Man-pages: discovery, feedback loops, and the perfect kernel commit message

Michael Kerrisk
man7.org Training and Consulting
http://man7.org/training/

Outline

Outline

Who am I?

- Contributor to Linux *man-pages* project since 2000
- Maintainer since 2004
- Lots of testing, lots of bug reports
  - Much kernel reading; a very small number of kernel patches
- Author of a book on the Linux programming interface
- IOW: looking at Linux APIs a lot and for a long time
  - I.e., kernel-user-space APIs and libc APIs
- *http://man7.org/*

man7.org

The Linux *man-pages* project

- Documents kernel-user-space and C library APIs
  - Mostly pages in Sections 2 (syscalls) and 3 (library functions)
  - Some pages in Sections 4 (devices) and 5 (file formats)
  - Also: many overview pages in Section 7
  - *https://www.kernel.org/doc/man-pages/*
- Passed 1000-page mark in July 2016
  - ≈2200 interfaces documented
  - ≈146k lines (≈2500 pages) of rendered text

man7.org

Outline

Two interlinked topics:

- *man-pages* project
  - History, current state, challenges
- How can we get API design right (or at least better)?
  - Why API design is challenging
  - Mitigations
  - The problem of discovery
  - The feedback loop
  - The perfect kernel commit message

man7.org

Outline

Brief history of *man-pages*

- Founded in 1993
- Release 1.0: 305 pages
  - Pages put together mostly by $\approx$6 authors
  - Often rather short pages (average rendered length: 50 lines)
- Initial maintainer: Rik Faith
  - 1.0 to 1.5 (1993 - Feb 1995)
- Subsequently: Andries Brouwer
  - 1.6 to 1.70 (1995 - Oct 2004)
- Since Nov 2004: Michael Kerrisk
  - 2.00 onward
  - As at 4.07 (July 2016): 1002 pages
- (Two lengthy spells of maintainership $\Rightarrow$ good continuity!)

man7.org

Some statistics pre/post 2004

Attribute	Pre 2.00	2.00 and later
Timespan	1993-2004 (11 yrs)	2004–2016 (12 yrs)
# of releases	71	174 [4, 5] *
Avg diff/yr [1]	24k [2]	75k [3] *
Pages added	765	≈262 [4]
Pages removed	-	≈25 [4]
Avg rendered page length	95 lines (1.70)	145 lines [4] *

* I like to believe that I've improved the state of the project

- Much higher level of activity

- More, longer, better pages

[1] Diff stats exclude POSIX man pages and COLOPHON sections

[2] Includes initial release (1.0)

[3] Especially in man-pages-2.xx: vast numbers of typo, spelling (US), notational, and consistency fixes

[4] As at man-pages-4.07, July 2016

[5] ≈16k commits

man7.org

Why things are better

- I've put a lot of energy into the project
- Some of that was to turn *man-pages* into a **visible** project
- Before 2004, *man-pages* was nearly invisible:
    - No regular release announcements to any mailing list
    - No version control(!) or change logs (no history :-( )
    - No public infrastructure
    - No in-page info on how to report bugs
- Fixes
    - Regular release notes on LKML since start of 2006
    - Nov 2004: private SVN; from 2008: public Git
    - Late 2007: added project mailing list, website, bug tracker, blog, online rendered pages
    - Dec 2007: ⇒ `COLOPHON` on each page describes how to report bugs (a feedback loop!)

Outline

Challenges: participation (or: the lies I tell)

- man-pages-4.05
  - 467 commits, 440 pages changed
  - 74 "contributors" (a record)
  - The hidden truth:
    - Most contributions are comments or emailed bug reports
    - Few actual patches or reviews of patches
    - From 467 commits: I was author of 401 ($\approx$70%)
  - But, outside contribution is still **much** better than in 2004
    - A "good" release in 2005 might have seen input from 10 people
- Since 2004: $\approx$262 new man pages added
  - The hidden truth: I wrote 164 of those ($>$ 60%)
    - And cowrote many of the others
  - But, culture has slowly improved...
    - E.g., for all 4 syscalls added in Linux 3.17, devs drafted a man page

man7.org

Challenges: do I trust a patch?

- Many corners of interface where I'm not deeply knowledgeable...

- To detect bogus patches and bug reports in those corners, I need one of:
  - **Confidence in submitter**/reporter (usually based on past work; uncommon)
  - A **competent reviewer** (often difficult to find)
  - To **improve my own knowledge** sufficiently so that I can review (can be *very* time-consuming)

- Lacking any of above, reports+patches languish/get lost :-(
  - Sometimes revisit *much* later, and find I do now have requisite knowledge
    - Occasionally, reports get dealt with 5+ years later :-(

man7.org

Challenges: lack of contributors

- Lack of patches from others
  - Patches from kernel / libc devs are still the exception
  - 2009-present: I am author of $> 75\%$ of all patches
    - Yes, perhaps half of my patches are typo/wording/minor fixes, but still...
- Lack of reviewers ($\approx$100 `Reviewed-by:` tags in git log)
  - I am reviewer of last resort for vast majority of patches

man7.org

Challenges: me as the bottleneck

- Most of my work on *man-pages* has been voluntary
  - Except ≈8 months in 2008 in paid LF fellowship
- In addition to being maintainer, I am majority contributor
- Pace of project depends strongly on my energy/availability
  - Pace has varied wildly; for example (commits/year):
    - 2007: 1712
    - 2011: 296 (pretty burned out; nearly stepped away)
    - 2015: 3076
    - 2016: ≈2000 (expected)

man7.org

Ways to help: contribution

- Whenever you see someone changing the user-space API:
  - Remind them to CC *linux-api@vger.kernel.org*
    - *https://www.kernel.org/doc/man-pages/linux-api-ml.html*
  - Ask them to (in decreasing order of preference):
    - Write a patch for the man page
    - Send in some plain text describing API change
    - CC me + *linux-man@vger.kernel.org* on mail thread containing source code patch
      ⚠ But this is not a scalable solution...
- Write patches for man pages
  - *http://www.kernel.org/doc/man-pages/contributing.html*
- Review patches on *linux-man@vger.kernel.org*

man7.org

Ways to help: funding/finding a maintainer

- The situation where there is no paid maintainer for core documentation is ridiculous, right?
- I like to believe that current *man-pages* is a lot better than what I inherited
- But it could be so much better...
  - E.g., 250+ commits in *man-pages-4.04* to expand feeble *futex(2)* page from 169 to 1001 lines
    - But that work was > 5 years overdue
  - Long backlog of work:
    - ≈200 FIXMEs in man pages source files
    - *www.kernel.org/doc/man-pages/missing_pages.html* (a long list)

man7.org

Ways to help: funding/finding a maintainer

- There is (easily) enough work for a full-time maintainer
  - And I'm not necessarily saying it should/I want it to be me
- But, failing that, point people at *http://man7.org/training*
  - Help keep the current *man-pages* maintainer and family fed...

man7.org

Outline

Many kinds of APIs

- Pseudo-filesystems (`/proc`, `/sys`, `/dev/mqueue`, debugfs, configfs, etc.)
- Netlink
- Auxiliary vector
- Virtual devices
- Signals
- System calls ⇐ focus, for purposes of example
- Multiplexor syscalls (*ioctl()*, *prctl()*, *fcntl()*, *bpf()*, *perf_event_open()*, ...)

man7.org

Design goals for APIs

Properly designed and implemented API should:

- be bug free!

- be as simple as possible (but no simpler)

- be easy to use / difficult to misuse

- be consistent with related/similar APIs

- avoid need for compat layer, or gratuitous arch. differences

- integrate well with existing APIs
  - e.g., interactions with *fork()*, *exec()*, threads, signals, FDs

- be as general as possible

- allow for future extension

- adhere to relevant standards, where possible (e.g., POSIX)

- be at least as good as earlier APIs with similar functionality

- be maintainable over time (a multilayered question)

man7.org

We've failed repeatedly
on every one of those points

A few personal/recent favorites follow; for much more, see:

http://man7.org/conf/fosdem2016/designing_linux_kernel_APIs-fosdem-2016-Kerrisk.pdf
http://man7.org/conf/lca2013/Why_kernel_space_sucks-2013-02-01-printable.pdf

man7.org

Bugs

- Won't go into numerous examples...
- Suffice to say, kernel (and libc) APIs have repeatedly been released with bugs
    - "Show me a new Linux API, and I'll show you a bug"
    - (More recently, fuzzers such as *trinity* have helped get rid of many of more egregious cases)
- Frequently: insufficient prerelease testing
- Painful for userspace
    - User-space code may need to special case for kernel version

Design inconsistencies

From `arch/Kconfig`

```
#
# ABI hall of shame
#
config CLONE_BACKWARDS
        bool
        help
          Architecture has tls passed as the 4th argument of clone(2),
          not the 5th one.

config CLONE_BACKWARDS2
        bool
        help
          Architecture has the first two arguments of clone(2) swapped.

config CLONE_BACKWARDS3
        bool
        help
          Architecture has tls passed as the 3rd argument of clone(2),
          not the 5th one.
...
```

- And still more variations on ia64, SPARC, blackfin, m68k
- At least a half dozen *clone()* APIs...

Behavioral inconsistencies

- *mlock(start, length)*
    - Round *start* down to page boundary
    - Round *length* up to next page boundary
    - *mlock(4000, 6000)* affects bytes 0..12287
        - (Assuming page size of 4096B)
- *remap_file_pages(start, length, ...)*
    - Round *start* down to page boundary
    - Round *length* **down** to next page boundary
    - *remap_file_pages(4000, 6000)* affects?
        - Bytes **0 to 4096**
- Users expect similar looking APIs to behave similarly
    - Violate that assumption, and users write buggy code

Behavioral inconsistencies

- Various system calls allow one process to change attributes of another process
    - e.g., *setpriority()*, *ioprio_set()*, *migrate_pages()*, *prlimit()*
- Calls from unprivileged process require UID/GID match between caller and target
    - I.e., some combination UIDs or GIDs must match between caller and target ("*t-*")
- Let's make life interesting for user space:
    - *setpriority()*: *euid == t-ruid || euid == t-euid*
    - *ioprio_set()*: *ruid == t-ruid || euid == t-ruid*
    - *migrate_pages()*: *uid == t-ruid || uid == t-suid ||
      euid == t-ruid || euid == t-suid*
    - *prlimit()*:
      *(ruid == t-ruid && ruid == t-euid && ruid == t-suid) &&
      (rgid == t-rgid && rgid == t-guid && rgid == t-sgid)*

man7.org

Maintainability: a many faceted problem

- API maintainability has many aspects...

Maintainability: extensible APIs

- Many historical Linux APIs lacked a *flags* argument or other mechanism to allow extension of an API
  - Thus: *umount()* ⇒ *umount2()*; *preadv()* ⇒ *preadv2()*; *epoll_create()* ⇒ *epoll_create1()*; renameat() ⇒ *renameat2()*; and so on
  - *https://lwn.net/Articles/585415/*
- And many historical APIs that had *flags* argument failed to check for invalid flag bits
  - *sigaction(sa.sa_flags)*, *recv()*, *clock_nanosleep()*, *msgrcv()*, *semget()*, *semop(sops.sem_flg)*, *open()*, and many others
  - Problem 1: assigning meaning to previously unused bit may break user-space code that carelessly passed that bit
  - Problem 2: user-space has no way to check kernel support for a flag
  - *https://lwn.net/Articles/588444/*

man7.org

Maintainability: we don't do decentralized design well

- Decentralized development can fail badly when it comes to (coherent) design

Maintainability: we don't do decentralized design well

- Linux capabilities: divide power of root into **small** pieces
  - A compromised program that has capabilities is harder to exploit than a compromised set-UID program
  - Linux 4.8: 38 capabilities
- Kernel developer's dilemma for new "dangerous" feature:
  - Add a new capability? (But: avoid explosion of capabilities)
  - Or assign feature to existing capability silo?
- Adding to an existing silo is preferable...
  - "But which one?"
  - (Looks at *capabilities(7)*) "Hey! Sysadmins will do this!"
  - Welcome `CAP_SYS_ADMIN`, the new root
  - ≈40% of all capability checks in kernel (game over...)
    - *https://lwn.net/Articles/486306/*

man7.org

Maintainability: we don't do decentralized design well

- Cgroups v1...
- A mess of inconsistent interfaces, interpretation of "hierarchy", and more

man7.org

We're just traditionalists

- It's not just us...
- A long history of getting things wrong in UNIX APIs
  - Using syscall function result to both return info and indicate success/failure is a fundamental design error
    - Purposes can conflict: *getpriority()*, *fcntl(F_GETOWN)*
  - Design of System V IPC truly was awful
  - Semantics of POSIX record locks are broken by design
    - Linux now has a better replacement!
  - *select()* modifies FD sets in place, forcing reinitialization inside loops
  - UNIX domain socket *sun_path* null termination
    - Present since 1984
    - *http://man7.org/conf/fosdem2016/puzzle-slides--UNIX-domain-sockets-API-bug.pdf*

man7.org

API design is hard

And when we fail...

- (Usually) can't fix a broken API
  - Fix == ABI change
  - User-space will break
  - (By contrast, fixing non-user-facing bugs and performance issues is often much easier)
- *Thousands* of user-space programmers will live with a (bad) design for *decades*
- ⇒ We need to get API design right first time

man7.org

Outline

What can we do to ensure API design is better first time round?

Goals

- Make sure API is well designed, fit for purpose, and extensible
- Prevent ABI regressions
- Minimize bugs

Resources

- Review
- Testing
  - Mechanical testing has limited application
- Need to involve humans...
- As early as possible
  - (Usually can't fix an API after release)

Outline

Unit tests

- To state the obvious, unit tests:
  - **Prevent behavior regressions** in face of future refactoring of implementation
  - Provide **checks that API works as expected**/advertised
    - I.e., does it do what it says on the tin?
- Failures on both points have been surprisingly frequent
  - See my previous presentations

man7.org

Example (does it do what it says on the tin?)

- *recvmmsg()* system call (linux 2.6.33)
  - Performance: receive multiple datagrams via single syscall
  - *timeout* argument added late in implementation, after reviewer suggestion
- Intention versus implementation:
  - Apparent concept: place timeout on receipt of complete set of datagrams
  - Actual implementation: timeout *tested only after receipt of each datagram*
    - Renders timeout useless...
- Clearly, no serious testing of implementation

Where to put your tests?

- Historically, only real home was LTP (Linux Test Project), but:
    - Tests were out of kernel tree
    - Often added only after APIs were released
    - Coverage was only partial
    - *https://linux-test-project.github.io/*
- *kselftest* project (started in 2014) was created to improve matters:
    - Tests reside in kernel source tree
        - `make kselftest`
    - Paid maintainer: Shuah Khan
    - Wiki: *https://kselftest.wiki.kernel.org/*
    - Mailing list: *linux-kselftest@vger.kernel.org*

man7.org

But, how do you know what to test if there is no specification?

Outline

"Programming is not just an act of telling a computer what to do: it is also an act of telling other programmers what you wished the computer to do. Both are important, and the latter deserves care."

Andrew Morton, March 2012

man7.org

Fundamental problem behind
(e.g.) *recvmmsg() timeout* bugs:

no one wrote a specification
during development or review

A test needs a specification

recvmmsg() timeout argument needed a specification; something like:

- The *timeout* argument implements three cases:
    1. *timeout* is NULL: the call blocks until *vlen* datagrams are received.
    2. *timeout* points to *{0, 0}*: the call (immediately) returns up to *vlen* datagrams if they are available. If no datagrams are available, the call returns immediately, with the error EAGAIN.
    3. *timeout* points to a structure in which at least one of the fields is nonzero.
       The call blocks until either:
        - (a) the specified timeout expires
        - (b) *vlen* messages are received
       In case (a), if one or more messages has been received, the call returns the number of messages received; otherwise, if no messages were received, the call fails with the error EAGAIN.

- If, while blocking, the call is interrupted by a signal handler, then:
    - if 1 or more datagrams have been received, then those datagrams are returned (and interruption by a signal handler is not (directly) reported by this or any subsequent call to *recvmmsg()*.
    - if no datagrams have so far been received, then the call fails with the error EINTR.

man7.org

Specifications help

Specifications have numerous benefits:

- Provides target for implementer
- Without specification, how can we differentiate implementer's *intention* from actual *implementation*
    - IOW: **how do we know what is a bug?**
- Allow us to write unit tests
- Allow reviewers to more easily understand and critique API
    - $\Rightarrow$ will likely increase number of reviewers

man7.org

Where to put your specification?

- At a minimum: in the commit message
- To gain good karma: a *man-pages* patch
  - *https://www.kernel.org/doc/man-pages/patches.html*

man7.org

Outline

Example: inotify

- Filesystem event notification API
  - Detect file opens, closes, writes, renames, deletions, etc.
- *A Good Thing*™...
  - Improves on predecessor (*dnotify*)
  - Better than polling filesystems using *readdir()* and *stat()*
- But it should have been *A Better Thing*™

man7.org

Writing a "real" inotify application

- Back story: I thought I understood inotify
- Then I tried to write a "real" application...
  - **Mirror state of a directory tree** in application data structure
  - 1500 lines of C with (lots of) comments
    - *http://man7.org/tlpi/code/online/dist/inotify/inotify _dtree.c.html*
  - Written up on LWN (*https://lwn.net/Articles/605128/*)
- And understood all the work that inotify still leaves you to do
- **And what inotify could perhaps have done better**

man7.org

The limitations of inotify

A few among several tricky problems when using inotify:

- Event notifications don't include PID or UID
  - Can't determine who/what triggered event
  - It might even be you
  - *Why not supply PID / UID, at least to privileged programs?*
- Monitoring of directories is not recursive
  - Must add new watches for each subdirectory
    - (**Probably unavoidable** limitation of API)
  - Can be **expensive** for large directory tree ⇒ **see next point**

man7.org

The limitations of inotify

File renames generate MOVED_FROM+MOVED_TO event pair

- Useful: provides old and new name of file
- But two details combine to create a problem:
    - MOVED_FROM+MOVED_TO **not** guaranteed to be consecutive
    - No MOVED_TO if target directory is not monitored
        - Can't be sure if MOVED_FROM will be followed by MOVED_TO
- ⇒ matching MOVED_FROM+MOVED_TO must be done heuristically
    - Unavoidably racey, leading to possible matching failures
- Matching failures ⇒ treated as tree delete + tree re-create (expensive!)
- *User-space handling would have been much simpler, and deterministic, if MOVED_FROM+MOVED_TO had been guaranteed consecutive by kernel*

Only way to discover design problems in a new nontrivial API is by writing complete, real-world application(s)

(preferably more than one...)
(before the API is released in mainline kernel...)

API limitations should be rectified, or at least clearly documented, before API release...

Outline

Documentation is good for the health of APIs

- Inevitably, the process of **writing documentation makes you reflect about your design** more deeply
- Documentation:
  - **Makes it easier for others to understand** your API, think about it, and critique it
    - Lowers hurdle for involvement
  - **Broadens the audience that will understand** and critique your API
    - Do it well enough, and you might even get user-space programmers involved
- A well written man page is a pretty good vehicle, I'd say

man7.org

Man pages as a test specification

A well written man page often suffices as a test specification for finding real bugs:

- *utimensat()*: *http://linux-man-pages.blogspot.com/2008/06/whats-wrong-with-kernel-userland_30.html*
- *timerfd*: *http://thread.gmane.org/gmane.linux.kernel/613442*
    - (Gmane come back soon, we miss you)

man7.org

Outline

How do we discover when an API change has occurred?

- How do we discover when a kernel-user-space API change has occurred?
- No simple way...
- Personally (for *man-pages*):
    - I mostly don't have time to track LKML
    - Watching *linux-api@vger.kernel.org*
    - Scripting to find candidate API differences between successive kernel versions trees
        - Very imperfect...
    - LWN, KernelNewbies LinuxChanges
    - Sheer luck
        - Randomly notice something from reading kernel source, an online article/mail thread, f2f conversation, etc.
        - Occasionally, a *man-pages* patch out of the blue

man7.org

How do we discover when an API change has occurred?

- **Many** people are interested in this question, including:
  - User-space programmers
  - C library developers
  - *man-pages* project
  - *strace* project
  - Testing projects (LTP, trinity, ...)
  - LSB, KernelNewbies LinuxChanges, ...
- **Please** CC *linux-api@vger.kernel.org* on API/ABI changes...
- Discovery occurs at different times/rates for different groups
  - User-space programmers, as a group, are most affected
    - And often the last to know!

man7.org

Discoverability is even a problem for kernel developers

"Quite frankly, our most common ABI change is that
we don't even realize that something changed.
And then people may or may not notice it."
–Linus Torvalds, LKML, March 2012

- I.e., kernel developers are sometimes not even aware they are changing kernel-user-space API

Silent API changes

- So we get silent API changes
- Two (from many) examples:
  - Adjustments of POSIX MQ implementation in Linux 3.5 caused two user-space breakages
    - *mq_overview(7)*
  - Linux 2.6.12 silently changed semantics of *fcntl(F_SETOWN)* for MT programs
    - But only worked this out a few years later...
    - Too late to revert (maybe some apps depend on new behavior!)
    - Linux 2.6.32 added `F_SETOWN_EX` to provide old behavior
  - (Unit tests, anyone?)

man7.org

Outline

The problem

- Probably 6+ months before your API appears in distributions and starts getting used in real world
- Worst case: only then will bugs be reported and design faults become clear
  - As user-space programmers start to employ APIs in real-world applications
- But that's too late...
  - (Probably can't change ABI...)
- Need as much feedback as possible **before** API is released

Strive to shorten worst-case
feedback loop
$\Rightarrow$
Publicize API design
as widely + early as possible

Shortening the feedback loop

Ideally, do all of the following before API release (1/2):

- Write a **detailed specification**
    - Elaborate full range of inputs for all arguments
    - Elaborate consequent behavior and resulting output
    - Consider interactions with signals, threads, *fork()*, *execve()*
- Write **example programs** that fully demonstrate API
- Email relevant mailing lists and, especially, relevant people
- CC *linux-api@vger.kernel.org*
    - As per `Documentation/SubmitChecklist`...
    - Alerts interested parties of API changes:
        - C library projects, *man-pages*, LTP, trinity, kselftest, LSB, tracing projects, and user-space programmers
        - *https://www.kernel.org/doc/man-pages/linux-api-ml.html*

man7.org

Shortening the feedback loop

Ideally, do all of the following before API release (2/2):

- For good karma + more publicity: write an LWN.net article
  - Good way of **reaching end users** of your API
    - Ask readers for feedback
  - *http://lwn.net/op/AuthorGuide.lwn*

man7.org

Outline

The perfect kernel commit message?

Okay; perfection is in the eye of the beholder

Perfection = better documentation and better user-space APIs

Three iterations
of a patch series that I happened
to get interested in recently

Version 1

```
Subject: Show virtualized dentry root in mountinfo for cgroupfs
Date: Sun, 17 Apr 2016 15:04:30 -0500
From: Serge Hallyn

With the current cgroup namespace patches, the root dentry path of a
mount as shown in /proc/self/mountinfo is the full global cgroup
path.   It is common for userspace to use /proc/self/mountinfo to
search for cgroup mountpoints, and expect the root dentry path to
relate to the cgroup paths in /proc/self/cgroup.   Patch 2 in this
set therefore virtualizes the root dentry path relative to the
reader's cgroup namespace root.
```

- For a people in the know (perhaps a few in CC), the above might be clear

- For ~~idiots~~ me, it's far from clear what this is about

- There's value in assuming there are lots of ~~idiots~~ people short on time out there

  - Some of them might be able to help you

man7.org

Version 2

After some offlist conversations with Serge

```
Subject: [PATCH] mountinfo: implement show_path for kernfs and cgroup
Date: Thu, 5 May 2016 10:20:58 -0500
From: Serge Hallyn

Short explanation:

When showing a cgroupfs entry in mountinfo, show the path of the mount
root dentry relative to the reader's cgroup namespace root.

Long version:

When a uid 0 task which is in freezer cgroup /a/b, unshares a new cgroup
namespace, and then mounts a new instance of the freezer cgroup, the new
mount will be rooted at /a/b.  The root dentry field of the mountinfo
entry will show '/a/b'.
 [38 more lines omitted]
```

- Better, but...
- Short version doesn't really explain user-space problem that is being solved
- Long version could still break things down rather more clearly

man7.org

Version 3

After more conversation with Serge

```
Subject: [PATCH] mountinfo: implement show_path for kernfs and cgroup
Date: Mon, 9 May 2016 09:59:55 -0500
From: Serge Hallyn

Patch summary:

When showing a cgroupfs entry in mountinfo, show the path of the mount
root dentry relative to the reader's cgroup namespace root.

Short explanation (courtesy of mkerrisk):

If we create a new cgroup namespace, then we want both
/proc/self/cgroup and /proc/self/mountinfo to show cgroup paths that
are correctly virtualized with respect to the cgroup mount point.
Previous to this patch, /proc/self/cgroup shows the right info,
but /proc/self/mountinfo does not.

["Long version" As before]
```

- I.e., include a short summary of the user-space problem
  - Best tailored to an audience that is naïve about the domain
    - Short explanation here might even be enough to **give a random user-space programmer a clue** what this is about

Version 3

- But there's more

```
Subject: [PATCH] mountinfo: implement show_path for kernfs and cgroup
Date: Mon, 9 May 2016 09:59:55 -0500
[...]

Example (by mkerrisk):

[94 lines of shell sessions plus explanations]
```

A detailed example:

- Complete walk through starting from scratch: shell commands + explanations

- Demonstration of the problem as it exists without the patch

- Demonstration of the same command sequence on a patched kernel, showing how it fixes problem

- Did this to make sure **I** understand, but it's exactly the info many others need for understanding

man7.org

Overkill?

- You might argue that this is overkill
- I'd argue that it makes a whole lot of people's lives easier
  - Including mine
- And you (the kernel developer) probably made your own life easier too
  - More reviewers, more feedback, better/faster feedback
  - And when you come back to this later, **you** will be able to understand what you did and why

man7.org

Who should do this for each patch?

- You know the answer
  - It doesn't scale for me to do this
- One person has all the requisite knowledge: you, the original developer
  - You will have done all the thinking, and (hopefully) testing
  - Just need to elaborate that in writing
  - And the less knowledge you assume in your audience, the wider that audience can be

Summary: why you should be doing this

This is about:

- Making you think harder about the API
- Making you do careful walk-through testing
- Showing others what you mean in detail
- Lowering the bar to understanding
- Letting discovery happen earlier and more easily
- Broadening your reviewer base

Summary: when you should be doing this

- Feedback about API bugs that arrives after mainline release is usually too late...

- Many (most?) API changes that are interesting have a long gestation

  - I.e., many patch iterations
  - E.g., memory protection keys:
    - First patch submission in May 2015
    - Merged in Linux 4.9-rc1
    - Mainline release in December 2016

- The long development window that precedes release is an opportunity...

- Don't leave it to late patch iterations to make your commit message "rich"

  - **Lengthen the feedback window**: do it from the beginning

man7.org

Outline

Doing it right

Jeff Layton, OFD locks, Linux 3.15 (commit 5d50ffd7c31):

- "Open file description locks"
- Fix serious design problems with POSIX record locks
  - (POSIX record locks are essentially unreliable in the presence of any library that works with files)
- Did everything nearly perfectly, in terms of developing feature

man7.org

Doing it right

Jeff Layton, OFD locks, Linux 3.15 (commit 5d50ffd7c31):

- Clearly explained **rationale** and changes in commit message
- Provided example programs
- Publicized the API
  - Mailing lists
  - LWN.net article (*http://lwn.net/Articles/586904/*)
- Wrote a man pages patch
  - (Feedback led to renaming of constants and feature)
- Engaged with glibc developers (patches for glibc headers + manual)
  - Refined patches in face of review
  - Maintainers were unresponsive $\Rightarrow$ resubmitted *many* times
- Triggered work to get API into next POSIX standard
- Made it all look simple

Thanks!

mtk@man7.org
Slides at http://man7.org/conf/

Linux/UNIX system programming training (and more)
http://man7.org/training/

The Linux Programming Interface, http://man7.org/tlpi/

Outline

Version 2 (complete mail)

```
Subject: [PATCH] mountinfo: implement show_path for kernfs and cgroup
Date: Thu, 5 May 2016 10:20:58 -0500
From: Serge Hallyn

Short explanation:

When showing a cgroupfs entry in mountinfo, show the path of the mount
root dentry relative to the reader's cgroup namespace root.

Long version:

When a uid 0 task which is in freezer cgroup /a/b, unshares a new cgroup
namespace, and then mounts a new instance of the freezer cgroup, the new
mount will be rooted at /a/b. The root dentry field of the mountinfo
entry will show '/a/b'.

 cat > /tmp/do1 << EOF
 mount -t cgroup -o freezer freezer /mnt
 grep freezer /proc/self/mountinfo
 EOF

 unshare -Gm  bash /tmp/do1
 > 330 160 0:34 / /sys/fs/cgroup/freezer rw,nosuid,nodev,noexec,relatime
 > 355 133 0:34 /a/b /mnt rw,relatime - cgroup freezer rw,freezer

The task's freezer cgroup entry in /proc/self/cgroup will simply show
'/':
```

man7.org

Version 2 (complete mail)

```
 grep freezer /proc/self/cgroup
 9:freezer:/

If instead the same task simply bind mounts the /a/b cgroup directory,
the resulting mountinfo entry will again show /a/b for the dentry root.
However in this case the task will find its own cgroup at /mnt/a/b,
not at /mnt:

 mount --bind /sys/fs/cgroup/freezer/a/b /mnt
 130 25 0:34 /a/b /mnt rw,nosuid,nodev,noexec,relatime shared:21 - cgroup c

In other words, there is no way for the task to know, based on what is
in mountinfo, which cgroup directory is its own.

With this patch, the dentry root field in mountinfo is shown relative
to the reader's cgroup namespace. I.e.:

 unshare -Gm  bash /tmp/do1
 > 330 160 0:34 / /sys/fs/cgroup/freezer rw,nosuid,nodev,noexec,relatime -
 > 355 133 0:34 / /mnt rw,relatime - cgroup freezer rw,freezer

This way the task can correlate the paths in /proc/pid/cgroup to
/proc/self/mountinfo, and determine which cgroup directory (in any
mount which the reader created) corresponds to the task.
```

man7.org

Version 3 (complete mail)

```
Subject: [PATCH] mountinfo: implement show_path for kernfs and cgroup
Date: Mon, 9 May 2016 09:59:55 -0500
From: Serge Hallyn

Patch summary:

When showing a cgroupfs entry in mountinfo, show the path of the mount
root dentry relative to the reader's cgroup namespace root.

Short explanation (courtesy of mkerrisk):

If we create a new cgroup namespace, then we want both /proc/self/cgroup
and /proc/self/mountinfo to show cgroup paths that are correctly
virtualized with respect to the cgroup mount point.  Previous to this
patch, /proc/self/cgroup shows the right info, but /proc/self/mountinfo
does not.

Long version:

When a uid 0 task which is in freezer cgroup /a/b, unshares a new cgroup
namespace, and then mounts a new instance of the freezer cgroup, the new
mount will be rooted at /a/b.  The root dentry field of the mountinfo
entry will show '/a/b'.

 cat > /tmp/do1 << EOF
 mount -t cgroup -o freezer freezer /mnt
 grep freezer /proc/self/mountinfo
```

Version 3 (complete mail)

```
 EOF

 unshare -Gm  bash /tmp/do1
 > 330 160 0:34 / /sys/fs/cgroup/freezer rw,nosuid,nodev,noexec,relatime - c
 > 355 133 0:34 /a/b /mnt rw,relatime - cgroup freezer rw,freezer

The task's freezer cgroup entry in /proc/self/cgroup will simply show
'/':

 grep freezer /proc/self/cgroup
 9:freezer:/

If instead the same task simply bind mounts the /a/b cgroup directory,
the resulting mountinfo entry will again show /a/b for the dentry root.
However in this case the task will find its own cgroup at /mnt/a/b,
not at /mnt:

 mount --bind /sys/fs/cgroup/freezer/a/b /mnt
 130 25 0:34 /a/b /mnt rw,nosuid,nodev,noexec,relatime shared:21 - cgroup cp

In other words, there is no way for the task to know, based on what is
in mountinfo, which cgroup directory is its own.

Example (by mkerrisk):

First, a little script to save some typing and verbiage:

# cat cgroup_info.sh
```

Version 3 (complete mail)

```
#!/bin/sh
echo -e "\t/proc/self/cgroup:\t$(cat /proc/self/cgroup | grep freezer)"
cat /proc/self/mountinfo | grep freezer |
        awk '{print "\tmountinfo:\t\t" $4 "\t" $5}'
#

Create cgroup, place this shell into the cgroup, and look at the state
of the /proc files:

# mkdir -p /sys/fs/cgroup/freezer/a/b
# echo $$ > /sys/fs/cgroup/freezer/a/b/cgroup.procs
# echo $$
2653
# cat /sys/fs/cgroup/freezer/a/b/cgroup.procs
2653                              # Our shell
14254                            # cat(1)
# ./cgroup_info.sh
        /proc/self/cgroup:      10:freezer:/a/b
        mountinfo:              /          /sys/fs/cgroup/freezer

Create a shell in new cgroup and mount namespaces. The act of creating
a new cgroup namespace causes the process's current cgroups directories
to become its cgroup root directories. (Here, I'm using my own version
of the "unshare" utility, which takes the same options as the util-linux
version):

# ~mtk/tlpi/code/ns/unshare -Cm bash
```

man7.org

Version 3 (complete mail)

```
Look at the state of the /proc files:

# ./cgroup_info.sh
        /proc/self/cgroup:      10:freezer:/
        mountinfo:              /       /sys/fs/cgroup/freezer

The third entry in /proc/self/cgroup (the pathname of the cgroup inside
the hierarchy) is correctly virtualized w.r.t. the cgroup namespace,
which is rooted at /a/b in the outer namespace.

However, the info in /proc/self/mountinfo is not for this cgroup
namespace, since we are seeing a duplicate of the mount from the
old mount namespace, and the info there does not correspond to the
new cgroup namespace. However, trying to create a new mount still
doesn't show us the right information in mountinfo:

# mount --make-rslave /               # Prevent our mount operations
                                       # propagating to other mountns
# mkdir -p /mnt/freezer               # Create a new mount point
# umount /sys/fs/cgroup/freezer       # Discard old mount
# mount -t cgroup -o freezer freezer /mnt/freezer/
# ./cgroup_info.sh
        /proc/self/cgroup:      7:freezer:/
        mountinfo:              /a/b    /mnt/freezer

The act of creating a new cgroup namespace caused the process's
current freezer directory, "/a/b", to become its cgroup freezer root
directory. In other words, the pathname directory of the directory
```

man7.org

Version 3 (complete mail)

```
   within the newly mounted cgroup filesystem should be "/",
   but mountinfo wrongly shows us "/a/b". The consequence of this is
   that the process in the cgroup namespace cannot correctly construct
   the pathname of its cgroup root directory from the information in
   /proc/PID/mountinfo.

   With this patch, the dentry root field in mountinfo is shown relative
   to the reader's cgroup namespace.  So the same steps as above:

   # mkdir -p /sys/fs/cgroup/freezer/a/b
   # echo $$ > /sys/fs/cgroup/freezer/a/b/cgroup.procs
   # ./cgroup_info.sh
           /proc/self/cgroup:      10:freezer:/a/b
           mountinfo:              /           /sys/fs/cgroup/freezer
   # ~mtk/tlpi/code/ns/unshare -Cm bash
   # ./cgroup_info.sh
           /proc/self/cgroup:      10:freezer:/
           mountinfo:              /../..  /sys/fs/cgroup/freezer
   # mount --make-rslave /
   # mkdir -p /mnt/freezer
   # umount /sys/fs/cgroup/freezer
   # mount -t cgroup -o freezer freezer /mnt/freezer/
   # ./cgroup_info.sh
           /proc/self/cgroup:      10:freezer:/
           mountinfo:              /           /mnt/freezer

   # ls /mnt/freezer/
   cgroup.clone_children   freezer.parent_freezing   freezer.state      tasks
```

Version 3 (complete mail)

```
cgroup.procs           freezer.self_freezing    notify_on_release
# echo $$
3164
# cat /mnt/freezer/cgroup.procs
2653                   # First shell that placed in this cgroup
3164                   # Shell started by 'unshare'
14197                  # cat(1)
```

man7.org