reliability
  - Network file systems are surprisingly hard to get right
- > Long term stability
  - No frequent changes requiring user action
  - Some users still depend on SL5
    - > (~ 9 years old)
- > Security
  - Open environment, hundreds of users, high fluctuation
  - The best person's account can become compromised and abused
- > => Maintenance



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 3

# Linux

- > Strictly speaking, this is just the kernel
- > Started out 1991 as a hobby project, GPL license
  - Volunteers still contribute significantly
- > But today most of the work is done by people paid for it
  - > 80%
  - Among the top 30 contributing companies [1]:
    - > Intel, Red Hat, Samsung, IBM, SUSE, Texas instruments, Google, Oracle, AMD, Nvidia, Broadcom, ARM, Cisco, Fujitsu, Linux Foundation, QLogic
      - Volunteers still lead the list
    - > Long tail of > 400 companies
  - => Vendors of hardware or services, Linux distributors
- > Mainline releases ~ every 2 months
  - typically ~ 10k changes, ~ 250 *new* contributors (> 1000 total)



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 4

# Linux Particularities

- > Stable API/ABI provided to userland
  - System calls, used by libc or applications directly
- > No stable API internally
  - affects kernel extensions (modules)
- > Interfaces change all the time
  - in-kernel callers are adapted
  - external ones must cope
- > Interfaces are removed and replaced with new ones
  - sometimes the new one is GPLONLY
    - > can only be used from code under a GPL compatible license



# “Linux” Distributions

- > Kernel
  - + boot loader, libc, init system
  - + installer
  - + userland (applications, services)
  - + documentation & support
  - + maintenance
- > Some components are developed in-house
- > Most are taken from “upstream” projects - like the Linux kernel
  - often, distributors are major contributors and/or sponsors
    - > X.org, GNOME, freedesktop.org, ...



# Distribution Families

- > debian
  - Community driven
  - stable ← testing ← unstable
- > Red Hat
  - RHEL (product) ← Fedora (community)
- > SUSE
  - SLE (product) ← OpenSUSE (community)
- > Derivatives
  - debian: Ubuntu (backed by Canonical Ltd.)
  - RHEL: CentOS, SL, OEL
    - > Cloning from the free sources is possible, requires trademark removal
    - > CentOS, SL are not regarded as competition
      - CentOS maintainers hired by Red Hat, project sponsored



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 7

# Distributions: Release Cycles & Life Times

- > debian stable:
  - released ~ every 2 years (effectively, no commitment)
  - maintained for 1 year after next stable release
    - > new Long Term Support project: + 2 years (partly sponsored)
- > RHEL:
  - Major releases ~ every 4 years (effectively, no commitment)
    - > Minor releases ~ every 8 months
    - > Kernel ABI stable within minor release
    - > Whitelisted ABI stable within *major* release
  - maintained for 7 years, + 3 years with limitations
- > Ubuntu:
  - LTS releases every 2 years (non-LTS: every 6 months)
  - maintained for 5 years (non-LTS: 9 months)



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 8

# Ubuntu: Supported for 5 Years?

## > careful: many components are only for 3 years

### ■ quite a few only for 9 months

#### > anything from universe/multiverse: not at all

#### ■ someone *may* be caring for the package

#### > but no guarantees

```
root@mops03:~# ubuntu-support-status
Support status summary of 'mops03':
```

```
You have 20 packages (1.1%) supported until January 2017 (9m)
You have 36 packages (1.9%) supported until April 2019 (3y)
You have 4 packages (0.2%) supported until February 2017 (9m)
You have 1766 packages (93.9%) supported until April 2021 (5y)
```

```
You have 0 packages (0.0%) that can not/no-longer be downloaded
You have 54 packages (2.9%) that are unsupported
```



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 9

# Stable Kernel ?

## > A simple change from Linux 4.4:

```
--- a/fs/splice.c
+++ b/fs/splice.c
@@ -809,6 +809,13 @@
 */
static int splice_from_pipe_next(struct pipe_inode_info *pipe, struct splice_desc *sd)
{
+ /*
+  * Check for signal early to make process killable when there are
+  * always buffers available
+  */
+ if (signal_pending(current))
+   return -ERESTARTSYS;
+
  while (!pipe->nrbufs) {
    if (!pipe->writers)
      return 0;
```

## > made it into the “3.13” kernel in 14.04 LTS within weeks

## > and broke the OpenAFS client



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 10

## Remote Filesystems

- > Linux (distributions) come with a couple of clients (and servers)
- > For various reasons, we need(ed) more than those
  - dCache
  - AFS
  - Lustre
- > dCache uses the in-kernel NFS client
- > OpenAFS and Lustre clients need external kernel modules
  - in addition, OpenAFS' license is not GPL compatible



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 11

## OpenAFS

- > originated from “project Andrew” at CMU
- > AFS was then marketed by a spinoff: Transarc Corp.
- > Transarc was then acquired and became IBM Pittsburgh Labs
- > IBM then discontinued the AFS product
  - but made the source available under the IPL License
- > The “OpenAFS project” picked up this source and developed it
  - Project infrastructure: CMU, Stanford, MIT
  - Developers: all volunteers
    - > typically paid for AFS related work
      - Companies selling AFS support
      - Labs & Universities
      - Contractors working for military
  - Worked reasonably well - until recently



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 12

# OpenAFS on Open Hub

## Activity

<b>30 Day Summary</b> <small>Mar 17 2016 — Apr 16 2016</small> <b>5 Commits</b> <b>4 Contributors</b>	<b>12 Month Summary</b> <small>Apr 16 2015 — Apr 16 2016</small> <b>130 Commits</b> <small>Down -218 (62%) from previous 12 months</small> <b>19 Contributors</b> <small>Down -6 (24%) from previous 12 months</small>
--	---

### Commits per Month



## Community

### Ratings

9 users rate this project:  
 ★★★★★ 4.9/5.0


Click to add your rating  
 ☆☆☆☆☆

[Review this Project!](#)

### Most Recent Contributors

- |                     |                 |
|---------------------|-----------------|
| Jeffrey Eric Altman | Stephan Wiesand |
| Michael Meffie      | Benjamin Kaduk  |
| Mark Vitale         | Marcio Barbosa  |

### Contributors per Month



<https://www.openhub.net/p/openafs/>



# OpenAFS on Open Hub

## Contributors: Listing

ⓘ Analyzed 1 day ago, based on code collected 1 day ago.

Showing page 1 of 14      Search / Filter on:       Sort by: All Time Commits

Name	Kudos	12 Month Commits	All Time Commits	5 Year Trend	Primary Language	First Commit	Last Commit
Jeffrey Eric Altman <small>( OpenAFS Gatekeeper, Developer, Elder Emeritus )</small>		34	3496		C	over 12 years ago	about 1 month ago
Daria Phoebe Brashear <small>( Elder, Gatekeeper, Developer )</small>		1	1903		C	over 15 years ago	8 months ago
Simon Wilkinson <small>( inactive )</small>		0	1396		C	about 9 years ago	over 1 year ago
Andrew Deason <small>( inactive )</small>		2	1006		C	about 10 years ago	9 months ago
Marc Dionne		8	392		C	almost 14 years ago	9 months ago
Michael Meffie <small>( inactive )</small>		19	304		C	almost 7 years ago	about 1 month ago
Russ Allbery <small>( Developer )</small>		0	294		C	almost 15 years ago	almost 3 years ago
Chas Williams <small>( inactive )</small>		7	288		C	over 15 years ago	4 months ago

➤ Many significant long term contributors are no longer active



## What happened ?

- > Only a handful of contributors left
  - all part time, or even doing it as a hobby
- > Some contributors changed jobs, no longer deal with AFS
  - some voluntarily, some not
- > And/or they are frustrated
- > One of the two companies selling AFS support withdrew
  - almost completely (still collaborate on security issues)
  - marketing a proprietary fork, now see OpenAFS as competition
- > On the payroll: the developer who kept the Linux client afloat
  - and nobody stepped up to replace him
    - > which is hard



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 15

## Result

- > OpenAFS now copes with Linux 4.4
  - and Ubuntu 14.04/16.04 LTS
    - > but no fixed packages yet (except from a ppa)
- > Not yet with Linux 4.5 (work in progress)
  - nor 4.6, nor 4.7rc
- > How long until the next backports break the client again?
- > => Ubuntu + AFS home directories not sustainable a.t.m.
  - Showstopper for the next generation of Linux desktops as planned
- > EL + AFS: same problem - in principle
  - much less likely to strike anytime soon though



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 16



## And Lustre ?

- > CFS Inc. → SUN → Oracle → Whamcloud(/Xyratex) → Intel
- > Does not have the GPLONLY problem
- > But: “Enterprise Edition” product marketed more aggressively now
- > “Community Releases” bring features & fixes
  - no more maintenance releases for those
- > May require moving to new client release if kernel changes
  - clients are compatible with older servers
- > => no immediate danger, but clouds on the horizon



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 17

## Conclusions

- > Using Linux is a solid strategy
- > So is using a free clone of RHEL
  - RH is profitable
    - > SL & CentOS are part of the ecosystem, not competition
  - alas, it's not fun on the desktop
- > Ubuntu LTS is an attractive option for desktops
  - but right now, not if the OpenAFS client is required
  - and not as enterprise grade as EL
  - Is canonical profitable now? Probably not.
- > Lustre perspective slightly unclear
- > Open AFS perspective completely unclear
- > Strategically, try to stick with standards when possible



Stephan Wiesand | Linux, the Universe and Everything | 2016-05-24 | Page 18

## References

- > [1] <http://www.linuxfoundation.org/publications/linux-foundation/who-writes-linux-2015>

