

Thorsten Leemhuis

**What's Up in the
Land of the Linux
Kernel**

%pre

%pre | whoami



%pre | whoami @work



%pre | whoami @work



ct magazin für computer techniek 1 18.12.2015

Smartwatches im Test

Richtig einstellen mit wenig Aufwand

Die Sicherheits-Checkliste

Handy, Router, NAS, Smart-TV, Server, Browser, Facebook ...

Upgrade auf Windows 10 stoppen
VR-Brille Samsung Gear
USB-Sticks mit Typ C
Fairphone 2 im Test
Spiele-Highlights 2015

Audi macht Ernst mit VR

Die Kabel-Deutschland-Lücke
Web-Typographie
Linux-Prozessmanagement
GUI mit Python
Skylake übertakten

Alte und billige Tablets sinnvoll nutzen

Cooler Tablet-Projekte

Second Screen, Notenständer, Info-Display, Auto-Infotainment ... € 4,50

Kernel Log | Linux 4.4

Thorsten Leemhuis

Linux-Kernel 4.4

3D-Grafiktreiber für virtuelle Maschinen und mehr Sicherheit bei RAID 5

Der Linux-Kernel 4.4 wird einen Grafiktreiber für den Rapsi mitbringen. Performance-Verbesserungen im Netzwerksystem sollen DDoS-Angriffe erschweren. Zwei neue Ansätze versprechen High-End-SSDs mehr Leistung zu entfalten.

Das in der ersten Jahreshälfte erwartete Linux 4.4 erweitert die Fähigkeiten des Treibers. Virtuelle-GPU. Dadurch sollen Linux-Distributionen, die in einer mit KVM betriebenen virtuellen Maschine (VM) laufen, in Zukunft die 3D-Beschleunigung des Wirtsystems nutzen können. Bei diesem über mehrere Jahre als „VirtIO 3D“ entwickelten Ansatz nutzt ein Guest-3D-Treiber die OpenGL-Befehle zur Ausführung in den Host weiter. Das Verfahren erfordert neben Linux 4.4 die Grafikbibliothek Mesa 11.1 und den Systememulator Qemu 2.3, beide sollen noch im Dezember erscheinen.

Virtualisierungs-Software von VMware oder VirtualBox ermöglicht schon länger eine Nutzung der 3D-Beschleunigung in Linux-VMs. Die VMware-Produkte sind allerdings proprietär und VirtualBox erfordert oft die mühsame Einrichtung passender 3D-Gasttreiber. Anders als VirtIO 3D funktionieren die 3D-Support dieser Virtualisierungslösungen allerdings auch in VMs mit Windows. Selbigen gilt auch für das noch unfertige „JRMG“ von Intel, bei dem der Gast-Treiber direkt auf Funktionen von Intel-GPUs zugreift.

Rapsi-Treiber

Linux bringt mit Version 4.4 erstmals einen Kernel-Grafiktreiber für die Broadcom-Prozessoren mit, die auf dem verschiedenen Ausführungen des Raspberry Pi steuern. Dieser Treiber funktioniert weitgehend autark und ist nicht auf den Grafiktreiber in der proprietären Firmware angewiesen, wie es bei den derzeit zu meist eingesetzten Treibern der Fall ist. Der in 4.4 enthaltene und von Broadcom selbst vorgegebene Treiber befindet sich aber

bindung, keine 3D-Beschleunigung die sollen Verbesserungen ermöglichen, die in Linux 4.5 einfließen sollen.

Der MD-RAID-Code wird bei Software-RAIDs der Level 4, 5 und 6 ein Log führen können, das auf einem weiteren Datenstrang liegt und Datenverfälschungen bei Systemabstürzen verhindert. Das zugrundeliegende Verfahren ähnelt dem von Journaling Dateisystemen wie Ext4. Der Kernel schreibt jede Änderung zuerst in das Log und erst danach auf die am RAID beteiligten Datenstränge. Falls die Stromversorgung beim Schreiben auf die RAID-Datenstränge ausfällt, kann der Kernel die im Log hinterlegten Daten beim nächsten Start nutzen, um die Integrität innerhalb kurzer Zeit wiederherzustellen.

Das Log kann auch die Geschwindigkeiten/Wartung steigern, da es Änderungen kurz puffert. Die Log-Funktion für MD-RAID stammt von Facebook-Mitarbeitern, die bereits an Erweiterungen arbeiten, die das Log zu einem vollwertigen Writeback-Cache machen. Dabei puffert das Log länger und mehr, was der Geschwindigkeit zugute kommt.

Offene SSDs

Neu ist auch Unterstützung für ein LightNVM-generiertes Framework, das für „Open Channel SSD“ gedacht ist. Mit diesem Begriff bezeichnen die LightNVM-Entwickler einige vorteilhaft für Server gedachte SSDs, bei denen das Betriebssystem einige Arbeiten übernehmen kann, die normalerweise der RAID-Translation Layer (RTL) oder das Bad Block Management der SSD-Firmware erledigen. Das Delegieren ans Betriebssystem



ge-Algorithmen meist tun. RACK ist vorerst experimentell und stammt von Google. Das Unternehmen setzt den Algorithmus offenbar schon eine Weile ein und hat ihn bei der IETF zur Standardisierung eingereicht.

Ungeplante Anwendungen können entstehen mit dem eBPF (extended Berkeley Packet Filter) ausgeführte Programme in den Kernel laden, um damit Datenströme zu verarbeiten, die durch den Kernel fließen. Dadurch kann beispielsweise ein nicht von Host ausgehendes Topik im Kernel eBPF Filter beim Kernel installieren, damit nur die Netzwerkpakete an den Host weitergeht, die der Nutzer untersuchen will. Von ungeplanten Anwendungen stammende eBPF-Programme unterliegen allerdings einigen Einschränkungen, damit Angreifer den eBPF-Interpreter nicht missbrauchen.

Unabhängig davon haben die Entwickler die Performance-Analyse-Werkzeuge erweitert, damit es eBPF-Programme automatisch bewerten, prüfen und in den Kernel laden kann. Der Kernel kann mit solchen Programmen analysierte Events frühzeitig aufführen, um Overhead und den Einfluss der Analyse zu reduzieren.

Langzeit-Kernel

Facebook-Entwickler haben die Prozessoren beim Einsatz der BPF-Mechanismen ist, um es zu evaluieren. Zuerst helfen sie festzustellen, dass die darüber aktivierte Datenverteilungsmethode die Performance bei ihren Hardware-RAIDs der Level 1 und 6 erheblich verbessert.

Zu dem neu zum Kernel stehenden Treiber gehört einer für USB-WLAN-Chips von Realtek, für die es bislang nur einen Staging-Treiber gab, der größere Qualitätsanforderungen aufweist. Der Kernel 4.4 wird zudem einige per Firmware angesprochene Sound-Chips besser unterstützen und spricht per I.S. angestrichelte Audio-Chips von Skylake-Notebooks an. Der Webcam-Treiber steuert jetzt vier weitere Gerätebibliothek-Unterstützung an. Linux 4.4 wurde zudem vorab zu einem Langzeit-Kernel erklärt. Daher soll es nicht nur knapp drei Monate, sondern bis mindestens Juni 2016 mit Fehlerkorrekturen und kleineren Verbesserungen versorgt werden. (thl/ctd)

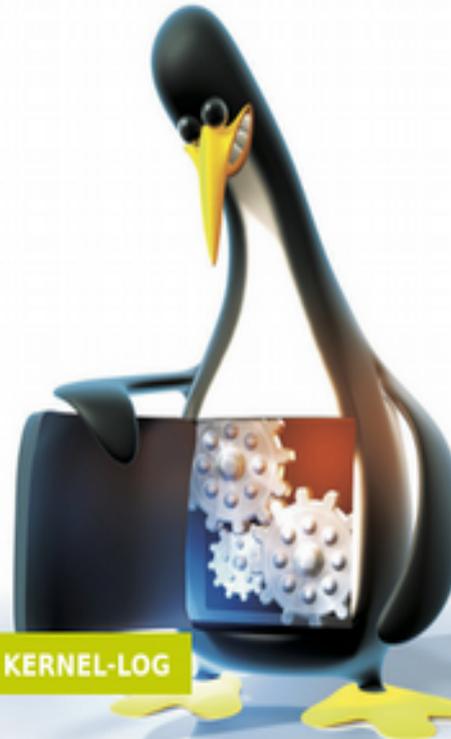
%pre | whoami @work

INFOS ZUM ARTIKEL

Kapitel

- 01 Geschwindigkeitswahl wird optimiert
- 02 Wartezeiten bei starkem Netzwerkverkehr reduzieren
- 03 EFI-Firmware spielt BIOS-Updates selbst ein
- 04 Parallele Verzeichnisabfrage
- 05 Kopierbeschleuniger für NFS
- 06 Knautschzone
- 07 Kernel abdichten
- 08 Unterstützung für rund fünfhundert weitere Geräte
- 09 Treiber für Thunderbolt, Xbox-Controller, Audio-Codecs, ...
- 10 Farbmanagement beim Intel-Treiber
- 12 Grafiktreiber: GeForce 830M, Raspi, 5K-Display, ...
- 13 Unterstützung für Platten mit Shingled Magnetic Recording (SMR)
- 14 Mehrere Dateisystem-Bäume bei Ceph
- 15 Direktzugriff auf schnelle

Die Neuerungen von Linux 4.7 UPDATE



TRENDS & NEWS | KERNEL-LOG

Thorsten Leemhuis

30.05.2016

Kernel-Log, Linux, Linux 4.7, Linux-Kernel

Die neue Kernel-Version unterstützt AMDs neue Grafikchips. Ferner soll Linux 4.7 das Stromsparerpotenzial moderner Prozessoren stärker ausschöpfen und Wartezeiten vermeiden, die bislang bei hoher Netzwerklast auftraten.

Linus Torvalds hat Linux 4.7 freigegeben. Eine der wichtigsten Verbesserungen der neuen Kernel-Version: Der Amdgpu-Treiber unterstützt jetzt die Polaris-Grafikprozessoren von AMD (u. a. [1](#) [2](#)

www.heise.de/thema/Kernel_Log

%pre | whoami @work

The screenshot shows the homepage of 'The H Open' website. At the top left is the logo 'THE H OPEN' in red and black. To its right is a search bar with the text 'Search The H Open' and a 'Search' button. Below the logo is a navigation bar with links for 'Last 7 days', 'News Archive', and 'Features'. The main content area features a large article titled 'Kernel Log: Coming in 3.10 (Part 1) Networking' by Thorsten Leemhuis, dated 07 June 2013, 17:41. The article text discusses Linux 3.10 improvements and the release of the fourth pre-release version. To the right of the article is a sidebar with 'OPEN HEADLINES' and another 'Kernel Log' article titled '4: Drivers'. The website has a green and white color scheme.

THE H OPEN

Open

In association with helse online

Search The H Open Search

Last 7 days News Archive Features

07 June 2013, 17:41 1 2 next »

Kernel Log: Coming in 3.10 (Part 1)

Networking

by Thorsten Leemhuis

Linux 3.10 sees improvements in the way lost packets at the end of TCP transactions are handled, speeding up HTTP data transfer. It also sees the addition of support for VLAN stacking and Realtek's RTL8188EE wireless chip.

Last weekend, Linus Torvalds released [the fourth pre-release version of Linux 3.10](#). Technical problems meant that it was initially only available via the git source code management system, though Torvalds notes in the email announcing the release that he considers git to be the easiest way of getting hold of it anyway. Linux 3.10-rc4 is now also available as a source code archive [via kernel.org](#).

As usual, Torvalds and his fellow developers merged all of the major changes for the 3.10 kernel in the two weeks following [the release of version 3.9](#). Linux 3.10 thus now finds itself in the stabilisation phase, where large-scale changes are undertaken only in exceptional cases. Therefore, we are now able to offer a fuller overview of the most important changes to be anticipated in the 3.10 kernel, which is due for release in late June/early July. We will be presenting a series of articles looking at various areas of the kernel.

Kernel Log

COMING IN LINUX 3.10

1: Networking

[Networking enhancements](#)
[Minor gems](#)

THE H OPEN HEADLINES

- The H is closing down
- Hardware Hacks: Fire, alarms, touchable boards and NFC rings
- GitHub gets smart over open source licences
- NSS 3.15.1 brings TLS 1.2 support to Firefox
- Second Android signature attack disclosed
- One month left for the EclipseCon Europe 2013 call for papers

THE H OPEN

Kernel Log: Coming In 3.10 (Part 4) - Drivers

Kernel Log

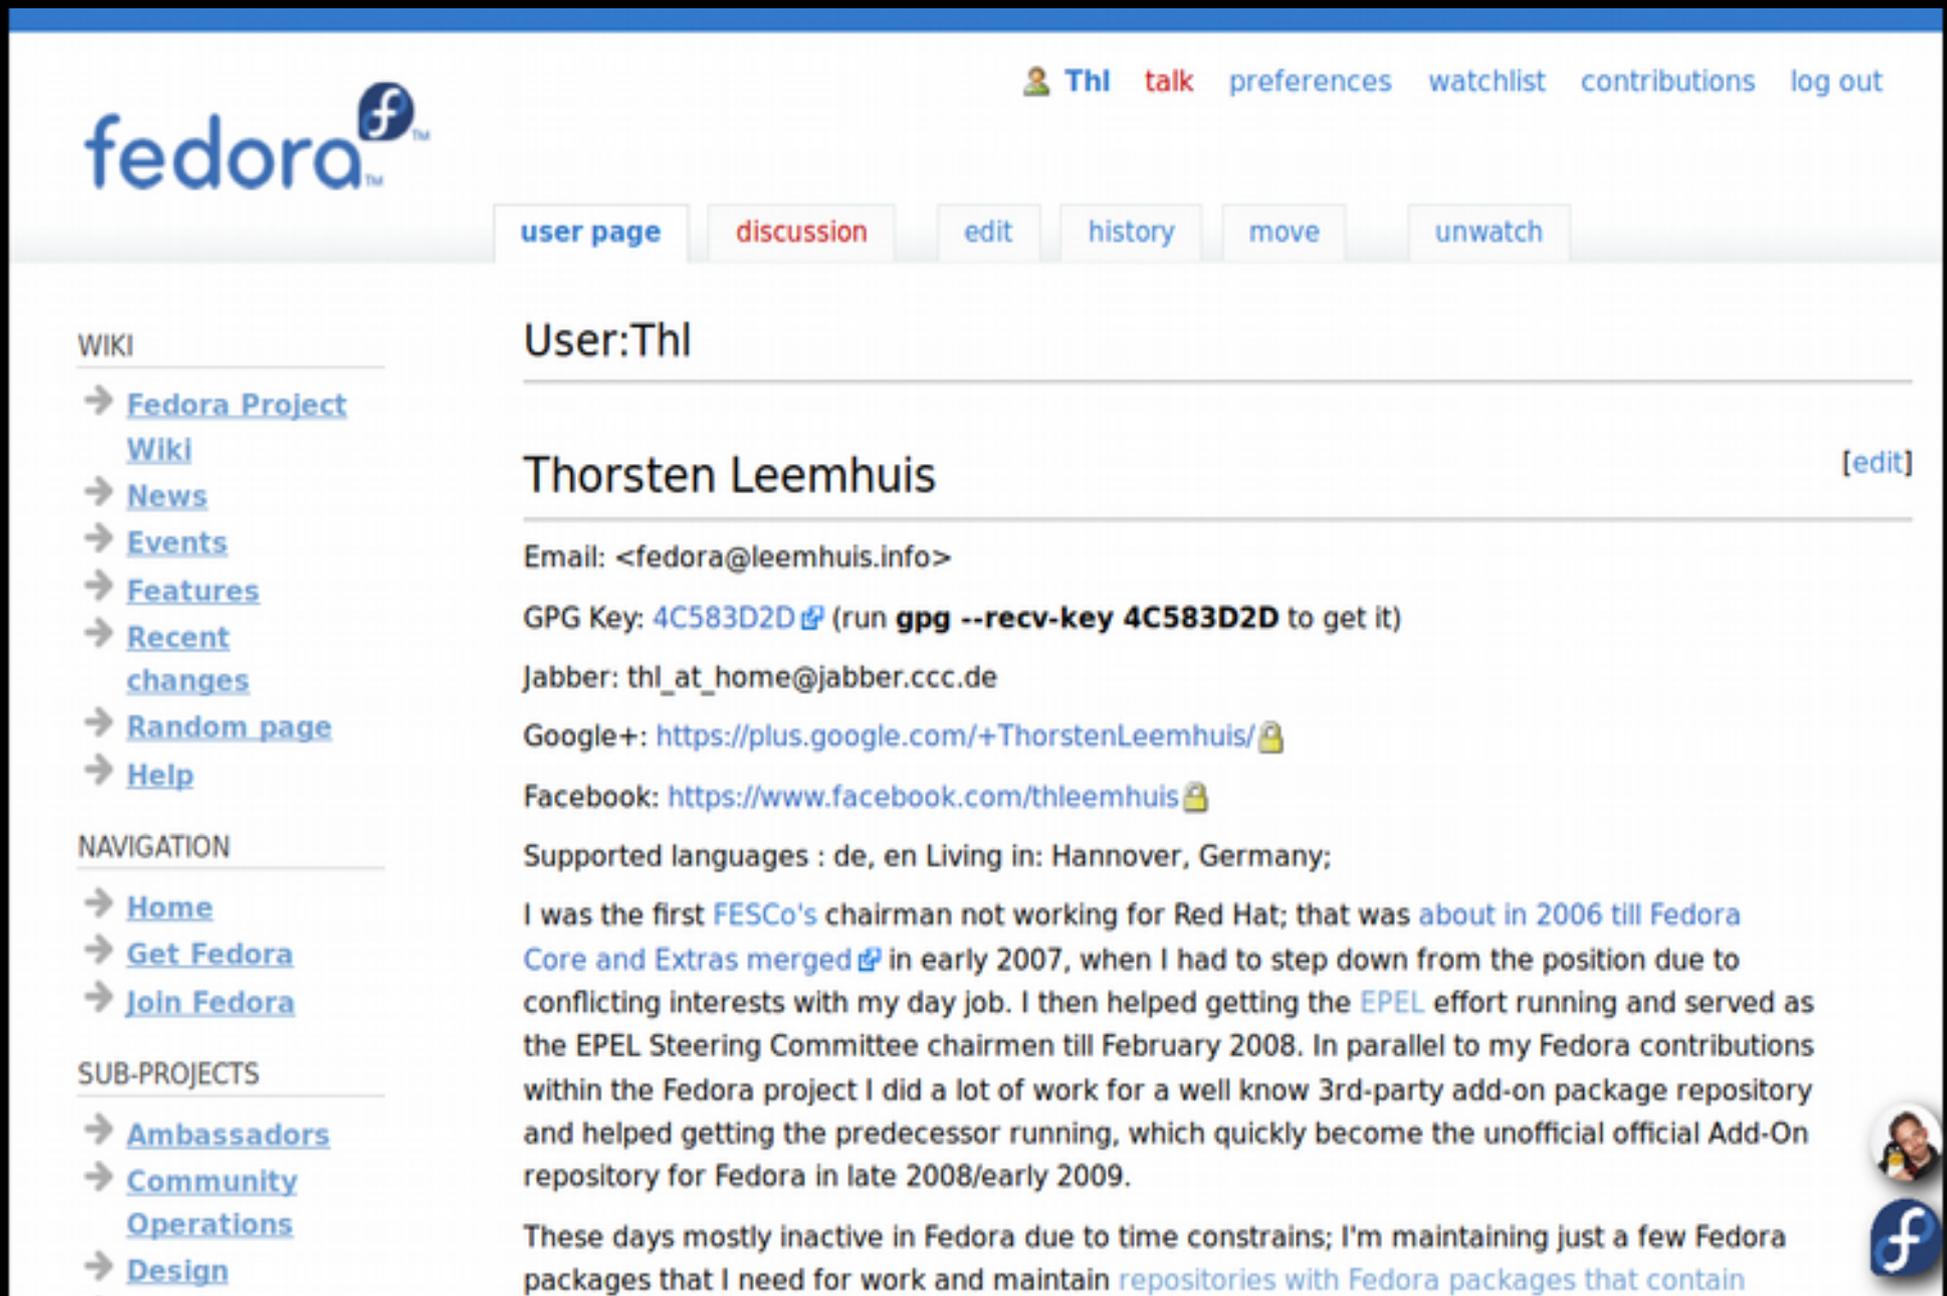
COMING IN LINUX 3.10

4: Drivers

Linux 3.10 will be able to use the video acceleration features offered by Radeon graphics cores. Systems with Intel graphics will wake from standby faster. Linux now has an input device driver for Apple's infrared receiver [more »](#)

closed down July 2013 :-)

%pre | whoami @Fedora



The screenshot shows a Fedora user profile page for 'Thl'. The page has a blue header with the Fedora logo and navigation links: 'Thl', 'talk', 'preferences', 'watchlist', 'contributions', and 'log out'. Below the header is a navigation bar with buttons for 'user page', 'discussion', 'edit', 'history', 'move', and 'unwatch'. The main content area is titled 'User:Thl' and contains a bio for Thorsten Leemhuis, including his email, GPG key, Jabber ID, Google+, and Facebook links. It also lists supported languages and a detailed bio paragraph. A sidebar on the left contains 'WIKI' and 'NAVIGATION' sections with various links. A small circular profile picture and a Facebook icon are visible in the bottom right corner.

fedora™

Thl talk preferences watchlist contributions log out

user page discussion edit history move unwatch

WIKI

- [Fedora Project Wiki](#)
- [News](#)
- [Events](#)
- [Features](#)
- [Recent changes](#)
- [Random page](#)
- [Help](#)

NAVIGATION

- [Home](#)
- [Get Fedora](#)
- [Join Fedora](#)

SUB-PROJECTS

- [Ambassadors](#)
- [Community Operations](#)
- [Design](#)

User:Thl

Thorsten Leemhuis [\[edit\]](#)

Email: <fedora@leemhuis.info>

GPG Key: [4C583D2D](#) (run **gpg --recv-key 4C583D2D** to get it)

Jabber: thl_at_home@jabber.ccc.de

Google+: <https://plus.google.com/+ThorstenLeemhuis/>

Facebook: <https://www.facebook.com/thleemhuis>

Supported languages : de, en Living in: Hannover, Germany;

I was the first FESCo's chairman not working for Red Hat; that was about in 2006 till Fedora Core and Extras merged in early 2007, when I had to step down from the position due to conflicting interests with my day job. I then helped getting the EPEL effort running and served as the EPEL Steering Committee chairmen till February 2008. In parallel to my Fedora contributions within the Fedora project I did a lot of work for a well know 3rd-party add-on package repository and helped getting the predecessor running, which quickly become the unofficial official Add-On repository for Fedora in late 2008/early 2009.

These days mostly inactive in Fedora due to time constrains; I'm maintaining just a few Fedora packages that I need for work and maintain repositories with Fedora packages that contain

%pre | whoami @Fedora

- [Documentation](#)
- [EPEL](#)
- [Infrastructure](#)
- [Internationalization](#)
- [Localization](#)
- [Marketing](#)
- [Magazine](#)
- [Package Maintainers](#)
- [Quality Assurance](#)
- [Websites](#)
- [All projects](#)

SEARCH

Go

Search

TOOLS

- [What links here](#)
- [Related changes](#)
- [Upload file](#)
- [Special pages](#)
- [Printable version](#)
- [Permanent link](#)

Package repository with Linux vanilla kernels for Fedora [\[edit\]](#)

This page contains information about a [set of repositories](#) which contain RPM packages with Linux vanilla kernels built for Fedora. 'Vanilla' in this scope means 'unmodified'. In other words: the sources used to compile those kernels come straight from kernel.org and do not contain any of the enhancements which the official Fedora kernels contain.

How to use these repos [\[edit\]](#)

How to use, the quick (aka TLDR) verison [\[edit\]](#)

Download the definitions for the Kernel vanilla repositories:

```
curl -s https://repos.fedorapeople.org/repos/thl/kernel-vanilla.repo | sudo
```

Run this to get the latest development kernel:

```
sudo dnf --enablerepo=kernel-vanilla-mainline update
```

You don't want to run a development kernel and want the latest stable kernel instead? Then run this:

```
sudo dnf --enablerepo=kernel-vanilla-stable update
```

Reboot. That's it - at least most of the time, as sometimes it's not that easy:

- Is UEFI Secure Boot active on your system? Then you have to disable it in your BIOS Setup to run kernels from these repos, as they are not signed with a key that a default Secure Boot setup



%pre | whoami @linux

not a Linux kernel
developer

%pre | whoami @linux

good overview

%pre | whoami @linux

thanks to my constant
reporting

%pre | whoami @linux

thanks to working as
regression tracker for
Linux 4.7 and 4.8

%pre | whoami @linux

and thx to lwn.net

%pre | this talk

1. short status overview
2. recent developments
3. important changes
4. things in the works
5. meta view on development

%pre | this talk

- 1. short status overview -
*Where we are and what is
happening right now*

%pre | this talk

- 2. recent developments -
*features in the
newest kernel version*

%pre | this talk

- 3. important changes -
*important changes in the
past 12 to 18 months*

%pre | this talk

- 4. work in progress -
trying to look ahead a bit

%pre | this talk

- 5. meta view -

*development pace, growth,
number of developers, current
problems, ...*

%pre | this talk

**BTW: slides will continue
to change this frequently
*(this is already slide #21)***

%pre | this talk

therefore better write
questions down!

part 1:
short status overview

brief overview

everything works
smoothly as usual

this is a good thing!

brief overview

current kernel version:

Linux 4.8

brief overview

released yesterday

brief overview

Ubuntu 16.10 will use it
and many others, too

brief overview

like each new version, it
brings support for
~500 more devices or
device classes

brief overview

100 to 200 of them are for
ACPI-, USB-, or PCI/PCIe
devices

brief overview

that's important because
hardware vendors release
new devices frequently

brief overview

a bunch of new &
improved security and
hardening techniques

quite a lot happens in this area!

brief overview

ETA for Linux 4.9:

December, 5th

or one week later

brief overview

new versions typically get
released every

9 or 10 weeks

one week less than

a few years ago

brief overview

you can bet on it!

*you'll lose sometimes,
but win most of the time*

brief overview

every new version consists
of ~12.500 Commits

brief overview

from about
~1.500 developers

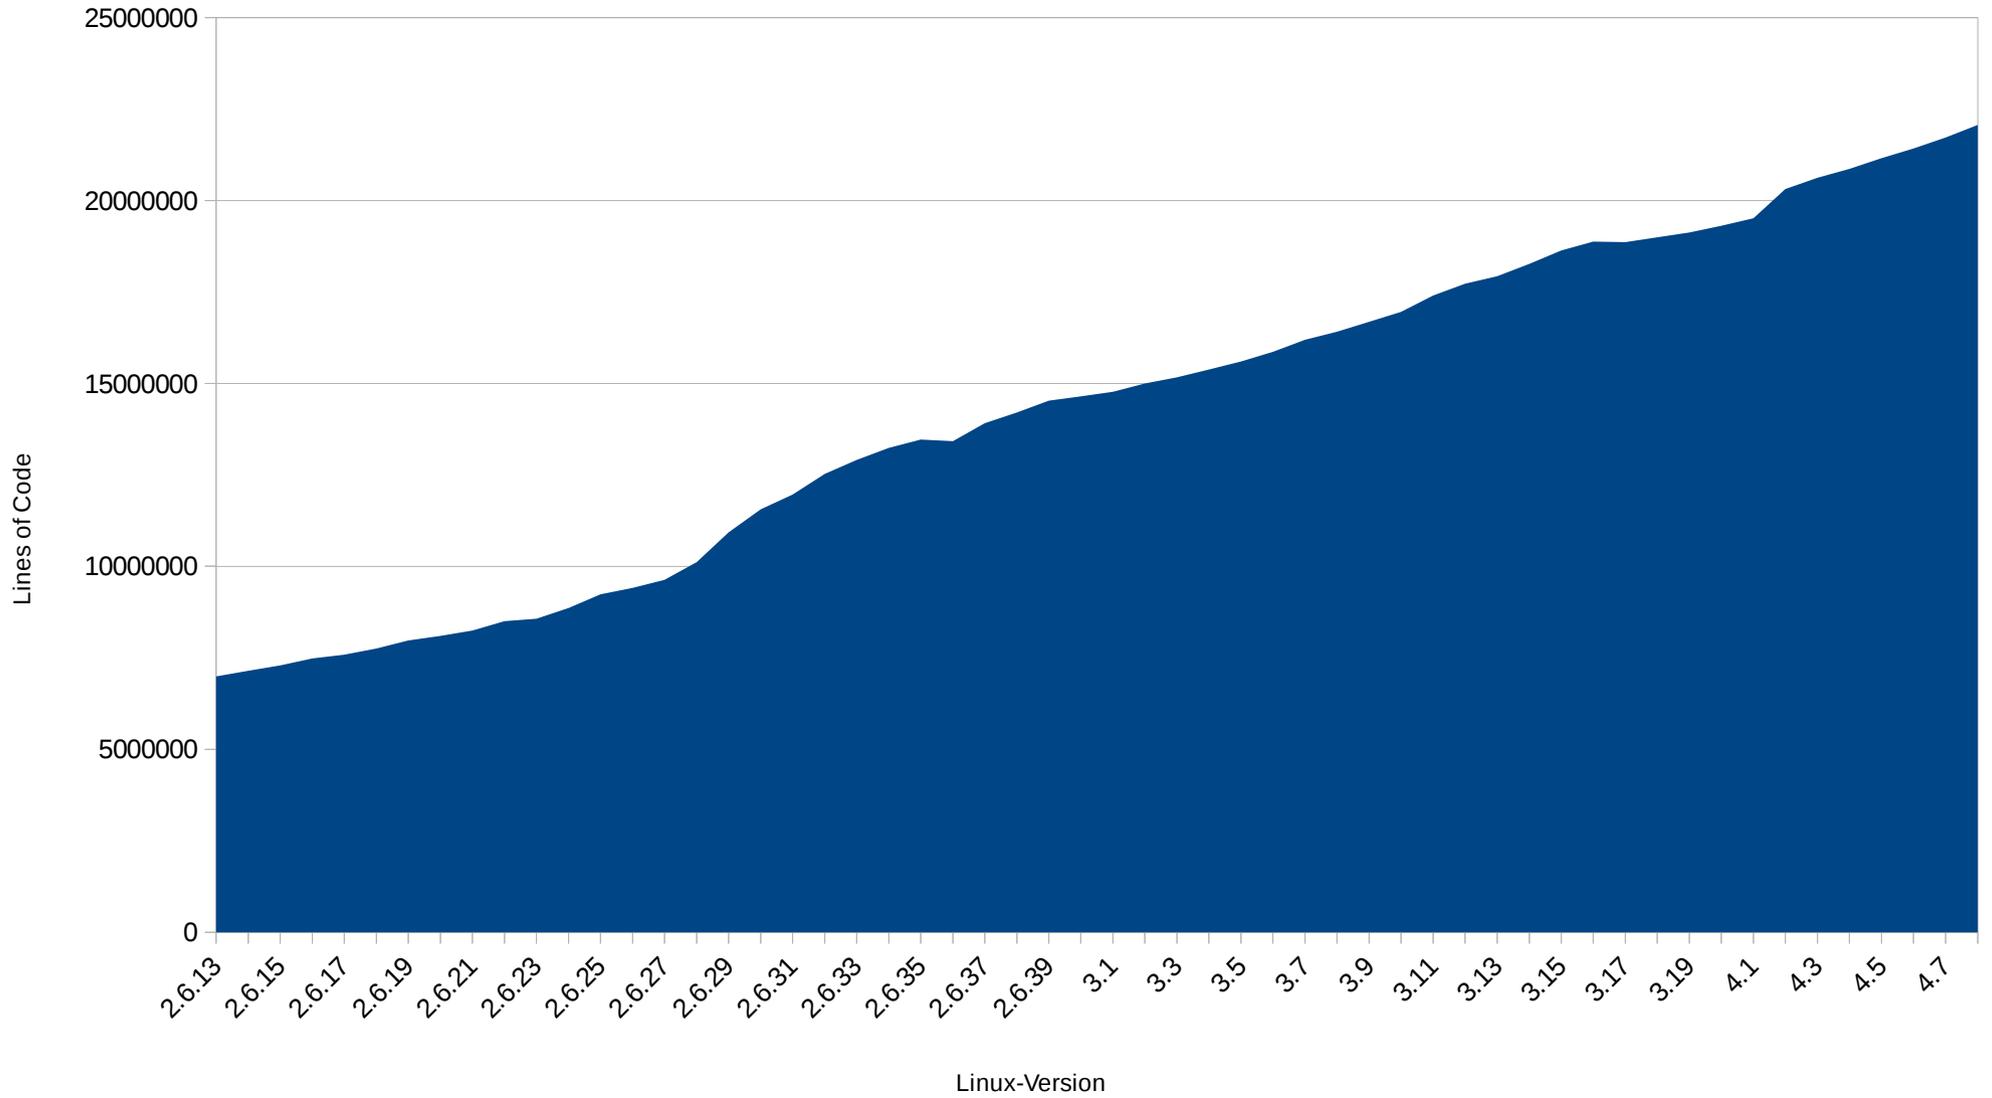
brief overview

containing about
700.000 lines of
new or modified code

brief overview

sources grow by
~290.000 lines
every release

brief overview



brief overview

22 million lines of code

*(including documentation,
comments, blank lines, ...)*

brief overview

between 4.7 and 4.8-rc1

14 days with

12.423 commits containing

647.518 insertions &

307.408 deletions

brief overview

one change every
97 seconds

brief overview

despite that speed,
everything works very well
nothing to worry about

brief overview



Linux Kernel Development

*How Fast It is Going, Who is Doing It,
What They Are Doing and
Who is Sponsoring the Work*

25th Anniversary Edition

A Publication of The Linux Foundation
August 2016

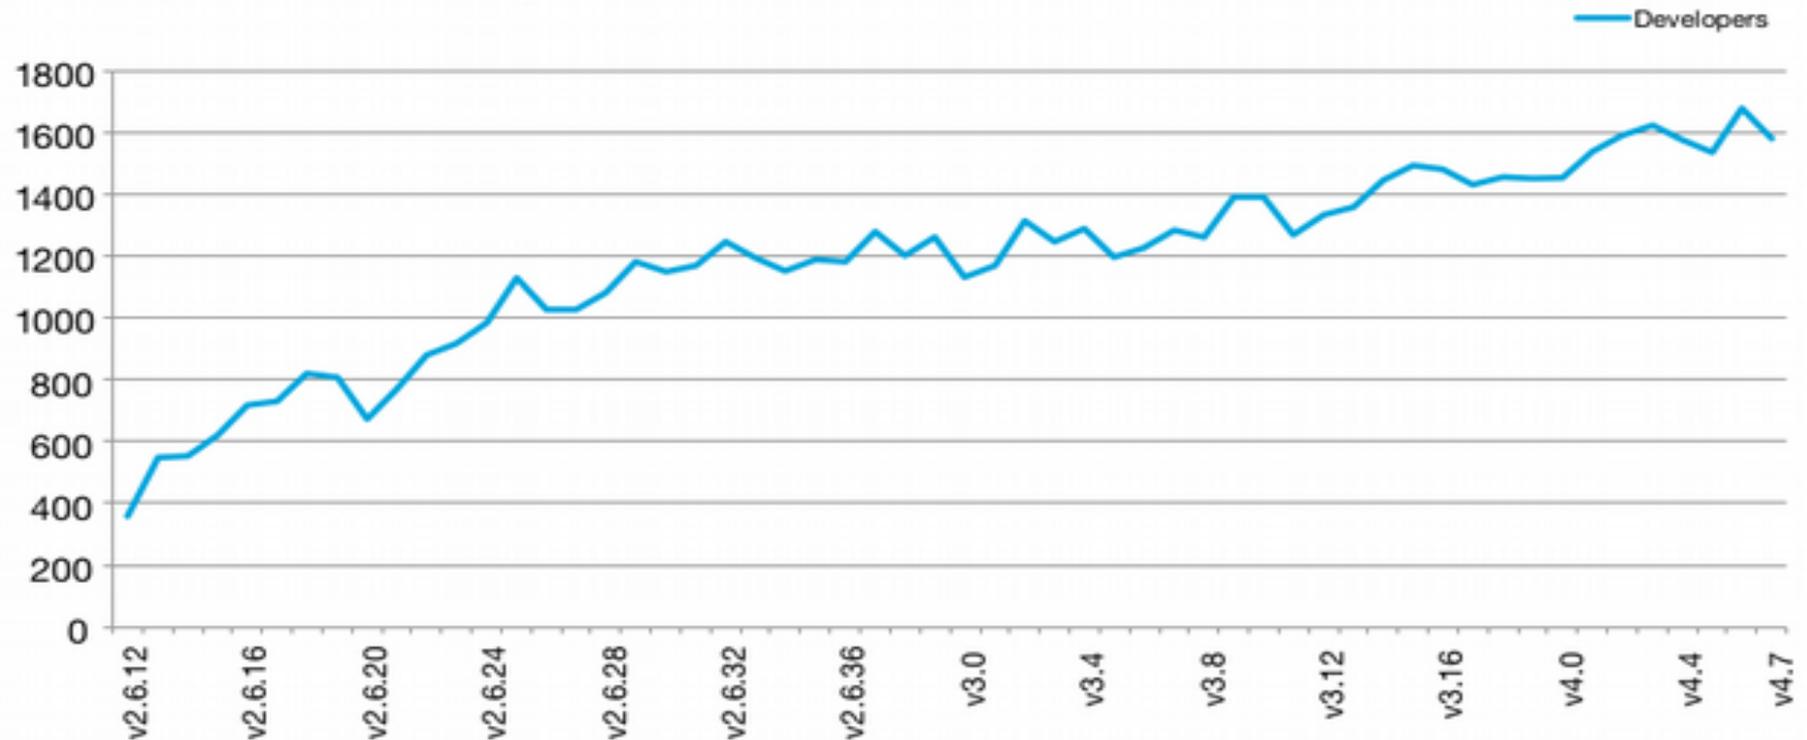
AUTHORS

Jonathan Corbet, LWN.net
Greg Kroah-Hartman, The Linux Foundation

www.linuxfoundation.org

brief overview

Developers Contributing to Each Kernel Version



Since the beginning of the Git era (the 2.6.11 release in 2005), a total of 13,594 developers have contributed to the Linux kernel; those developers worked for a minimum of 1,340 companies. The number of companies supporting work on the kernel appears to be stable and not growing like the number of developers — but the 251 companies observed to have supported work on 4.2 was an all-time record.

brief overview



News from the source

Content
[Weekly Edition](#)
[Archives](#)
[Search](#)
[Kernel](#)
[Security](#)
[Distributions](#)
[Events calendar](#)
[Unread comments](#)

[LWN FAQ](#)
[Write for us](#)

Edition
[Return to the
Kernel page](#)

A look at the 4.8 development cycle

[LWN subscriber-only content]

By **Jonathan Corbet**
September 28, 2016

As of this writing, the 4.8 development cycle is nearing its end. Linus has let it be known that a relatively unusual [-rc8 release candidate](#) will be required before the final release, but that still means that the cycle will only require 70 days, fitting into the usual pattern. A look at the development statistics for this release also fits the pattern about now.

With regard to the release cycle, it has become boringly regular in recent years. The 3.8 kernel, released on February 18, 2013, came out on a Sunday, as has every subsequent release with the exception of 3.11, which was released on Monday, September 2, 2013. In these last few years, the only cycle that has taken longer than 70 days was 3.13, which required 77 days. The extra week that time around was forced by Linus's travels, rather than anything inherent in that cycle itself. Since then, every cycle has taken 63 or 70 days, with the sole exception of 3.16, which showed up in 56 (and one could quibble that it was really a 63-day cycle as well — that was the time Linus experimented with opening the merge window before the previous final release had been made).

In this 70-day cycle, we have seen the addition of 13,253 non-merge changesets from 1,578 developers — so far; the numbers will increase slightly before the end. It is thus a busy cycle, though the record for the busiest (3.15, with 13,722 commits) remains unchallenged. Those developers grew the kernel by 350,000 lines this time around. The most active developers in this cycle were:

lwn.net/Articles/701650/

brief overview



News from the source

Content

[Weekly Edition](#)

[Archives](#)

[Search](#)

[Kernel](#)

[Security](#)

[Distributions](#)

[Events calendar](#)

[Unread comments](#)

[LWN FAQ](#)

[Write for us](#)

Edition

[Return to the](#)

[Kernel page](#)

Work on the 4.8 kernel was supported by 217 employers that we were able to identify. The most active employers this time around were:

Most active 4.8 employers

By changesets			By lines changed		
Intel	1960	14.8%	Samsung	120693	14.4%
Red Hat	1143	8.6%	Intel	104291	12.4%
(Unknown)	806	6.1%	(None)	102848	12.3%
(None)	746	5.6%	Red Hat	48563	5.8%
Linaro	662	5.0%	IBM	42298	5.0%
IBM	654	4.9%	Mellanox	29226	3.5%
Samsung	637	4.8%	(Unknown)	27671	3.3%
SUSE	338	2.6%	Linaro	22960	2.7%
Google	294	2.2%	Broadcom	18040	2.2%
AMD	281	2.1%	Cisco	17868	2.1%
Oracle	259	2.0%	MediaTek	16292	1.9%
Texas Instruments	258	1.9%	QLogic	15986	1.9%
Mellanox	243	1.8%	ARM	14397	1.7%
Renesas Electronics	223	1.7%	Renesas	14283	1.7%

lwn.net/Articles/701650/

that was part 1

next up

part 2:
recent developments

recently | Linux 4.7

Linux 4.8

recently | Linux 4.8

XFS RMAP

Intel vGPU

driver for GeForce 1000

eXpress Data Path (XDP)

recently | Linux 4.8

the first three
are still WIP :-/

*(and that's why they get
mentioned later in this talk)*

recently | Linux 4.8

and even XDP is still in its
early phase

recently | Linux 4.8

XDP is a programmable
and high performance
networking data path

recently | Linux 4.8

IOW: a bypass within the kernel to handle some of the network traffic

recently | Linux 4.8

picks up packages right
after the NIC got them

recently | Linux 4.8

and process them with BPF
programs you write

recently | Linux 4.8

less features

recently | Linux 4.8

but handles some tasks
quicker and with
less overhead

recently | Linux 4.8

supported right now:
early drop,
forward & rewrite

recently | Linux 4.8

DDoS mitigation/defense
more efficient forwarding

recently | Linux 4.8

not a kernel bypass,
like DPDK

recently | Linux 4.8

Advantages of XDP over DPDK

- Allows option of busy polling or interrupt driven networking
- No need to allocate huge pages
- Dedicated CPUs are not required, user has many options on how to structure the work between CPUs
- No need to inject packets into the kernel from a third party userspace application
- No special hardware requirements
- No need to define a new security model for accessing networking HW
- No third party code/licensing required

recently | Linux 4.8

4.8: what else?

recently | 4.8: even more

optimized WLAN transfers

improved Raspi 3 support

architecture emulation

containers

drivers for new graphic cores

from ARM, Intel, Nvidia

recently | 4.8: even more

Software driver for RDMA over
Converged Ethernet (RoCE)

bunch of security improvements

new tools and formats for
writing kernel documentation

recently | 4.8: even more

more details about 4.8:

heise.de/-3283402

lwn.net/Kernel/Index/#Releases-4.8

kernelnewbies.org/Linux_4.8 (WIP)

recently | 4.8: even more

4.8 will soon be used in
Arch Linux, Ubuntu 16.10,
OpenSuse Tumbleweed

*within a few weeks also in
Fedora 23/24/25...*

that was part 2

next up

part 3:
important changes

important changes

Kernel Self Protection /
kernel hardening

changes | kernel hardening

security bugs just happen

changes | kernel hardening

it often takes 5 years
until they are found

changes | kernel hardening

need to protect the kernel
against common flaws

changes | kernel hardening

just like today's cars
are prepared for a
potential crash

changes | kernel hardening

the Kernel Self Protection
Projekt (aka KSPP) drives
these efforts

changes | kernel hardening

they improved or merged
a big bunch of security and
hardening techniques

changes | kernel hardening

KASLR improvements,
SLUB freelist randomization,
GCC Plugins,
Hardened usercopy,
ro_after_init for modules,
improved RNG

changes | kernel hardening

that's just what's
new in 4.8!

*note: (1) most, but not all were
driven by the KSPP; (2) some
were ported from grsecurity/PAX*

changes | kernel hardening

many similar things
in earlier versions

changes | kernel hardening

more to come

Kernel stack hardening in 4.9?

changes | kernel hardening

BTW: yes, this hardening
sometimes has a negative
impact on performance

changes | kernel hardening

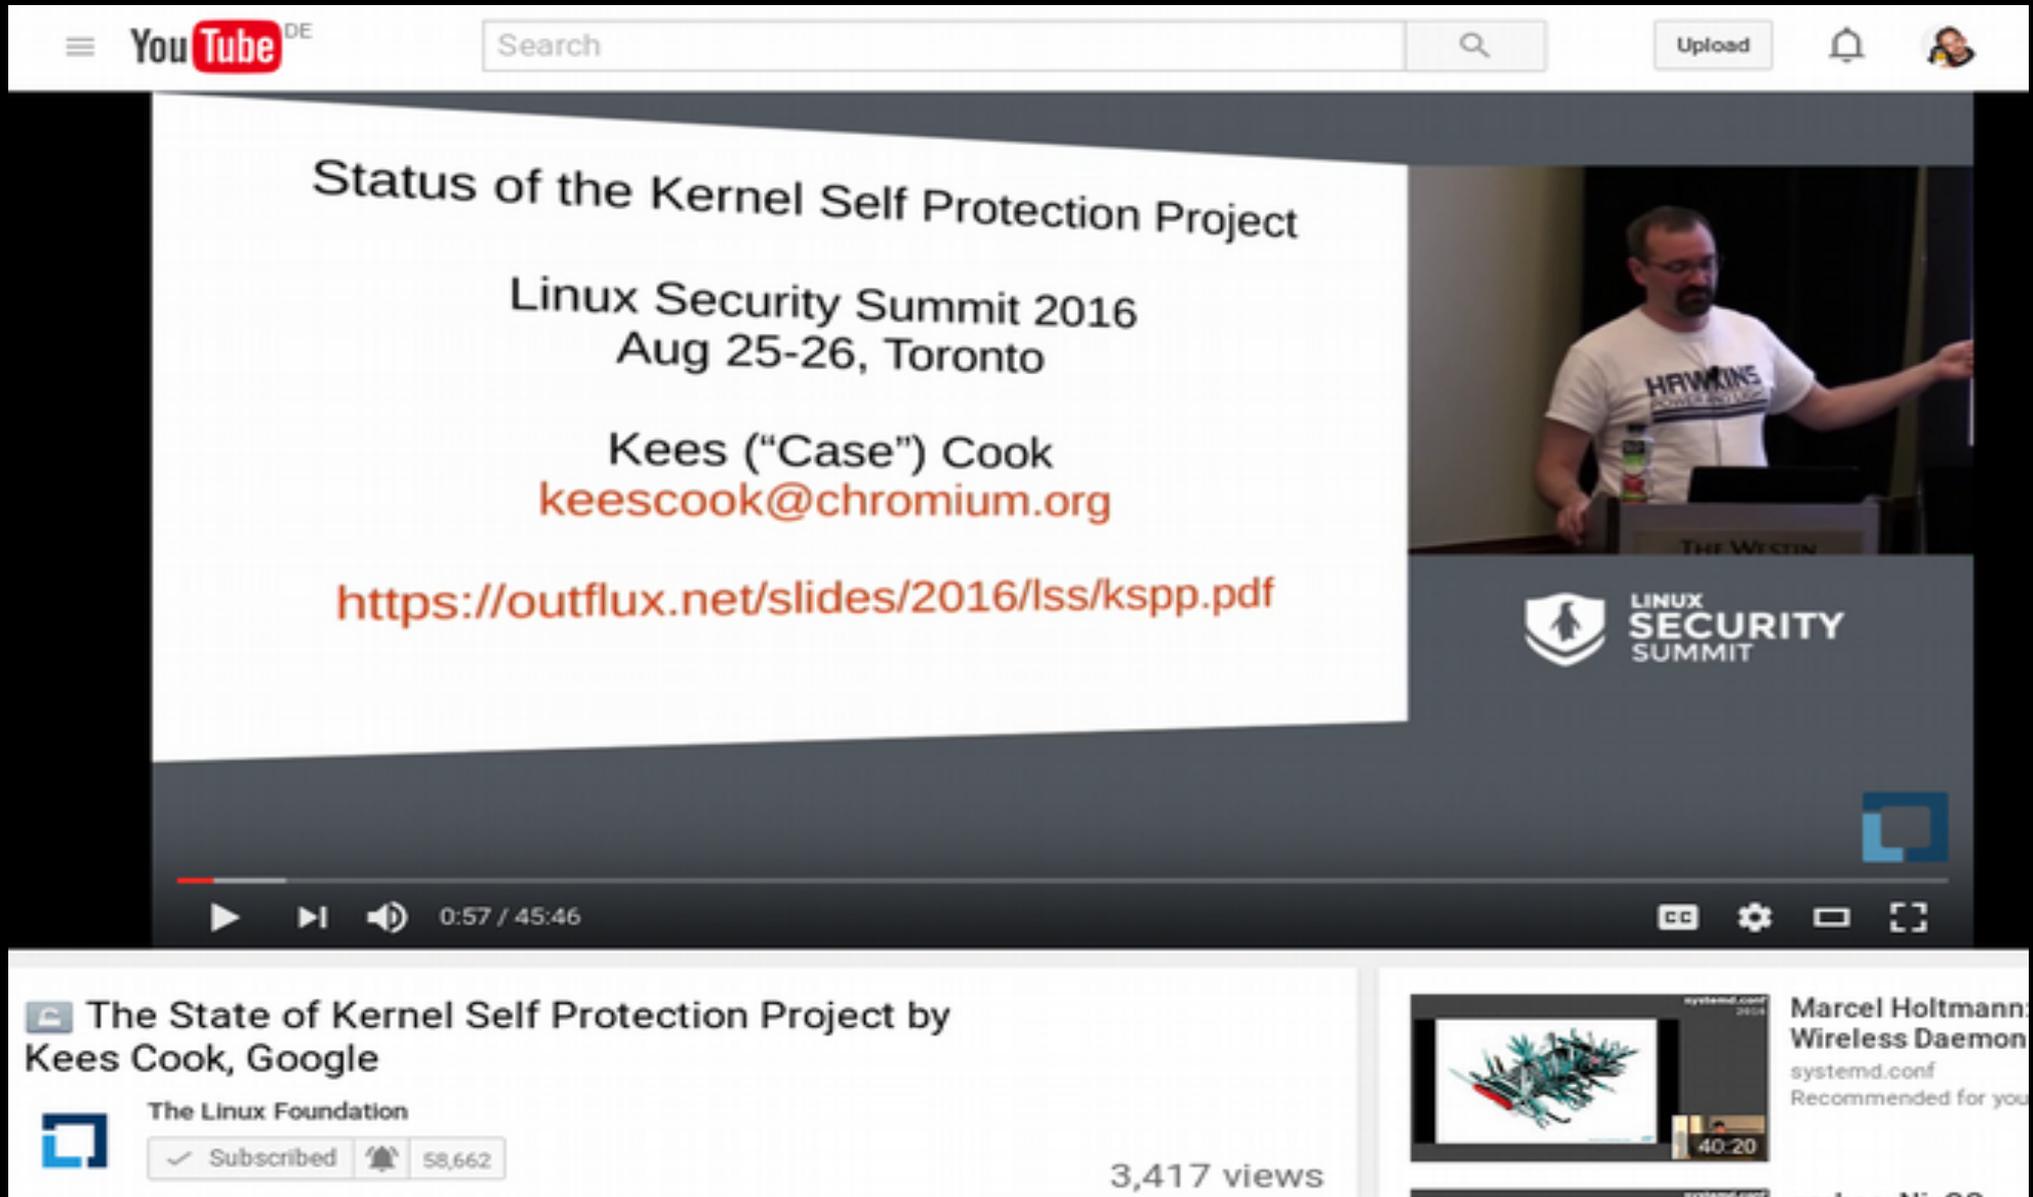
just like protections
techniques make cars
heavier and thus slower

get used to it

changes | kernel hardening

more details in videos and
slides from the recently
held security summit

changes | kernel hardening



The image shows a YouTube video player interface. The video content is a presentation slide titled "Status of the Kernel Self Protection Project". The slide text includes: "Linux Security Summit 2016 Aug 25-26, Toronto", "Kees ('Case') Cook", "keescook@chromium.org", and a URL "https://outflux.net/slides/2016/lss/kspp.pdf". The video player shows a progress bar at 0:57 / 45:46. Below the video, the title "The State of Kernel Self Protection Project by Kees Cook, Google" is visible, along with the channel name "The Linux Foundation" and a subscriber count of 58,662. The view count is 3,417 views. To the right, a recommended video thumbnail for "Marcel Holtmann: Wireless Daemon" is shown.

YouTube DE Search Upload

Status of the Kernel Self Protection Project

Linux Security Summit 2016
Aug 25-26, Toronto

Kees ("Case") Cook
keescook@chromium.org

<https://outflux.net/slides/2016/lss/kspp.pdf>

LINUX SECURITY SUMMIT

0:57 / 45:46

The State of Kernel Self Protection Project by Kees Cook, Google

The Linux Foundation

Subscribed 58,662

3,417 views

Marcel Holtmann: Wireless Daemon
systemd.conf
Recommended for you

www.youtube.com/watch?v=aMkCKeZ8xZw

important changes

what else?

important changes

BPF got improved and
became more important
(and there is more to come)

WIP | BPF

Berkeley Packet Filter/BPF:
What tcpdump uses to just
get the relevant packages

WIP | BPF

was improved (Extended
BPF/eBPF) and became
flexible & programmable
in-kernel VM

WIP | BPF

VM == "abstract
computing machine"

*(think of a Java VM
and not KVM/XEN/VMware/...)*

WIP | BPF

making the BPF more
capable and useful

(the e in eBPF got dropped lately)

WIP | BPF

BPF is used in more
and more areas
*to improve things or
realize new features*

WIP | BPF

network traffic control,
performance monitoring,
tracing, syscall filtering, ...

WIP | BPF

for example used in the
eXpress Data Path/XDP

WIP | BPF

dynamically analyzing
performance improved a

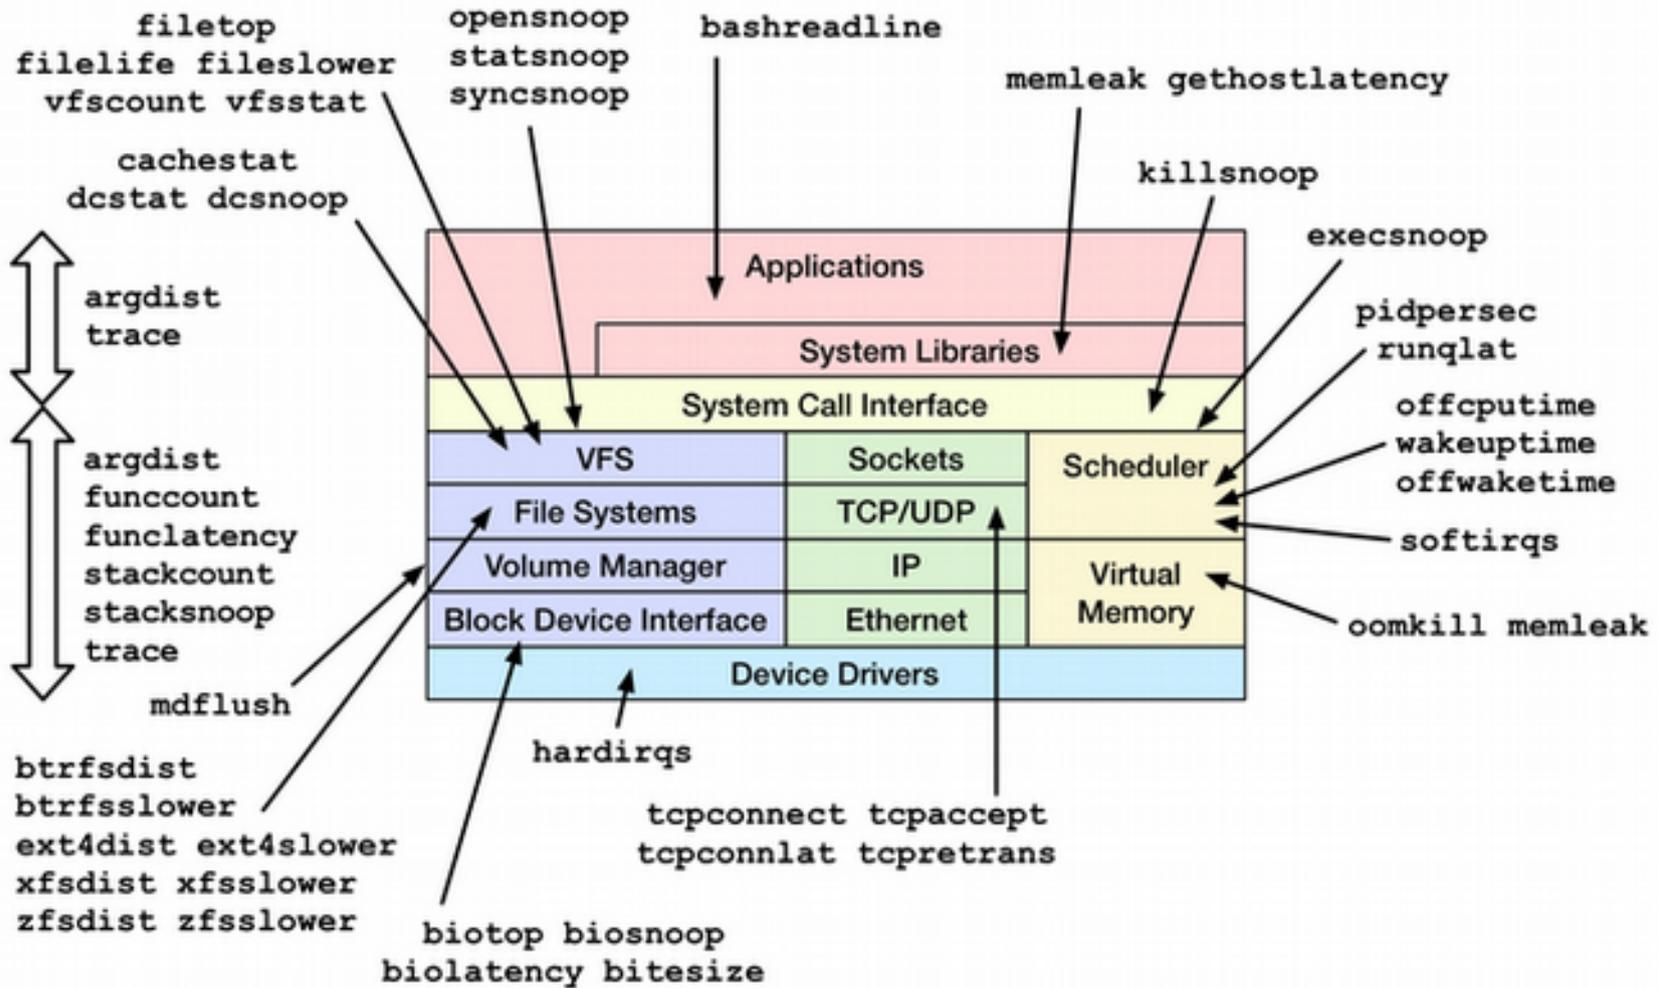
lot thx to eBPF

became/becomes

more D-Trace like

WIP | BPF

Linux bcc/BPF Tracing Tools



<https://github.com/iovisor/bcc#tools> 2016

github.com/iovisor/bcc

WIP | BPF

Linux BPF Superpowers

05 Mar 2016

Last month I spoke at Facebook's [Performance @Scale](#) event about Linux BPF Superpowers. These are coming to Linux in the 4.x series, and I've been using them in new open source performance tools.

Video is on [Facebook](#) (30 mins):

[Linux 4.x Performance: Using BPF Superpowers \(Brendan Gregg\)](#)

Posted by [At Scale](#) on Friday, February 26, 2016

Slides are on [slideshare](#):



important changes

what else?

important changes

the graphics drivers
improved quite a lot

changes | GPU

not only those in the
kernel, but and the
userland drivers which are
build on top of it

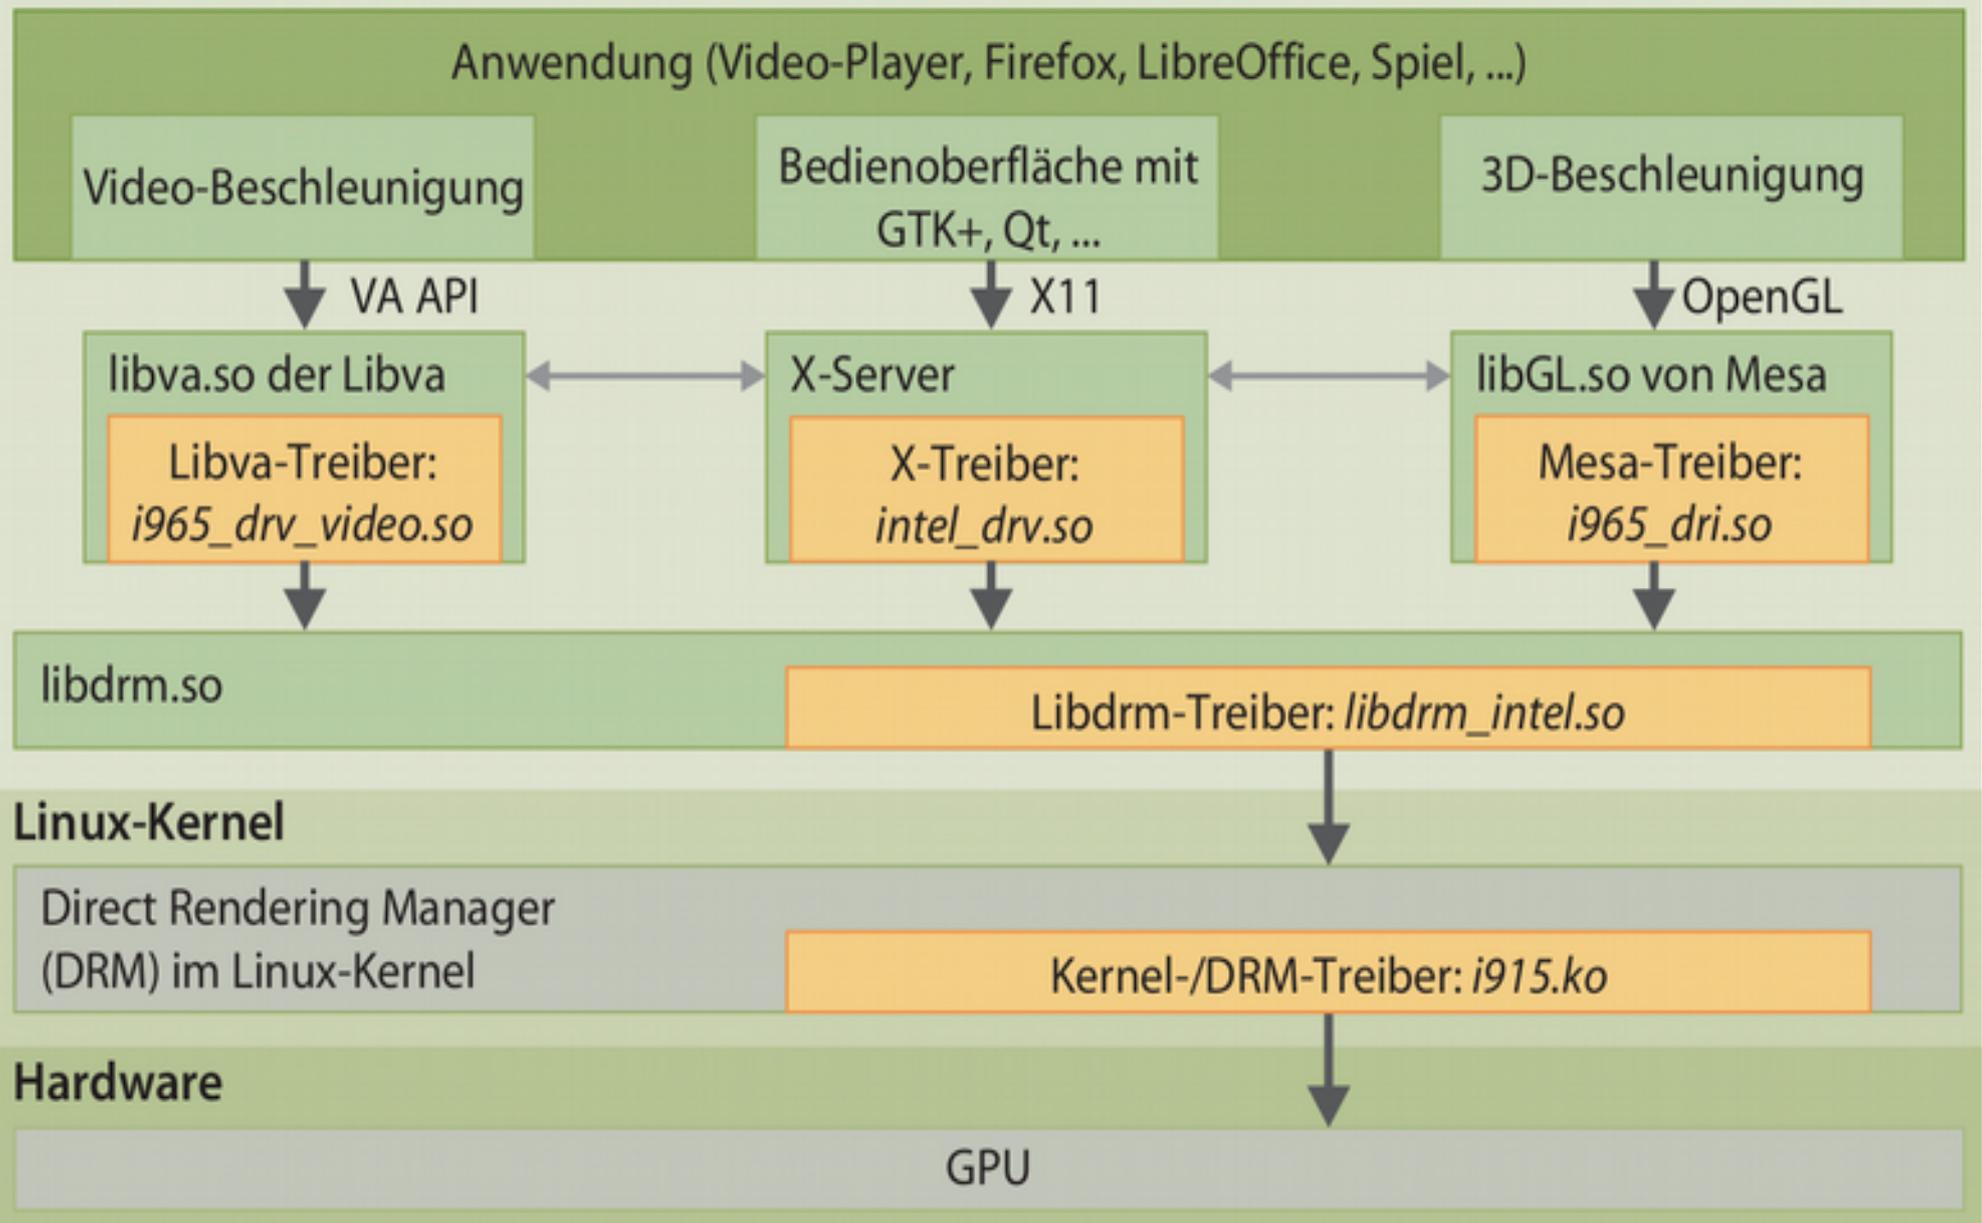
changes | GPU

worth its own talk :-/

(what follows is a brief overview)

changes | GPU

Userspace



changes | GPU

that's why some of the
improvements I'll mention
are in userland drivers

changes | GPU

Intel GPUs

changes | GPU Intel

finally supports OpenGL
4.3 with Mesa 12.0

(run more games, #yeah!)

changes | GPU Intel

Mesa 12.1 (13.0) will soon
bring support for 4.4 or 4.5

changes | GPU Intel

Intel's drivers are not as good as their reputation

(Skylake support for example came late and had stability problems in the beginning)

changes | GPU

next up

changes | GPU

AMD Radeon

changes | GPU AMD

FLOSS drivers got
a whole lot better in the
past two years!

(thx AMD!)

changes | GPU AMD

most things just work

*Video acceleration, power
management, DP MST, HiDPI,
HDMI&DP Audio, ...*

changes | GPU AMD

and 3D performance
improved quite a lot

changes | GPU AMD

sometimes it's close to or
even better than AMDs old
proprietary drivers
(for the older GPUs)

changes | GPU AMD

Mesa 12.0 brought
OpenGL 4.3 support, too

changes | GPU AMD

4.7 brought support for
the latest GPUs

Rx 400 series

changes | GPU AMD

IOW: AMDs FLOSS drivers
are more than good
enough for most things
and better than their reputation

changes | GPU AMD

some problems as well

DAL, performance on newer

GPUs, ...

changes | GPU

next up

changes | GPU

Nvidia GeForce

changes | GPU Nvidia

FLOSS drivers still slow
and lack a lot of features
*nevertheless they are good
enough for some use cases*

changes | GPU Nvidia

it didn't help much that
Nvidia participates a little
in Nouveau development
(since ~2 years now thx to Tegra)

changes | GPU Nvidia

acceleration on recent
GPUs requires a
signed firmware

900 and 1000er series

changes | GPU Nvidia

took more than a year for
the 900 series

changes | GPU Nvidia

support merged
in Linux 4.6

changes | GPU Nvidia

**4.8 brought a driver for
GeForce 1070 & 1080**

changes | GPU Nvidia

still waiting for Nvidia to
release the Firmware :-/

important changes

what else?

changes | Btrfs

BTRFS:

no big changes recently

changes | Btrfs

but a few small
improvements and fixes

changes | Btrfs

is this "Next generation
file system for Linux"
finally stable?

changes | Btrfs

that questions is a bit like
"can I swim here?"

changes | Btrfs

the answer depends on
the local conditions

changes | Btrfs

but also on your
swimming abilities

changes | Btrfs

for Ext4 it's a bit like this

changes | Btrfs



changes | Btrfs

for Btrfs it's more like this

changes | Btrfs



changes | Btrfs

that's for the common usage of Btrfs; there are a few areas more like this

changes | Btrfs



changes | Btrfs

as some parts of Btrfs are
still WIP and unstable

changes | Btrfs

for example the
RAID 5/6 support

changes | Btrfs

Feature	Status	Notes
Subvolumes, snapshots	OK	
Trim (aka. discard)	OK	<i>fstrim</i> and mounted with <i>-o discard</i> (has performance implications)
Autodefrag	OK	
Defrag	mostly OK	extents get unshared (<i>see below</i>)
Auto-repair	OK	automatically repair from a correct spare copy if possible (dup, raid1, raid10)
Compression	mostly OK	(<i>needs verification and source</i>) auto-repair and compression may crash
Scrub	OK	
Scrub + RAID56	Unstable	will verify but not repair
Filesystem resize	OK	shrink, grow
Send	OK	corner cases may still exist
Receive	OK	
Single (block group profile)	OK	
DUP (block group profile)	OK	
RAID0	OK	
RAID1	OK	
RAID10	OK	
RAID56	Unstable	write hole still exists, parity not checksummed
Seeding	OK	should be better documented

btrfs.wiki.kernel.org/index.php/Status

changes | Btrfs

one of the reasons why
Btrfs still requires a
learning phase

changes | Btrfs

which is only worth it
if you benefit from Btrfs
features

important changes

what else?

important changes

quite a lot of improvements in the network area
(way to many to mention the all)

changes | network

some made Linux ready
for 100 Gigabit Ethernet
(just one example)

changes | network

a lot of improvements
for tunneling

#cloud

changes | network

nft/nftables

matures slowly

(iptables successor)

important changes

what else?

changes

virtio-gpu/Virgl 3D

changes | virtio-gpu

3D in Linux VMs running
on a Linux host

changes | virtio-gpu

3D driver in the guest
passed commands to the
host which executes them
with its 3D driver

changes | virtio-gpu

everything required in
the latest distros

*(e.g. Linux >= 4.4 &&
Mesa >= 11.1 &&
Qemu >= 2.5 && ...)*

changes

what else?

changes

Raspberry-Pi-Support

changes | raspi

out of the box support for
the various raspi variants
improved a lot

changes | raspi

and we now have floss
drivers in kernel and Mesa
called "vc4"

changes

what else?

changes

a lot of other things

changes | much more

copy offloading for NFS

OOM killer got optimized

ARM64 support matures

parallel NFS (pnfs)

Cgroup v2

changes | much more

NVDIMMs and other types of
Persistent Memory (pmem)

CRIU (Checkpoint & Restart)

Cgroup Namespaces

userfaultfd: Post-copy live
migration for Qemu/KVM

that was part 3

next up

part 4:
things in the works
(aka "WIP")

WIP

what the next few
month might bring

WIP

not kdbus

it's dead, Jim

WIP

developers started
"Bus1" instead

WIP | Bus1

still pre-alpha and not
officially announced

WIP | Bus1

design inspired by
Androids "binder" and
other IPC techniques

WIP | Bus1

universal IPC

*can transport D-Bus messages,
but can also be used for other
protocols*

WIP

what else?

WIP

Intels Graphics
Virtualization Technology
(GVT-g)

WIP

use parts of your iGPU
in VMs

WIP

3D acceleration in VMs
(Linux and Windows)

WIP

initial bits in 4.8

WIP

what else?

WIP

XFS Reverse-mapping

WIP | XFS RMAP

also in 4.8, but not used
for anything yet

WIP | XFS RMAP

foundation for many
improvements currently
developed

WIP | XFS RMAP

of free extents. I.e. it is coherent with the free space trees we already maintain and never overlaps with them.

This reverse mapping infrastructure is the building block of several upcoming features - reflink, copy-on-write data, dedupe, online metadata and data scrubbing, highly accurate bad sector/data loss reporting to users, and significantly improved reconstruction of damaged and corrupted filesystems. There's a lot of new stuff coming along in the next couple of cycles, and it all builds in the rmap infrastructure.

As such, it's a huge chunk of new code with new on-disk format features and internal infrastructure. It warns at mount time as an

WIP | XFS RMAP

IOV: many features that people like in Btrfs & ZFS
(but integrated volume management is not among them)

WIP | XFS RMAP

so will XFS in the end be
our "next gen file system
for Linux"?

WIP | XFS RMAP

we'll see, but it certainly
will be interesting to watch

WIP

what else?

WIP

MD-RAID:

Clustering, Log-Device

WIP | Mdadm

clustering support for
MDDRAID1 slowly getting
production ready

WIP | Mdadm

Log-Device:
Journaling for RAID 5/6

WIP | Mdadm

avoids the "write hole"

*and improves the performance a
little bit; more to come*

WIP

what else?

WIP

a lot more

WIP | much more

live patching
improvements

year 2038

realtime &
realtime KVM

WIP | much more

richacl

support for SMR hard discs

Tinification

schedutil

that was part 4

next up

Teil 5: meta view

meta

a look at kernel
development itself

meta

a new longterm kernel
every January

(longterm == 2 years support)

meta | Longterm

makes it more predictable
(for users and distros)

meta | Longterm

Linux Kernel Monkey Log

Random bits from Greg Kroah-Hartman

Blog Archives

Search

RSS

SEP 6TH, 2016

4.9 == Next LTS Kernel

As I [briefly mentioned](#) a few weeks ago on my G+ page, the plan is for the 4.9 Linux kernel release to be the next “Long Term Supported” (LTS) kernel.

Last year, at the Linux Kernel Summit, [we discussed](#) just how to pick the LTS kernel. Many years ago, we tried to let everyone know ahead of time what the kernel version would be, but that caused a lot of problems as people threw crud in there that really wasn't ready to be merged, just to make it easier for their “day job”. That was many years ago, and people insist they aren't going to do this again, so let's see what happens.

I reserve the right to not pick 4.9 and support it for two years, if it's a major pain because people abused this notice. If so, I'll possibly drop back to 4.8, or just wait for 4.10 to be released. I'll let everyone know by updating the [kernel.org releases page](#) when it's time
(many months from now.)

Recent Posts

[4.9 == Next LTS Kernel](#)

[Kdbus Details](#)

[Binary Blobs to C Structures](#)

[Bootting a Self-signed Linux Kernel](#)

[Longterm Kernel 3.10](#)

About Greg

Greg is a Fellow at the [Linux Foundation](#) and is responsible for the Linux kernel stable releases. He is also the maintainer of a variety of different kernel subsystems (USB, char/misc, staging, etc.) and has written a few books about Linux kernel development.

meta view

what else?

meta view

there are about
two or three security
issues every week

meta | security

and each year there are a
few that are really critical

meta | security

better be prepared!

meta | security

the big distros know that

meta | security

many vendors do not
understand this yet:-/

(esp. in the embedded world)

meta | security

that will become a big
problem with the
Internet-of-Things (IoT)

meta | security

The image shows a screenshot of a Twitter post from Octave Klaba / Oles (@olesovhcom) dated September 19, 2016. The tweet reports a DDoS attack: "@Dominik28111 we got 2 huge multi DDoS: 1156Gbps then 901Gbps". A screenshot of a terminal window is embedded in the tweet, showing a list of IP addresses and their corresponding bandwidths in pps (packets per second) and bps (bits per second). The list includes:

- 141101 | 7991000pps | 92120371320bps
- 141822 | 961266pps | 10164065688bps
- 7039 | 36447333pps | 310431776768bps
- | 11518142pps | 98140493136bps
- 900 | 3450300pps | 29380814296bps
- 040 | 22434666pps | 191048318976bps
- 007039 | 93766762pps | 799069437952bps
- 41900 | 3450300pps | 29380814296bps
- 92 | 16026379pps | 136649443464bps
- 7045 | 25634000pps | 218305615184bps
- | 11529383pps | 98233078032bps
- 959 | 7555266pps | 64350800832bps
- 044 | 14566000pps | 124009818792bps
- 007045 | 72241333pps | 615385180840bps
- 41959 | 7555266pps | 64350800832bps
- 51 | 11529383pps | 98233078032bps

The tweet has 155 retweets and 136 likes. The user's profile information includes: "founder, owner, chairman, cto #Ovh. marié et père de 1+2 filles.", location "Roubaix Valley", website "ovh.com", "280K Vine Loops", and "Joined May 2011".

it seems it already is :-/

meta view

what else?

meta view

GPL enforcement

meta view | GPL enforcement

a long debate recently

indirectly triggered by

"hch vs VMware"

meta view | GPL enforcement

some of the most
important kernel
developers are against it

meta view | GPL enforcement

will be a topic on this years
kernel summit

meta view

what else?

meta view

stable & longterm kernels
not working as hoped

meta view | stable issues

(1) to many regressions?

meta view | stable issues

(2) ignored by
manufactures of
embedded hardware?

meta view | stable issues

many recent android
phones running quite
old kernels

meta view | stable issues

this is a mess and been
like that for a long time

meta view | stable issues

kernel developer once
again look how matters
can be improved

meta view

what else?

meta view

testing and regression
tracking are a problem

my pet peeve ;-)

meta view | testing

bug reporting is hard

bugzilla.kernel.org

doesn't work to well

to hard to run mainline

scares testers away

meta view | testing

lot's of room for
improvements and
community participation

meta view

what else?

meta view

more and more tools are
developed and used to
find bugs early

meta | bug hunt tools

fuzzing tools:

Trinity, Kcov/Syzkaller

meta | bug hunt tools

CI: kernelci.org,
kbuild test robot aka "0-
day"

meta | bug hunt tools

Kselftest, GCC plugins

meta view

what else?

meta view

new Linux competitors
emerging for use in IoT

e.g. Zephyr & Magenta

meta view

Linux too heavy for IoT?
Remains to be seen...

meta | Linux & IoT

maybe IoT hardware
becomes powerful enough
for Linux soon

meta | Linux & IoT

and maybe tinification
makes Linux more suitable

meta view

what else?

meta view

issues with
kernel development?

meta | development

developers getting old?

not enough fresh blood

not enough reviewers

LKML tone too unfriendly,

hostile and anti-woman?

meta changes | development

really complicated &
complex topics

meta changes | development

a lot is not as bad as the
journalists make it sound

meta changes | development

a lot of things are
way better then
5 or 10 years ago

meta | development

but yes, some things
could be better
like so often in life

that was part 5

%post

%post

more details?

%post | more details

use a search engine, Luke!

*slides and recording from talks
and articles available*

%post | more details

ask me if you can't find
anything

%post

three more things

%post

(1) help testing the kernel

%post | testing

especially if you own not
that widespread Hardware

%post | testing

or use your distro in a
unusual way

%post | testing

if you do not test it,
nobody might

%post | testing

bugs/regressions only get
noticed when they're old

%post | testing

making them

way harder

to track down and fix

%post | testing

which will annoy **you**
in the end

%post | testing

in your interest to test

%post

(2) tell me what you
think about this talk
or my articles

%post | feedback

there is always something
to improve

*I might do it again just like this if
you do not provide feedback*

%post

(3) want to be kept
up2date?

%post

heise online > Themen > Kernel-Log

Kernel-Log

Thema Kernel-Log



Alle Beiträge zu: Kernel-Log

Linux 4.8 bringt neue Treiber und verbessert die Sicherheit



Das am 3. Oktober erwartete Linux 4.8 bringt zahlreiche neue und überarbeitete Treiber, durch die der Kernel 500 weitere Geräte oder Geräteklassen unterstützt. Ferner gab es allerlei Verbesserungen, die das Ausnutzen von Sicherheitslücken erschweren.

19. September, 08:01 Uhr  59

Linux 4.8: Verbieten neuer XFS-Funktionen und Schnellstraße im Netzwerk-Stack



Das XFS-Dateisystem soll einige Features lernen, die Anwender an Btrfs und ZFS schätzen. Der Netzwerk-Code bekommt einen Express-Pfad, um bestimmte Aufgaben erheblich effizienter erledigen zu können.

Anzeige



it'sa 2016
Die IT-Security Messe und Kongress

Lösungen und Know-how
für mehr IT-Security
18. - 20. Oktober // Nürnberg

Get Your Upgrade, NOW!
Find new solutions

> mehr auf
heise online

Artikel zum Thema

Die Neuerungen von Linux 4.8

Das Anfang Oktober erwartete Linux 4.8 bringt zahlreiche neue und überarbeitete Treiber, durch die der Kernel fünfhundert weitere Geräte



www.heise.de/thema/Kernel_Log

%post



News from the source

Content

[Weekly Edition](#)

[Archives](#)

[Search](#)

[Kernel](#)

[Security](#)

[Distributions](#)

[Events calendar](#)

[Unread comments](#)

[LWN FAQ](#)

[Write for us](#)

[LWN.net Weekly Edition for September 29, 2016 \(One big page\)](#)

- [Front](#): GTK+ version numbering; Vulkan driver anatomy; OpenType 1.8 and style attributes.
- [Security](#): The trouble with new TLS version numbers; New vulnerabilities in bash, firefox, imagemagick, openssl, ...
- [Kernel](#): 4.8 development statistics; A low-level hibernation bug hunt.
- [Distributions](#): ARC++; RIP Kristoffer H. Rose, Ubuntu Yakkety Yak, Firefox OS, ...
- [Development](#): Systemd programming, 30 months later; MIT License, ...
- [Announcements](#): KDE Advisory Board, Lenovo laptops, ...

[LWN.net Weekly Edition for September 22, 2016 \(One big page\)](#)

- [Front](#): ATypI font special.
- [Security](#): On the way to safe containers; New vulnerabilities in chromium, graphicsmagick, kernel mozilla, ...
- [Kernel](#): Btrfs encryption; Stable kernel creation; BBR congestion control.
- [Distributions](#): The NTP pool system; Debian 8.6, ...
- [Development](#): Font build chains; Emacs 25.1; CouchDB 2.0; The Python packaging ecosystem; ...
- [Announcements](#): PGConf US cfp, ...

[LWN.net Weekly Edition for September 15, 2016 \(One big page\)](#)

- [Front](#): Automating hinting for every script; Backports and long-term stable kernels.
- [Security](#): Filesystem images & unprivileged containers; Minisail; New vulnerabilities in libarchive

lwn.net

%post

Google+ Search Sign in

The Linux Kernel Archives

Thorsten Leemhuis

The Linux kernel

#linux #kernel

2,357 followers · 169 posts · Public

FOLLOW

Thorsten Leemhuis 2d

The second pre-release of the #Linux #kernel 4.8 is out: <https://lkml.org/lkml/2016/8/14/874> Side note: The first regression report for Linux 4.8 is out as well: <https://lkml.org/lkml/2016/8/14/484> And I sent out another one for Linux 4.7, too: <https://lkml.org/lkml/2016/8/14/483>

LKML.ORG

Messages in this thread

• First message in thread

• Linux Torvalds

From: Linux Torvalds <>

Date: Sun, 14 Aug 2016 20:25:10 -0700

Subject: Linux 4.8-rc2

We've had a week since the merge window closed, and rc2 is out. Go test it.

The diffstat for rc2 looks unusual, in that only about 1/5th is drivers (normally drivers account for about half the bulk of the updates). Instead, it's dominated by arch updates, and fs/ and mfd. But that's most likely just because the fixes to the bigger driver areas haven't started trickling in yet - rc2 tends to be a quiet period after the craziness that is the merge window.

And the reason mfd stands out is mainly just because of a pull request from the merge window that I had delayed until after rc1 in order to look more at it.

So I expect that things will be back to normal this upcoming week.

Nothing really strange seems to be going on, so please just go out and test it and report any problems you encounter. It's obviously fairly early in the rc series, but I don't think there was anything particularly worrisome this merge window, so don't be shy.

Linux

Thorsten Leemhuis: +Marius Orcsik I haven't heard much about DAL recently, so I assumed things are still stucked, as

Thorsten Leemhuis 4d

Originally shared by Greg Kroah-Hartman - 27 comments

4.9 == next LTS kernel.

I need to get my blog back and running (stupid ruby dependencies), so until then, might as well say it here...

Thorsten Leemhuis 6d

Hellwig Announces He Will Appeal VMware Ruling After Evidentiary Set Back in Lower Court <http://sfconservancy.org/news/2016/aug/09/vmware-appeal/> [...] In a statement on his website, Christoph Hellwig announced today that he will appeal the ruling of the Hamburg District Court, which recently dismissed his case against VMware. As Christoph underscores in his statement, the ruling concerned German evidence law and the Court did not rule on the merits of the case. The ruling centered around German evidentiary rules related to documenting Christoph's contributi...

software freedom conservancy

plus.google.com/+ThorstenLeemhuis

%post

that's it – questions?

(TWIMC: this is slide #264)

Thorsten Leemhuis

mail: linux@leemhuis.info, thl@ct.de

GPG Key: 0x72B6E6EF4C583D2D

IRC@freenode.net: knurd

plus.google.com/ThorstenLeemhuis

#EOF