



How to Participate in the Linux Community

A Guide To The Kernel Development Process

by Jonathan Corbet, corbet@lwn.net

The purpose of this document is to help developers (and their managers) work with the development community with a minimum of frustration. It is an attempt to document how this community works in a way which is accessible to those who are not intimately familiar with Linux kernel development (or, indeed, free software development in general). While there is some technical material here, this is very much a process-oriented discussion which does not require a deep knowledge of kernel programming to understand.



August 2008

1.1: Executive Summary

The rest of this section covers the scope of the kernel development process and the kinds of frustrations that developers and their employers can encounter there. There are a great many reasons why kernel code should be merged into the official (“mainline”) kernel, including automatic availability to users, community support in many forms, and the ability to influence the direction of kernel development. Code contributed to the Linux kernel must be made available under a GPLcompatible license.

Section 2 introduces the development process, the kernel release cycle, and the mechanics of the merge window. The various phases in the patch development, review, and merging cycle are covered. There is some discussion of tools and mailing lists. Developers wanting to get started with kernel development are encouraged to track down and fix bugs as an initial exercise.

Section 3 covers early-stage project planning, with an emphasis on involving the development community as soon as possible.

Section 4 is about the coding process; several pitfalls which have been encountered by other developers are discussed. Some requirements for patches are covered, and there is an introduction to some of the tools which can help to ensure that kernel patches are correct.

Section 5 talks about the process of posting patches for review. To be taken seriously by the development community, patches must be properly formatted and described, and they must be sent to the right place. Following the advice in this section should help to ensure the best possible reception for your work.

Section 6 covers what happens after posting patches; the job is far from done at that point. Working with reviewers is a crucial part of the development process; this section offers a number of tips on how to avoid problems at this important stage. Developers are cautioned against assuming that the job is done when a patch is merged into the mainline.

Section 7 introduces a couple of “advanced” topics: managing patches with git and reviewing patches posted by others.

Section 8 concludes the document with pointers to sources for more information on kernel development.

1.2: What This Document Is About

The Linux kernel, at over 6 million lines of code and well over 1000 active contributors, is one of the largest and most active free software projects in existence. Since its humble beginning in 1991, this kernel has evolved into a best-of-breed operating system component which runs on pocket-sized digital music players, desktop PCs, the largest supercomputers in existence, and all types of systems in between. It is a robust, efficient, and scalable solution for almost any situation.

With the growth of Linux has come an increase in the number of developers (and companies) wishing to participate in its development. Hardware vendors want to ensure that Linux supports their products well, making those products attractive to Linux users. Embedded systems vendors, who use Linux as a component in an integrated product, want Linux to be as capable and well-suited to the task at hand as possible. Distributors and other software vendors who base their products on Linux have a clear interest in the capabilities, performance, and reliability of the Linux kernel. And end users, too, will often wish to change Linux to make it better suit their needs.

One of the most compelling features of Linux is that it is accessible to these developers; anybody with the requisite skills can improve Linux and influence the direction of its development. Proprietary products cannot offer this kind of openness, which is a characteristic of the free software process. But, if anything, the kernel is even more open than most other free software projects. A typical three-month kernel development cycle can involve over 1000 developers working for more than 100 different companies (or for no company at all).

Working with the kernel development community is not especially hard. But, that notwithstanding, many potential contributors have experienced difficulties when trying to do kernel work. The kernel community has evolved its own distinct ways of operating which allow it to function smoothly (and produce a high-quality product) in an environment where thousands of lines of code are being changed every day. So it is not surprising that Linux kernel development process differs greatly from proprietary development methods.

The kernel's development process may come across as strange and intimidating to new developers, but there are good reasons and solid experience behind it. A developer who does not understand the kernel community's ways (or, worse, who tries to flout or circumvent them) will have a frustrating experience in store. The development community, while being helpful to those who are trying to learn, has little time for those who will not listen or who do not care about the development process.

It is hoped that those who read this document will be able to avoid that frustrating experience. There is a lot of material here, but the effort involved in reading it will be repaid in short order. The development community is always in need of developers who will help to make the kernel better; the following text should help you – or those who work for you – join our community.

1.3: Credits

This document was written by Jonathan Corbet, corbet@lwn.net. It has been improved by comments from James Berry, Alex Chiang, Roland Dreier, Randy Dunlap, Jake Edge, Jiri Kosina, Matt Mackall, Amanda McPherson, Andrew Morton, and Jochen Voß.

This work was supported by the Linux Foundation; thanks especially to Amanda McPherson, who saw the value of this effort and made it all happen.

1.4: The Importance Of Getting Code Into The Mainline

Some companies and developers occasionally wonder why they should bother learning how to work with the kernel community and get their code into the mainline kernel (the “mainline” being the kernel maintained by Linus Torvalds and used as a base by Linux distributors). In the short term, contributing code can look like an avoidable expense; it seems easier to just keep the code separate and support users directly. The truth of the matter is that keeping code separate (“out of tree”) is a false economy.

As a way of illustrating the costs of out-of-tree code, here are a few relevant aspects of the kernel development process; most of these will be discussed in greater detail later in this document. Consider:

- Code which has been merged into the mainline kernel is available to all Linux users. It will automatically be present on all distributions which enable it. There is no need for driver disks, downloads, or the hassles of supporting multiple versions of multiple distributions; it all just works, for the developer and for the user. Incorporation into the mainline solves a large number of distribution and support problems.
- While kernel developers strive to maintain a stable interface to user space, the internal kernel API is in constant flux. The lack of a stable internal interface is a deliberate design decision; it allows fundamental improvements to be made at any time and results in higher-quality code. But one result of that policy is that any out-of-tree code requires constant upkeep if it is to work with new kernels. Maintaining out-of-tree code requires significant amounts of work just to keep that code working.

Code which is in the mainline, instead, does not require this work as the result of a simple rule requiring developers to fix any code which breaks as the result of an API change. So code which has been merged into the mainline has significantly lower maintenance costs.

- Beyond that, code which is in the kernel will often be improved by other developers. Surprising results can come from empowering your user community and customers to improve your product.
- Kernel code is subjected to review, both before and after merging into the mainline. No matter how strong the original developer’s skills are, this review process invariably finds ways in which the code can be improved. Often review finds severe bugs and security problems. This is especially true for code which has been developed in an closed environment; such code benefits strongly from review by outside developers. Out-of-tree code is lower-quality code.
- Participation in the development process is your way to influence the direction of kernel development. Users who complain from the sidelines are heard, but active developers have a stronger voice – and the ability to implement changes which make the kernel work better for their needs.
- When code is maintained separately, the possibility that a third party will contribute a different implementation of a similar feature always exists. Should that happen, getting your code merged will become much harder – to the point of impossibility. Then you will be faced with the unpleasant alternatives of either (1) maintaining a nonstandard feature out of tree indefinitely, or (2) abandoning your code and migrating your users over to the in-tree version.
- Contribution of code is the fundamental action which makes the whole process work. By contributing your code you can add new functionality to the kernel and provide capabilities and examples which are of use to other kernel developers. If you have developed code for Linux (or are thinking about doing so), you clearly have an interest in the continued success of this platform; contributing code is one of the best ways to help ensure that success.

All of the reasoning above applies to any out-of-tree kernel code, including code which is distributed in proprietary, binary-only form.

There are, however, additional factors which should be taken into account before considering any sort of binary-only kernel code distribution. These include:

- The legal issues around the distribution of proprietary kernel modules are cloudy at best; quite a few kernel copyright holders believe that most binary-only modules are derived products of the kernel and that, as a result, their distribution is a violation of the GNU General Public license (about which more will be said below). Your author is not a lawyer, and nothing in this document can possibly be considered to be legal advice. The true legal status of closed-source modules can only be determined by the courts. But the uncertainty which haunts those modules is there regardless.
- Binary modules greatly increase the difficulty of debugging kernel problems, to the point that most kernel developers will not even try. So the distribution of binary-only modules will make it harder for your users to get support from the community.
- Support is also harder for distributors of binary-only modules, who must provide a version of the module for every distribution and every kernel version they wish to support. Dozens of builds of a single module can be required to provide reasonably comprehensive coverage, and your users will have to upgrade your module separately every time they upgrade their kernel.
- Everything that was said above about code review applies doubly to closed-source code. Since this code is not available at all, it cannot have been reviewed by the community and will, beyond doubt, have serious problems.

Makers of embedded systems, in particular, may be tempted to disregard much of what has been said in this section in the belief that they are shipping a self-contained product which uses a frozen kernel version and requires no more development after its release. This argument misses the value of widespread code review and the value of allowing your users to add capabilities to your product. But these products, too, have a limited commercial life, after which a new version must be released. At that point, vendors whose code is in the mainline and well maintained will be much better positioned to get the new product ready for market quickly.

1.5: Licensing

Code is contributed to the Linux kernel under a number of licenses, but all code must be compatible with version 2 of the GNU General Public License (GPLv2), which is the license covering the kernel distribution as a whole. In practice, that means that all code contributions are covered either by GPLv2 (with, optionally, language allowing distribution under later versions of the GPL) or the three-clause BSD license. Any contributions which are not covered by a compatible license will not be accepted into the kernel.

Copyright assignments are not required (or requested) for code contributed to the kernel. All code merged into the mainline kernel retains its original ownership; as a result, the kernel now has thousands of owners.

One implication of this ownership structure is that any attempt to change the licensing of the kernel is doomed to almost certain failure. There are few practical scenarios where the agreement of all copyright holders could be obtained (or their code removed from the kernel). So, in particular, there is no prospect of a migration to version 3 of the GPL in the foreseeable future.

It is imperative that all code contributed to the kernel be legitimately free software. For that reason, code from anonymous (or pseudonymous) contributors will not be accepted. All contributors are required to “sign off” on their code, stating that the code can be distributed with the kernel under the GPL. Code which has not been licensed as free software by its owner, or which risks creating copyright-related problems for the kernel (such as code which derives from reverse-engineering efforts lacking proper safeguards) cannot be contributed.

Questions about copyright-related issues are common on Linux development mailing lists. Such questions will normally receive no shortage of answers, but one should bear in mind that the people answering those questions are not lawyers and cannot provide legal advice. If you have legal questions relating to Linux source code, there is no substitute for talking with a lawyer who understands this field. Relying on answers obtained on technical mailing lists is a risky affair.

2: How The Development Process Works

Linux kernel development in the early 1990's was a pretty loose affair, with relatively small numbers of users and developers involved. With a user base in the millions and with some 2,000 developers involved over the course of one year, the kernel has since had to evolve a number of processes to keep development happening smoothly. A solid understanding of how the process works is required in order to be an effective part of it.

2.1: The Big Picture

The kernel developers use a loosely time-based release process, with a new major kernel release happening every two or three months. The recent release history looks like this:

2.6.26	July 13, 2008
2.6.25	April 16, 2008
2.6.24	January 24, 2008
2.6.23	October 9, 2007
2.6.22	July 8, 2007
2.6.21	April 25, 2007
2.6.20	February 7, 2007

Every 2.6.x release is a major kernel release with new features, internal API changes, and more. A typical 2.6 release can contain over 10,000 changesets with changes to several hundred thousand lines of code. 2.6 is thus the leading edge of Linux kernel development; the kernel uses a rolling development model which is continually integrating major changes.

A relatively straightforward discipline is followed with regard to the merging of patches for each release. At the beginning of each development cycle, the "merge window" is said to be open. At that time, code which is deemed to be sufficiently stable (and which is accepted by the development community) is merged into the mainline kernel. The bulk of changes for a new development cycle (and all of the major changes) will be merged during this time, at a rate approaching 1,000 changes ("patches," or "changesets") per day.

(As an aside, it is worth noting that the changes integrated during the merge window do not come out of thin air; they have been collected, tested, and staged ahead of time. How that process works will be described in detail later on).

The merge window lasts for two weeks. At the end of this time, Linus Torvalds will declare that the window is closed and release the first of the "rc" kernels. For the kernel which is destined to be 2.6.26, for example, the release which happens at the end of the merge window will be called 2.6.26-rc1. The -rc1 release is the signal that the time to merge new features has passed, and that the time to stabilize the next kernel has begun.

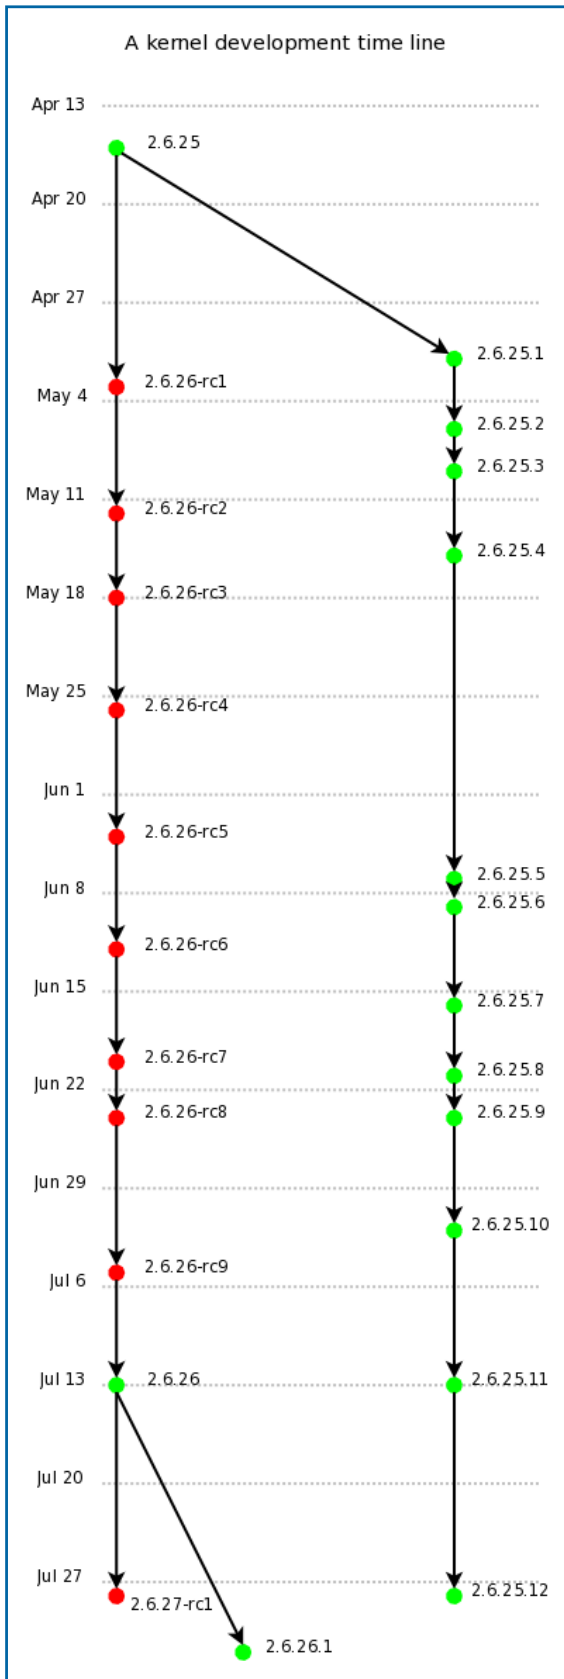
Over the next six to ten weeks, only patches which fix problems should be submitted to the mainline. On occasion a more significant change will be allowed, but such occasions are rare; developers who try to merge new features outside of the merge window tend to get an unfriendly reception.

As a general rule, if you miss the merge window for a given feature, the best thing to do is to wait for the next development cycle. (An occasional exception is made for drivers for previously-unsupported hardware; if they touch no in-tree code, they cannot cause regressions and should be safe to add at any time).

As fixes make their way into the mainline, the patch rate will slow over time. Linus releases new -rc kernels about once a week; a normal series will get up to somewhere between -rc6 and -rc9 before the kernel is considered to be sufficiently stable and the final 2.6.x release is made. At that point the whole process starts over again.

As an example, here is how the 2.6.25 development cycle went (all dates in 2008):

January 24	2.6.24 stable release
February 10	2.6.25-rc1, merge window closes
February 15	2.6.25-rc2
February 24	2.6.25-rc3
March 4	2.6.25-rc4
March 9	2.6.25-rc5
March 16	2.6.25-rc6
March 25	2.6.25-rc7
April 1	2.6.25-rc8
April 11	2.6.25-rc9
April 16	2.6.25 stable release



How do the developers decide when to close the development cycle and create the stable release? The most significant metric used is the list of regressions from previous releases. No bugs are welcome, but those which break systems which worked in the past are considered to be especially serious. For this reason, patches which cause regressions are looked upon unfavorably and are quite likely to be reverted during the stabilization period.

The developers' goal is to fix all known regressions before the stable release is made. In the real world, this kind of perfection is hard to achieve; there are just too many variables in a project of this size. There comes a point where delaying the final release just makes the problem worse; the pile of changes waiting for the next merge window will grow larger, creating even more regressions the next time around. So most 2.6.x kernels go out with a handful of known regressions though, hopefully, none of them are serious.

Once a stable release is made, its ongoing maintenance is passed off to the "stable team," currently comprised of Greg Kroah-Hartman and Chris Wright. The stable team will release occasional updates to the stable release using the 2.6.x.y numbering scheme. To be considered for an update release, a patch must (1) fix a significant bug, and (2) already be merged into the mainline for the next development kernel. Continuing our 2.6.25 example, the history (as of this writing) is:

- May 1 2.6.25.1
- May 6 2.6.25.2
- May 9 2.6.25.3
- May 15 2.6.25.4
- June 7 2.6.25.5
- June 9 2.6.25.6
- June 16 2.6.25.7
- June 21 2.6.25.8
- June 24 2.6.25.9

Stable updates for a given kernel are made for approximately six months; after that, the maintenance of stable releases is solely the responsibility of the distributors which have shipped that particular kernel.

2.2: The Lifecycle Of A Patch

Patches do not go directly from the developer's keyboard into the mainline kernel. There is, instead, a somewhat involved (if somewhat informal) process designed to ensure that each patch is reviewed for quality and that each patch implements a change which is desirable to have in the mainline. This process can happen quickly for minor fixes, or, in the case of large and controversial changes, go on for years. Much developer frustration comes from a lack of understanding of this process or from attempts to circumvent it.

In the hopes of reducing that frustration, this document will describe how a patch gets into the kernel. What follows below is an introduction which describes the process in a somewhat idealized way.

A much more detailed treatment will come in later sections. The stages that a patch goes through are, generally:

- **Design.** This is where the real requirements for the patch – and the way those requirements will be met – are laid out. Design work is often done without involving the community, but it is better to do this work in the open if at all possible; it can save a lot of time redesigning things later.
- **Early review.** Patches are posted to the relevant mailing list, and developers on that list reply with any comments they may have. This process should turn up any major problems with a patch if all goes well.
- **Wider review.** When the patch is getting close to ready for mainline inclusion, it will be accepted by a relevant subsystem maintainer – though this acceptance is not a guarantee that the patch will make it all the way to the mainline. The patch will show up in the maintainer's subsystem tree and into the staging trees (described below). When the process works, this step leads to more extensive review of the patch and the discovery of any problems resulting from the integration of this patch with work being done by others.
- **Merging into the mainline.** Eventually, a successful patch will be merged into the mainline repository managed by Linus Torvalds. More comments and/or problems may surface at this time; it is important that the developer be responsive to these and fix any issues which arise.

- **Stable release.** The number of users potentially affected by the patch is now large, so, once again, new problems may arise.
- **Long-term maintenance.** While it is certainly possible for a developer to forget about code after merging it, that sort of behavior tends to leave a poor impression in the development community. Merging code eliminates some of the maintenance burden, in that others will fix problems caused by API changes. But the original developer should continue to take responsibility for the code if it is to remain useful in the longer term.

One of the largest mistakes made by kernel developers (or their employers) is to try to cut the process down to a single “merging into the mainline” step. This approach invariably leads to frustration for everybody involved.

2.3: How Patches Get Into The Kernel

There is exactly one person who can merge patches into the mainline kernel repository: Linus Torvalds. But, of the over 12,000 patches which went into the 2.6.25 kernel, only 250 (around 2%) were directly chosen by Linus himself. The kernel project has long since grown to a size where no single developer could possibly inspect and select every patch unassisted. The way the kernel developers have addressed this growth is through the use of a lieutenant system built around a chain of trust. The kernel code base is logically broken down into a set of subsystems: networking, specific architecture support, memory management, video devices, etc. Most subsystems have a designated maintainer, a developer who has overall responsibility for the code within that subsystem.

These subsystem maintainers are the gatekeepers (in a loose way) for the portion of the kernel they manage; they are the ones who will (usually) accept a patch for inclusion into the mainline kernel.

Subsystem maintainers each manage their own version of the kernel source tree, usually (but certainly not always) using the git source management tool. Tools like git (and related tools like quilt or mercurial) allow maintainers to track a list of patches, including authorship information and other metadata. At any given time, the maintainer can identify which patches in his or her repository are not found in the mainline.

When the merge window opens, top-level maintainers will ask Linus to “pull” the patches they have selected for merging from their repositories. If Linus agrees, the stream of patches will flow up into his repository, becoming part of the mainline kernel. The amount of attention that Linus pays to specific patches received in a pull operation varies. It is clear that, sometimes, he looks quite closely. But, as a general rule, Linus trusts the subsystem maintainers to not send bad patches upstream.

Subsystem maintainers, in turn, can pull patches from other maintainers. For example, the networking tree is built from patches which accumulated first in trees dedicated to network device drivers, wireless networking, etc. This chain of repositories can be arbitrarily long, though it rarely exceeds two or three links. Since each maintainer in the chain trusts those managing lower-level trees, this process is known as the “chain of trust.”

Clearly, in a system like this, getting patches into the kernel depends on finding the right maintainer. Sending patches directly to Linus is not normally the right way to go.

2.4: Staging Trees

The chain of subsystem trees guides the flow of patches into the kernel, but it also raises an interesting question: what if somebody wants to look at all of the patches which are being prepared for the next merge window? Developers will be interested in what other changes are pending to see whether there are any conflicts to worry about; a patch which changes a core kernel function prototype, for example, will conflict with any other patches which use the older form of that function. Reviewers and testers want access to the changes in their integrated form before all of those changes land in the mainline kernel. One could pull changes from all of the interesting subsystem trees, but that would be a big and error-prone job.

The answer comes in the form of staging trees, where subsystem trees are collected for testing and review. The older of these trees, maintained by Andrew Morton, is called “-mm” (for memory management, which is how it got started). The -mm tree integrates patches from a long list of subsystem trees; it also has some patches aimed at helping with debugging.

Beyond that, -mm contains a significant collection of patches which have been selected by Andrew directly. These patches may have been posted on a mailing list, or they may apply to a part of the kernel for which there is no designated subsystem tree. As a result, -mm operates as a sort of subsystem tree of last resort; if there is no other obvious path for a patch into the mainline, it is likely to end up in -mm. Miscellaneous patches which accumulate in - mm will eventually either be forwarded on to an appropriate subsystem tree or be sent directly to Linus. In a typical development cycle, approximately 10% of the patches going into the mainline get there via -mm.

The current -mm patch can always be found from the front page of <http://kernel.org/>

Those who want to see the current state of -mm can get the “-mm of the moment” tree, found at: <http://userweb.kernel.org/~akpm/mmotm/>

Use of the MMOTM tree is likely to be a frustrating experience, though; there is a definite chance that it will not even compile. The other staging tree, started more recently, is linux-next, maintained by Stephen Rothwell. The linux-next tree is, by design, a snapshot of what the mainline is expected to look like after the next merge window closes. Linux-next trees are announced on the linux-kernel and linux-next mailing lists when they are assembled; they can be downloaded from:

<http://www.kernel.org/pub/linux/kernel/people/sfr/linux-next/>

Some information about linux-next has been gathered at:

<http://linux.f-seidel.de/linux-next/pmwiki/>

How the linux-next tree will fit into the development process is still changing. As of this writing, the first full development cycle involving linux-next (2.6.26) is coming to an end; thus far, it has proved to be a valuable resource for finding and fixing integration problems before the beginning of the merge window. See <http://lwn.net/Articles/287155/> for more information on how linux-next has worked to set up the 2.6.27 merge window.

Some developers have begun to suggest that linux-next should be used as the target for future development as well. The linux-next tree does tend to be far ahead of the mainline and is more representative of the tree into which any new work will be merged.

The downside to this idea is that the volatility of linux-next tends to make it a difficult development target. See <http://lwn.net/Articles/289013/> for more information on this topic, and stay tuned; much is still in flux where linux-next is involved.

2.5: Tools

As can be seen from the above text, the kernel development process depends heavily on the ability to herd collections of patches in various directions. The whole thing would not work anywhere near as well as it does without suitably powerful tools. Tutorials on how to use these tools are well beyond the scope of this document, but there is space for a few pointers.

By far the dominant source code management system used by the kernel community is git. Git is one of a number of distributed version control systems being developed in the free software community. It is well tuned for kernel development, in that it performs quite well when dealing with large repositories and large numbers of patches. It also has a reputation for being difficult to learn and use, though it has gotten better over time. Some sort of familiarity with git is almost a requirement for kernel developers; even if they do not use it for their own work, they'll need git to keep up with what other developers (and the mainline) are doing.

Git is now packaged by almost all Linux distributions. There is a home page at <http://git.or.cz/>

That page has pointers to documentation and tutorials. One should be aware, in particular, of the Kernel Hacker's Guide to git, which has information specific to kernel development:

<http://linux.yyz.us/git-howto.html>

Among the kernel developers who do not use git, the most popular choice is almost certainly Mercurial:

<http://www.selenic.com/mercurial/>

Mercurial shares many features with git, but it provides an interface which many find easier to use.

The other tool worth knowing about is Quilt:

<http://savannah.nongnu.org/projects/quilt/>

Quilt is a patch management system, rather than a source code management system. It does not track history over time; it is, instead, oriented toward tracking a specific set of changes against an evolving code base. Some major subsystem maintainers use quilt to manage patches intended to go upstream. For the management of certain kinds of trees (-mm, for example), quilt is the best tool for the job.

2.6: Mailing Lists

A great deal of Linux kernel development work is done by way of mailing lists. It is hard to be a fully-functioning member of the community without joining at least one list somewhere. But Linux mailing lists also represent a potential hazard to developers, who risk getting buried under a load of electronic mail, running afoul of the conventions used on the Linux lists, or both.

Most kernel mailing lists are run on vger.kernel.org; the master list can be found at:

<http://vger.kernel.org/vger-lists.html>

There are lists hosted elsewhere, though; a number of them are at lists.redhat.com.

The core mailing list for kernel development is, of course, [linux-kernel](mailto:linux-kernel@vger.kernel.org). This list is an intimidating place to be; volume can reach 500 messages per day, the amount of noise is high, the conversation can be severely technical, and participants are not always concerned with showing a high degree of politeness. But there is no other place where the kernel development community comes together as a whole; developers who avoid this list will miss important information.

There are a few hints which can help with [linux-kernel](mailto:linux-kernel@vger.kernel.org) survival:

- Have the list delivered to a separate folder, rather than your main mailbox. One must be able to ignore the stream for sustained periods of time.
- Do not try to follow every conversation – nobody else does. It is important to filter on both the topic of interest (though note that long-running conversations can drift away from the original subject without changing the email subject line) and the people who are participating.
- Do not feed the trolls. If somebody is trying to stir up an angry response, ignore them.
- When responding to [linux-kernel](mailto:linux-kernel@vger.kernel.org) email (or that on other lists) preserve the Cc: header for all involved. In the absence of a strong reason (such as an explicit request), you should never remove recipients. Always make sure that the person you are responding to is in the Cc: list. This convention also makes it unnecessary to explicitly ask to be copied on replies to your postings.

- Search the list archives (and the net as a whole) before asking questions. Some developers can get impatient with people who clearly have not done their homework.
- Avoid top-posting (the practice of putting your answer above the quoted text you are responding to). It makes your response harder to read and makes a poor impression.
- Ask on the correct mailing list. Linux-kernel may be the general meeting point, but it is not the best place to find developers from all subsystems.

The last point – finding the correct mailing list – is a common place for beginning developers to go wrong. Somebody who asks a networking-related question on linux-kernel will almost certainly receive a polite suggestion to ask on the netdev list instead, as that is the list frequented by most networking developers. Other lists exist for the SCSI, video4linux, IDE, filesystem, etc. subsystems. The best place to look for mailing lists is in the MAINTAINERS file packaged with the kernel source.

2.7: Getting Started With Kernel Development

Questions about how to get started with the kernel development process are common – from both individuals and companies. Equally common are missteps which make the beginning of the relationship harder than it has to be.

Companies often look to hire well-known developers to get a development group started. This can, in fact, be an effective technique. But it also tends to be expensive and does not do much to grow the pool of experienced kernel developers. It is possible to bring in-house developers up to speed on Linux kernel development, given the investment of a bit of time. Taking this time can endow an employer with a group of developers who understand the kernel and the company both, and who can help to train others as well.

Over the medium term, this is often the more profitable approach. Individual developers are often, understandably, at a loss for a place to start. Beginning with a large project can be intimidating; one often wants to test the waters with something smaller first. This is the point where some developers jump into the creation of patches fixing spelling errors or minor coding style issues. Unfortunately, such patches create a level of noise which is distracting for the development community as a whole, so, increasingly, they are looked down upon. New developers wishing to introduce themselves to the community will not get the sort of reception they wish for by these means.

Andrew Morton gives this advice for aspiring kernel developers:

The #1 project for all kernel beginners should surely be “make sure that the kernel runs perfectly at all times on all machines which you can lay your hands on”. Usually the way to do this is to work with others on getting things fixed up (this can require persistence!) but that’s fine – it’s a part of kernel development. (<http://lwn.net/Articles/283982/>)

In the absence of obvious problems to fix, developers are advised to look at the current lists of regressions and open bugs in general. There is never any shortage of issues in need of fixing; by addressing these issues, developers will gain experience with the process while, at the same time, building respect with the rest of the development community.

3: Early-Stage Planning

When contemplating a Linux kernel development project, it can be tempting to jump right in and start coding. As with any significant project, though, much of the groundwork for success is best laid before the first line of code is written. Some time spent in early planning and communication can save far more time later on.

3.1: Specifying The Problem

Like any engineering project, a successful kernel enhancement starts with a clear description of the problem to be solved. In some cases, this step is easy: when a driver is needed for a specific piece of hardware, for example. In others, though, it is tempting to confuse the real problem with the proposed solution, and that can lead to difficulties.

Consider an example: some years ago, developers working with Linux audio sought a way to run applications without dropouts or other artifacts caused by excessive latency in the system. The solution they arrived at was a kernel module intended to hook into the Linux Security Module (LSM) framework; this module could be configured to give specific applications access to the realtime scheduler. This module was implemented and sent to the linux-kernel mailing list, where it immediately ran into problems.

To the audio developers, this security module was sufficient to solve their immediate problem. To the wider kernel community, though, it was seen as a misuse of the LSM framework (which is not intended to confer privileges onto processes which they would not otherwise have) and a risk to system stability. Their preferred solutions involved realtime scheduling access via the rlimit mechanism for the short term, and ongoing latency reduction work in the long term.

The audio community, however, could not see past the particular solution they had implemented; they were unwilling to accept alternatives. The resulting disagreement left those developers feeling disillusioned with the entire kernel development process; one of them went back to an audio list and posted this:

There are a number of very good Linux kernel developers, but they tend to get outshouted by a large crowd of arrogant fools. Trying to communicate user requirements to these people is a waste of time. They are much too “intelligent” to listen to lesser mortals.
(<http://lwn.net/Articles/131776/>)

The reality of the situation was different; the kernel developers were far more concerned about system stability, long-term maintenance, and finding the right solution to the problem than they were with a specific module. The moral of the story is to focus on the problem – not a specific solution – and to discuss it with the development community before investing in the creation of a body of code.

So, when contemplating a kernel development project, one should obtain answers to a short set of questions:

- What, exactly, is the problem which needs to be solved?
- Who are the users affected by this problem? Which use cases should the solution address?
- How does the kernel fall short in addressing that problem now?

Only then does it make sense to start considering possible solutions.

3.2: Early Discussion

When planning a kernel development project, it makes great sense to hold discussions with the community before launching into implementation. Early communication can save time and trouble in a number of ways:

- It may well be that the problem is addressed by the kernel in ways which you have not understood. The Linux kernel is large and has a number of features and capabilities which are not immediately obvious. Not all kernel capabilities are documented as well as one might like, and it is easy to miss things. Your author has seen the posting of a complete driver which duplicated an existing driver that the new author had been unaware of. Code which reinvents existing wheels is not only wasteful it will also not be accepted into the mainline kernel.

- There may be elements of the proposed solution which will not be acceptable for mainline merging. It is better to find out about problems like this before writing the code.
- It's entirely possible that other developers have thought about the problem; they may have ideas for a better solution, and may be willing to help in the creation of that solution.

Years of experience with the kernel development community have taught a clear lesson: kernel code which is designed and developed behind closed doors invariably has problems which are only revealed when the code is released into the community. Sometimes these problems are severe, requiring months or years of effort before the code can be brought up to the kernel community's standards. Some examples include:

- The Devicescape network stack was designed and implemented for single-processor systems. It could not be merged into the mainline until it was made suitable for multiprocessor systems. Retrofitting locking and such into code is a difficult task; as a result, the merging of this code (now called mac80211) was delayed for over a year.
- The Reiser4 filesystem included a number of capabilities which, in the core kernel developers' opinion, should have been implemented in the virtual filesystem layer instead. It also included features which could not easily be implemented without exposing the system to user-caused deadlocks. The late revelation of these problems – and refusal to address some of them – has caused Reiser4 to stay out of the mainline kernel.
- The AppArmor security module made use of internal virtual filesystem data structures in ways which were considered to be unsafe and unreliable. This code has since been significantly reworked, but remains outside of the mainline.

In each of these cases, a great deal of pain and extra work could have been avoided with some early discussion with the kernel developers.

3.3: Who Do You Talk To?

When developers decide to take their plans public, the next question will be: where do we start? The answer is to find the right mailing list(s) and the right maintainer. For mailing lists, the best approach is to look in the MAINTAINERS file for a relevant place to post. If there is a suitable subsystem list, posting there is often preferable to posting on linux-kernel; you are more likely to reach developers with expertise in the relevant subsystem and the environment may be more supportive.

Finding maintainers can be a bit harder. Again, the MAINTAINERS file is the place to start. That file tends to not always be up to date, though, and not all subsystems are represented there. The person listed in the MAINTAINERS file may, in fact, not be the person who is actually acting in that role currently. So, when there is doubt about who to contact, a useful trick is to use git (and “git log” in particular) to see who is currently active within the subsystem of interest. Look at who is writing patches, and who, if anybody, is attaching Signed-off-by lines to those patches. Those are the people who will be best placed to help with a new development project.

If all else fails, talking to Andrew Morton can be an effective way to track down a maintainer for a specific piece of code.

3.4: When To Post?

If possible, posting your plans during the early stages can only be helpful. Describe the problem being solved and any plans that have been made on how the implementation will be done. Any information you can provide can help the development community provide useful input on the project.

One discouraging thing which can happen at this stage is not a hostile reaction, but, instead, little or no reaction at all. The sad truth of the matter is (1) kernel developers tend to be busy, (2) there is no shortage of people with grand plans and little code (or even prospect of code) to back them up, and (3) nobody is obligated to review or comment on ideas posted by others. If a request-for-comments posting yields little in the way of comments, do not assume that it means there is no interest in the project. Unfortunately, you also cannot assume that there are no problems with your idea. The best thing to do in this situation is to proceed, keeping the community informed as you go.

3.5: Getting Official Buy-In

If your work is being done in a corporate environment – as most Linux kernel work is – you must, obviously, have permission from suitably empowered managers before you can post your company’s plans or code to a public mailing list. The posting of code which has not been cleared for release under a GPL-compatible license can be especially problematic; the sooner that a company’s management and legal staff can agree on the posting of a kernel development project, the better off everybody involved will be.

Some readers may be thinking at this point that their kernel work is intended to support a product which does not yet have an officially acknowledged existence. Revealing their employer’s plans on a public mailing list may not be a viable option. In cases like this, it is worth considering whether the secrecy is really necessary; there is often no real need to keep development plans behind closed doors.

That said, there are also cases where a company legitimately cannot disclose its plans early in the development process. Companies with experienced kernel developers may choose to proceed in an open-loop manner on the assumption that they will be able to avoid serious integration problems later. For companies without that sort of in-house expertise, the best option is often to hire an outside developer to review the plans under a non-disclosure agreement. The Linux Foundation operates an NDA program designed to help with this sort of situation; more information can be found at:

http://www.linuxfoundation.org/en/NDA_program

This kind of review is often enough to avoid serious problems later on without requiring public disclosure of the project.

4: Getting The Code Right

While there is much to be said for a solid and community-oriented design process, the proof of any kernel development project is in the resulting code. It is the code which will be examined by other developers and merged (or not) into the mainline tree. So it is the quality of this code which will determine the ultimate success of the project.

This section will examine the coding process. We’ll start with a look at a number of ways in which kernel developers can go wrong. Then the focus will shift toward doing things right and the tools which can help in that quest.

4.1: Pitfalls

Coding Style

The kernel has long had a standard coding style, described in Documentation/CodingStyle. For much of that time, the policies described in that file were taken as being, at most, advisory. As a result, there is a substantial amount of code in the kernel which does not meet the coding style guidelines. The presence of that code leads to two independent hazards for kernel developers.

The first of these is to believe that the kernel coding standards do not matter and are not enforced. The truth of the matter is that adding new code to the kernel is very difficult if that code is not coded according to the standard; many developers will request that the code be reformatted before they will even review it. A code base as large as the kernel requires some uniformity of code to make it possible for developers to quickly understand any part of it. So there is no longer room for strangely-formatted code.

Occasionally, the kernel's coding style will run into conflict with an employer's mandated style. In such cases, the kernel's style will have to win before the code can be merged. Putting code into the kernel means giving up a degree of control in a number of ways – including control over how the code is formatted.

The other trap is to assume that code which is already in the kernel is urgently in need of coding style fixes. Developers may start to generate reformatting patches as a way of gaining familiarity with the process, or as a way of getting their name into the kernel changelogs – or both. But pure coding style fixes are seen as noise by the development community; they tend to get a chilly reception. So this type of patch is best avoided. It is natural to fix the style of a piece of code while working on it for other reasons, but coding style changes should not be made for their own sake.

The coding style document also should not be read as an absolute law which can never be transgressed. If there is a good reason to go against the style (a line which becomes far less readable if split to fit within the 80-column limit, for example), just do it.

Abstraction Layers

Computer Science professors teach students to make extensive use of abstraction layers in the name of flexibility and information hiding. Certainly the kernel makes extensive use of abstraction; no project involving several million lines of code could do otherwise and survive.

But experience has shown that excessive or premature abstraction can be just as harmful as premature optimization. Abstraction should be used to the level required and no further.

At a simple level, consider a function which has an argument which is always passed as zero by all callers. One could retain that argument just in case somebody eventually needs to use the extra flexibility that it provides. By that time, though, chances are good that the code which implements this extra argument has been broken in some subtle way which was never noticed – because it has never been used.

Or, when the need for extra flexibility arises, it does not do so in a way which matches the programmer's early expectation. Kernel developers will routinely submit patches to remove unused arguments; they should, in general, not be added in the first place.

Abstraction layers which hide access to hardware – often to allow the bulk of a driver to be used with multiple operating systems – are especially frowned upon. Such layers obscure the code and may impose a performance penalty; they do not belong in the Linux kernel.

On the other hand, if you find yourself copying significant amounts of code from another kernel subsystem, it is time to ask whether it would, in fact, make sense to pull out some of that code into a separate library or to implement that functionality at a higher level. There is no value in replicating the same code throughout the kernel.

#ifdef and Preprocessor Use In General

The C preprocessor seems to present a powerful temptation to some C programmers, who see it as a way to efficiently encode a great deal of flexibility into a source file. But the preprocessor is not C, and heavy use of it results in code which is much harder for others to read and harder for the compiler to check for correctness. Heavy preprocessor use is almost always a sign of code which needs some cleanup work.

Conditional compilation with `#ifdef` is, indeed, a powerful feature, and it is used within the kernel. But there is little desire to see code which is sprinkled liberally with `#ifdef` blocks. As a general rule, `#ifdef` use should be confined to header files whenever possible. Conditionally-compiled code can be confined to functions which, if the code is not to be present, simply become empty. The compiler will then quietly optimize out the call to the empty function. The result is far cleaner code which is easier to follow.

C preprocessor macros present a number of hazards, including possible multiple evaluation of expressions with side effects and no type safety. If you are tempted to define a macro, consider creating an inline function instead. The code which results will be the same, but inline functions are easier to read, do not evaluate their arguments multiple times, and allow the compiler to perform type checking on the arguments and return value.

Inline Functions

Inline functions present a hazard of their own, though. Programmers can become enamored of the perceived efficiency inherent in avoiding a function call and fill a source file with inline functions. Those functions, however, can actually reduce performance. Since their code is replicated at each call site, they end up bloating the size of the compiled kernel.

That, in turn, creates pressure on the processor's memory caches, which can slow execution dramatically. Inline functions, as a rule, should be quite small and relatively rare. The cost of a function call, after all, is not that high; the creation of large numbers of inline functions is a classic example of premature optimization.

In general, kernel programmers ignore cache effects at their peril. The classic time/space tradeoff taught in beginning data structures classes often does not apply to contemporary hardware. Space *is* time, in that a larger program will run slower than one which is more compact.

Locking

In May, 2006, the "Devicescape" networking stack was, with great fanfare, released under the GPL and made available for inclusion in the mainline kernel. This donation was welcome news; support for wireless networking in Linux was considered substandard at best, and the Devicescape stack offered the promise of fixing that situation. Yet, this code did not actually make it into the mainline until June, 2007 (2.6.22). What happened?

This code showed a number of signs of having been developed behind corporate doors. But one large problem in particular was that it was not designed to work on multiprocessor systems. Before this networking stack (now called mac80211) could be merged, a locking scheme needed to be retrofitted onto it.

Once upon a time, Linux kernel code could be developed without thinking about the concurrency issues presented by multiprocessor systems. Now, however, this document is being written on a dual-core laptop. Even on single-processor systems, work being done to improve responsiveness will raise the level of concurrency within the kernel. The days when kernel code could be written without thinking about locking are long past.

Any resource (data structures, hardware registers, etc.) which could be accessed concurrently by more than one thread must be protected by a lock. New code should be written with this requirement in mind; retrofitting locking after the fact is a rather more difficult task. Kernel developers should take the time to understand the available locking primitives well enough to pick the right tool for the job. Code which shows a lack of attention to concurrency will have a difficult path into the mainline.

Regressions

One final hazard worth mentioning is this: it can be tempting to make a change (which may bring big improvements) which causes something to break for existing users. This kind of change is called a "regression," and regressions have become most unwelcome in the mainline kernel. With few exceptions, changes which cause regressions will be backed out if the regression cannot be fixed in a timely manner. Far better to avoid the regression in the first place.

It is often argued that a regression can be justified if it causes things to work for more people than it creates problems for. Why not make a change if it brings new functionality to ten systems for each one it breaks? The best answer to this question was expressed by Linus in July, 2007:

So we don't fix bugs by introducing new problems. That way lies madness, and nobody ever knows if you actually make any real progress at all. Is it two steps forwards, one step back, or one step forward and two steps back?
(<http://lwn.net/Articles/243460/>)

An especially unwelcome type of regression is any sort of change to the user-space ABI. Once an interface has been exported to user space, it must be supported indefinitely. This fact makes the creation of user-space interfaces particularly challenging: since they cannot be changed in incompatible ways, they must be done right the first time. For this reason, a great deal of thought, clear documentation, and wide review for user-space interfaces is always required.

4.2: Code Checking Tools

For now, at least, the writing of error-free code remains an ideal that few of us can reach. What we can hope to do, though, is to catch and fix as many of those errors as possible before our code goes into the mainline kernel. To that end, the kernel developers have put together an impressive array of tools which can catch a wide variety of obscure problems in an automated way. Any problem caught by the computer is a problem which will not afflict a user later on, so it stands to reason that the automated tools should be used whenever possible.

The first step is simply to heed the warnings produced by the compiler. Contemporary versions of gcc can detect (and warn about) a large number of potential errors. Quite often, these warnings point to real problems. Code submitted for review should, as a rule, not produce any compiler warnings. When silencing warnings, take care to understand the real cause and try to avoid “fixes” which make the warning go away without addressing its cause.

Note that not all compiler warnings are enabled by default. Build the kernel with “make EXTRA_CFLAGS=-W” to get the full set.

The kernel provides several configuration options which turn on debugging features; most of these are found in the “kernel hacking” submenu. Several of these options should be turned on for any kernel used for development or testing purposes. In particular, you should turn on:

- `ENABLE_WARN_DEPRECATED`, `ENABLE_MUST_CHECK`, and `FRAME_WARN` to get an extra set of warnings for problems like the use of deprecated interfaces or ignoring an important return value from a function. The output generated by these warnings can be verbose, but one need not worry about warnings from other parts of the kernel.
- `DEBUG_OBJECTS` will add code to track the lifetime of various objects created by the kernel and warn when things are done out of order. If you are adding a subsystem which creates (and exports) complex objects of its own, consider adding support for the object debugging infrastructure.
- `DEBUG_SLAB` can find a variety of memory allocation and use errors; it should be used on most development kernels.
- `DEBUG_SPINLOCK`, `DEBUG_SPINLOCK_SLEEP`, and `DEBUG_MUTEXES` will find a number of common locking errors.

There are quite a few other debugging options, some of which will be discussed below. Some of them have a significant performance impact and should not be used all of the time. But some time spent learning the available options will likely be paid back many times over in short order.

One of the heavier debugging tools is the locking checker, or “lockdep.” This tool will track the acquisition and release of every lock (spinlock or mutex) in the system, the order in which locks are acquired relative to each other, the current interrupt environment, and more. It can then ensure that locks are always acquired in the same order, that the same interrupt assumptions apply in all situations, and so on. In other words, lockdep can find a number of scenarios in which the system could, on rare occasion, deadlock. This kind of problem can be painful (for both developers and users) in a deployed system; lockdep allows them to be found in an automated manner ahead of time. Code with any sort of non-trivial locking should be run with lockdep enabled before being submitted for inclusion.

As a diligent kernel programmer, you will, beyond doubt, check the return status of any operation (such as a memory allocation) which can fail. The fact of the matter, though, is that the resulting failure recovery paths are, probably, completely untested. Untested code tends to be broken code; you could be much more confident of your code if all those error-handling paths had been exercised a few times.

The kernel provides a fault injection framework which can do exactly that, especially where memory allocations are involved. With fault injection enabled, a configurable percentage of memory allocations will be made to fail; these failures can be restricted to a specific range of code. Running with fault injection enabled allows the programmer to see how the code responds when things go badly. See [Documentation/fault-injection/fault-injection.text](#) for more information on how to use this facility.

Other kinds of errors can be found with the “sparse” static analysis tool. With sparse, the programmer can be warned about confusion between user-space and kernel-space addresses, mixture of big-endian and small-endian quantities, the passing of integer values where a set of bit flags is expected, and so on. Sparse must be installed separately (it can be found at <http://www.kernel.org/pub/software/devel/sparse/> if your distributor does not package it); it can then be run on the code by adding “C=1” to your make command.

Other kinds of portability errors are best found by compiling your code for other architectures. If you do not happen to have an S/390 system or a Blackfin development board handy, you can still perform the compilation step. A large set of cross compilers for x86 systems can be found at <http://www.kernel.org/pub/tools/crosstool/>

Some time spent installing and using these compilers will help avoid embarrassment later.

4.3: Documentation

Documentation has often been more the exception than the rule with kernel development. Even so, adequate documentation will help to ease the merging of new code into the kernel, make life easier for other developers, and will be helpful for your users. In many cases, the addition of documentation has become essentially mandatory.

The first piece of documentation for any patch is its associated changelog. Log entries should describe the problem being solved, the form of the solution, the people who worked on the patch, any relevant effects on performance, and anything else that might be needed to understand the patch.

Any code which adds a new user-space interface – including new sysfs or /proc files – should include documentation of that interface which enables user-space developers to know what they are working with. See Documentation/ABI/README for a description of how this documentation should be formatted and what information needs to be provided.

The file Documentation/kernel-parameters.txt describes all of the kernel's boot-time parameters. Any patch which adds new parameters should add the appropriate entries to this file.

Any new configuration options must be accompanied by help text which clearly explains the options and when the user might want to select them.

Internal API information for many subsystems is documented by way of specially-formatted comments; these comments can be extracted and formatted in a number of ways by the “kernel-doc” script. If you are working within a subsystem which has kerneldoc comments, you should maintain them and add them, as appropriate, for externally-available functions. Even in areas which have not been so documented, there is no harm in adding kerneldoc comments for the future; indeed, this can be a useful activity for beginning kernel developers. The format of these comments, along with some information on how to create kerneldoc templates can be found in the file Documentation/kernel-doc-nano-HOWTO.txt.

Anybody who reads through a significant amount of existing kernel code will note that, often, comments are most notable by their absence. Once again, the expectations for new code are higher than they were in the past; merging uncommented code will be harder. That said, there is little desire for verbosely-commented code. The code should, itself, be readable, with comments explaining the more subtle aspects.

Certain things should always be commented. Uses of memory barriers should be accompanied by a line explaining why the barrier is necessary. The locking rules for data structures generally need to be explained somewhere. Major data structures need comprehensive documentation in general.

Non-obvious dependencies between separate bits of code should be pointed out. Anything which might tempt a code janitor to make an incorrect “cleanup” needs a comment saying why it is done the way it is. And so on.

4.4: Internal Api Changes

The binary interface provided by the kernel to user space cannot be broken except under the most severe circumstances. The kernel's internal programming interfaces, instead, are highly fluid and can be changed when the need arises. If you find yourself having to work around a kernel API, or simply not using a specific functionality because it does not meet your needs, that may be a sign that the API needs to change. As a kernel developer, you are empowered to make such changes.

There are, of course, some catches. API changes can be made, but they need to be well justified. So any patch making an internal API change should be accompanied by a description of what the change is and why it is necessary. This kind of change should also be broken out into a separate patch, rather than buried within a larger patch.

The other catch is that a developer who changes an internal API is generally charged with the task of fixing any code within the kernel tree which is broken by the change. For a widely-used function, this duty can lead to literally hundreds or thousands of changes – many of which are likely to conflict with work being done by other developers. Needless to say, this can be a large job, so it is best to be sure that the justification is solid.

When making an incompatible API change, one should, whenever possible, ensure that code which has not been updated is caught by the compiler. This will help you to be sure that you have found all in-tree uses of that interface. It will also alert developers of out-of-tree code that there is a change that they need to respond to. Supporting out-of-tree code is not something that kernel developers need to be worried about, but we also do not have to make life harder for out-of-tree developers than it needs to be.

5: Posting Patches

Sooner or later, the time comes when your work is ready to be presented to the community for review and, eventually, inclusion into the mainline kernel. Unsurprisingly, the kernel development community has evolved a set of conventions and procedures which are used in the posting of patches; following them will make life much easier for everybody involved. This document will attempt to cover these expectations in reasonable detail; more information can also be found in the files `SubmittingPatches`, `SubmittingDrivers`, and `SubmitChecklist` in the kernel documentation directory.

5.1: When To Post

There is a constant temptation to avoid posting patches before they are completely “ready.” For simple patches, that is not a problem. If the work being done is complex, though, there is a lot to be gained by getting feedback from the community before the work is complete. So you should consider posting in-progress work, or even making a git tree available so that interested developers can catch up with your work at any time.

When posting code which is not yet considered ready for inclusion, it is a good idea to say so in the posting itself. Also mention any major work which remains to be done and any known problems. Fewer people will look at patches which are known to be half-baked, but those who do will come in with the idea that they can help you drive the work in the right direction.

5.2: Before Creating Patches

There are a number of things which should be done before you consider sending patches to the development community. These include:

- Test the code to the extent that you can. Make use of the kernel’s debugging tools, ensure that the kernel will build with all reasonable combinations of configuration options, use cross-compilers to build for different architectures, etc.
- Make sure your code is compliant with the kernel coding style guidelines.
- Does your change have performance implications? If so, you should run benchmarks showing what the impact (or benefit) of your change is; a summary of the results should be included with the patch.
- Be sure that you have the right to post the code. If this work was done for an employer, the employer likely has a right to the work and must be agreeable with its release under the GPL.

As a general rule, putting in some extra thought before posting code almost always pays back the effort in short order.

5.3: Patch Preparation

The preparation of patches for posting can be a surprising amount of work, but, once again, attempting to save time here is not generally advisable even in the short term.

Patches must be prepared against a specific version of the kernel. As a general rule, a patch should be based on the current mainline as found in Linus’s git tree. It may become necessary to make versions against `-mm`, `linux-next`, or a subsystem tree, though, to facilitate wider testing and review. Depending on the area of your patch and what is going on elsewhere, basing a patch against these other trees can require a significant amount of work resolving conflicts and dealing with API changes.

Only the most simple changes should be formatted as a single patch; everything else should be made as a logical series of changes. Splitting up patches is a bit of an art; some developers spend a long time figuring out how to do it in the way that the community expects.

There are a few rules of thumb, however, which can help considerably:

- The patch series you post will almost certainly not be the series of changes found in your working revision control system. Instead, the changes you have made need to be considered in their final form, then split apart in ways which make sense. The developers are interested in discrete, self-contained changes, not the path you took to get to those changes.
- Each logically independent change should be formatted as a separate patch. These changes can be small (“add a field to this structure”) or large (adding a significant new driver, for example), but they should be conceptually small and amenable to a one-line description. Each patch should make a specific change which can be reviewed on its own and verified to do what it says it does.
- As a way of restating the guideline above: do not mix different types of changes in the same patch. If a single patch fixes a critical security bug, rearranges a few structures, and reformats the code, there is a good chance that it will be passed over and the important fix will be lost.
- Each patch should yield a kernel which builds and runs properly; if your patch series is interrupted in the middle, the result should still be a working kernel. Partial application of a patch series is a common scenario when the “git bisect” tool is used to find regressions; if the result is a broken kernel, you will make life harder for developers and users who are engaging in the noble work of tracking down problems.

- Do not overdo it, though. One developer recently posted a set of edits to a single file as 500 separate patches – an act which did not make him the most popular person on the kernel mailing list. A single patch can be reasonably large as long as it still contains a single *logical* change.
- It can be tempting to add a whole new infrastructure with a series of patches, but to leave that infrastructure unused until the final patch in the series enables the whole thing. This temptation should be avoided if possible; if that series adds regressions, bisection will finger the last patch as the one which caused the problem, even though the real bug is elsewhere. Whenever possible, a patch which adds new code should make that code active immediately.

Working to create the perfect patch series can be a frustrating process which takes quite a bit of time and thought after the “real work” has been done. When done properly, though, it is time well spent.

5.4: Patch Formatting

So now you have a perfect series of patches for posting, but the work is not done quite yet. Each patch needs to be formatted into a message which quickly and clearly communicates its purpose to the rest of the world. To that end, each patch will be composed of the following:

- An optional “From” line naming the author of the patch. This line is only necessary if you are passing on somebody else’s patch via email, but it never hurts to add it when in doubt.
- A one-line description of what the patch does. This message should be enough for a reader who sees it with no other context to figure out the scope of the patch; it is the line that will show up in the “short form” changelogs. This message is usually formatted with the relevant subsystem name first, followed by the purpose of the patch. For example:
gpio: fix build on CONFIG_GPIO_SYSFS=n
- A blank line followed by a detailed description of the contents of the patch. This description can be as long as is required; it should say what the patch does and why it should be applied to the kernel.

- One or more tag lines, with, at a minimum, one Signed-off-by: line from the author of the patch. Tags will be described in more detail below. The above three items should, normally, be the text used when committing the change to a revision control system. They are followed by:
- The patch itself, in the unified (“-u”) patch format. Using the “-p” option to diff will associate function names with changes, making the resulting patch easier for others to read.

You should avoid including changes to irrelevant files (those generated by the build process, for example, or editor backup files) in the patch. The file “dontdiff” in the Documentation directory can help in this regard; pass it to diff with the “-X” option.

The tags mentioned above are used to describe how various developers have been associated with the development of this patch. They are described in detail in the SubmittingPatches document; what follows here is a brief summary. Each of these lines has the format: tag: Full Name <email address> optional-other-stuff

The tags in common use are:

- Signed-off-by: this is a developer’s certification that he or she has the right to submit the patch for inclusion into the kernel. It is an agreement to the Developer’s Certificate of Origin, the full text of which can be found in Documentation/SubmittingPatches. Code without a proper signoff cannot be merged into the mainline.
- Acked-by: indicates an agreement by another developer (often a maintainer of the relevant code) that the patch is appropriate for inclusion into the kernel.
- Tested-by: states that the named person has tested the patch and found it to work.
- Reviewed-by: the named developer has reviewed the patch for correctness; see the reviewer’s statement in Documentation/SubmittingPatches for more detail.
- Reported-by: names a user who reported a problem which is fixed by this patch; this tag is used to give credit to the (often underappreciated) people who test our code and let us know when things do not work correctly.
- Cc: the named person received a copy of the patch and had the opportunity to comment on it.

Be careful in the addition of tags to your patches: only Cc: is appropriate for addition without the explicit permission of the person named.

5.5: Sending The Patch

Before you mail your patches, there are a couple of other things you should take care of:

- Are you sure that your mailer will not corrupt the patches? Patches which have had gratuitous white-space changes or line wrapping performed by the mail client will not apply at the other end, and often will not be examined in any detail. If there is any doubt at all, mail the patch to yourself and convince yourself that it shows up intact. Documentation/email-clients.txt has some helpful hints on making specific mail clients work for sending patches.
- Are you sure your patch is free of silly mistakes? You should always run patches through scripts/checkpatch.pl and address the complaints it comes up with. Please bear in mind that checkpatch.pl, while being the embodiment of a fair amount of thought about what kernel patches should look like, is not smarter than you. If fixing a checkpatch.pl complaint would make the code worse, don't do it.

Patches should always be sent as plain text. Please do not send them as attachments; that makes it much harder for reviewers to quote sections of the patch in their replies. Instead, just put the patch directly into your message.

When mailing patches, it is important to send copies to anybody who might be interested in it. Unlike some other projects, the kernel encourages people to err on the side of sending too many copies; don't assume that the relevant people will see your posting on the mailing lists. In particular, copies should go to:

- The maintainer(s) of the affected subsystem(s). As described earlier, the MAINTAINERS file is the first place to look for these people.
- Other developers who have been working in the same area – especially those who might be working there now. Using git to see who else has modified the files you are working on can be helpful.
- If you are responding to a bug report or a feature request, copy the original poster as well.
- Send a copy to the relevant mailing list, or, if nothing else applies, the linux-kernel list.
- If you are fixing a bug, think about whether the fix should go into the next stable update. If so, stable@kernel.org should get a copy of the patch. Also add a “Cc: stable@kernel.org” to the tags within the patch itself; that will cause the stable team to get a notification when your fix goes into the mainline.

When selecting recipients for a patch, it is good to have an idea of who you think will eventually accept the patch and get it merged. While it is possible to send patches directly to Linus Torvalds and have him merge them, things are not normally done that way. Linus is busy, and there are subsystem maintainers who watch over specific parts of the kernel. Usually you will be wanting that maintainer to merge your patches. If there is no obvious maintainer, Andrew Morton is often the patch target of last resort.

Patches need good subject lines. The canonical format for a patch line is something like:

[PATCH nn/mm] subsys: one-line description of the patch

where “nn” is the ordinal number of the patch, “mm” is the total number of patches in the series, and “subsys” is the name of the affected subsystem. Clearly, nn/mm can be omitted for a single, standalone patch.

If you have a significant series of patches, it is customary to send an introductory description as part zero. This convention is not universally followed though; if you use it, remember that information in the introduction does not make it into the kernel changelogs. So please ensure that the patches, themselves, have complete changelog information.

In general, the second and following parts of a multi-part patch should be sent as a reply to the first part so that they all thread together at the receiving end. Tools like git and quilt have commands to mail out a set of patches with the proper threading. If you have a long series, though, and are using git, please provide the –no-chain-reply-to option to avoid creating exceptionally deep nesting.

6: Followthrough

At this point, you have followed the guidelines given so far and, with the addition of your own engineering skills, have posted a perfect series of patches. One of the biggest mistakes that even experienced kernel developers can make is to conclude that their work is now done. In truth, posting patches indicates a transition into the next stage of the process, with, possibly, quite a bit of work yet to be done.

It is a rare patch which is so good at its first posting that there is no room for improvement. The kernel development process recognizes this fact, and, as a result, is heavily oriented toward the improvement of posted code. You, as the author of that code, will be expected to work with the kernel community to ensure that your code is up to the kernel's quality standards. A failure to participate in this process is quite likely to prevent the inclusion of your patches into the mainline.

6.1: Working With Reviewers

A patch of any significance will result in a number of comments from other developers as they review the code. Working with reviewers can be, for many developers, the most intimidating part of the kernel development process. Life can be made much easier, though, if you keep a few things in mind:

- If you have explained your patch well, reviewers will understand its value and why you went to the trouble of writing it. But that value will not keep them from asking a fundamental question: what will it be like to maintain a kernel with this code in it five or ten years later? Many of the changes you may be asked to make – from coding style tweaks to substantial rewrites – come from the understanding that Linux will still be around and under development a decade from now.
- Code review is hard work, and it is a relatively thankless occupation; people remember who wrote kernel code, but there is little lasting fame for those who reviewed it. So reviewers can get grumpy, especially when they see the same mistakes being made over and over again. If you get a review which seems angry, insulting, or outright offensive, resist the impulse to respond in kind. Code review is about the code, not about the people, and code reviewers are not attacking you personally.
- Similarly, code reviewers are not trying to promote their employers' agendas at the expense of your own. Kernel developers often expect to be working on the kernel years from now, but they understand that their employer could change. They truly are, almost without exception, working toward the creation of the best kernel they can; they are not trying to create discomfort for their employers' competitors.

What all of this comes down to is that, when reviewers send you comments, you need to pay attention to the technical observations that they are making. Do not let their form of expression or your own pride keep that from happening. When you get review comments on a patch, take the time to understand what the reviewer is trying to say. If possible, fix the things that the reviewer is asking you to fix. And respond back to the reviewer: thank them, and describe how you will answer their questions.

Note that you do not have to agree with every change suggested by reviewers. If you believe that the reviewer has misunderstood your code, explain what is really going on. If you have a technical objection to a suggested change, describe it and justify your solution to the problem. If your explanations make sense, the reviewer will accept them. Should your explanation not prove persuasive, though, especially if others start to agree with the reviewer, take some time to think things over again. It can be easy to become blinded by your own solution to a problem to the point that you don't realize that something is fundamentally wrong or, perhaps, you're not even solving the right problem.

One fatal mistake is to ignore review comments in the hope that they will go away. They will not go away. If you repost code without having responded to the comments you got the time before, you're likely to find that your patches go nowhere.

Speaking of reposting code: please bear in mind that reviewers are not going to remember all the details of the code you posted the last time around. So it is always a good idea to remind reviewers of previously raised issues and how you dealt with them; the patch changelog is a good place for this kind of information. Reviewers should not have to search through list archives to familiarize themselves with what was said last time; if you help them get a running start, they will be in a better mood when they revisit your code.

What if you've tried to do everything right and things still aren't going anywhere? Most technical disagreements can be resolved through discussion, but there are times when somebody simply has to make a decision. If you honestly believe that this decision is going against you wrongly, you can always try appealing to a higher power. As of this writing, that higher power tends to be Andrew Morton. Andrew has a great deal of respect in the kernel development community; he can often unjam a situation which seems to be hopelessly blocked. Appealing to Andrew should not be done lightly, though, and not before all other alternatives have been explored. And bear in mind, of course, that he may not agree with you either.

6.2: What Happens Next

If a patch is considered to be a good thing to add to the kernel, and once most of the review issues have been resolved, the next step is usually entry into a subsystem maintainer's tree. How that works varies from one subsystem to the next; each maintainer has his or her own way of doing things. In particular, there may be more than one tree – one, perhaps, dedicated to patches planned for the next merge window, and another for longer-term work.

For patches applying to areas for which there is no obvious subsystem tree (memory management patches, for example), the default tree often ends up being -mm. Patches which affect multiple subsystems can also end up going through the -mm tree.

Inclusion into a subsystem tree can bring a higher level of visibility to a patch. Now other developers working with that tree will get the patch by default. Subsystem trees typically feed into -mm and linux-next as well, making their contents visible to the development community as a whole. At this point, there's a good chance that you will get more comments from a new set of reviewers; these comments need to be answered as in the previous round.

What may also happen at this point, depending on the nature of your patch, is that conflicts with work being done by others turn up. In the worst case, heavy patch conflicts can result in some work being put on the back burner so that the remaining patches can be worked into shape and merged. Other times, conflict resolution will involve working with the other developers and, possibly, moving some patches between trees to ensure that everything applies cleanly. This work can be a pain, but count your blessings: before the advent of the linux-next tree, these conflicts often only turned up during the merge window and had to be addressed in a hurry. Now they can be resolved at leisure, before the merge window opens.

Some day, if all goes well, you'll log on and see that your patch has been merged into the mainline kernel. Congratulations! Once the celebration is complete (and you have added yourself to the MAINTAINERS file), though, it is worth remembering an important little fact: the job still is not done. Merging into the mainline brings its own challenges.

To begin with, the visibility of your patch has increased yet again. There may be a new round of comments from developers who had not been aware of the patch before. It may be tempting to ignore them, since there is no longer any question of your code being merged. Resist that temptation, though; you still need to be responsive to developers who have questions or suggestions.

More importantly, though: inclusion into the mainline puts your code into the hands of a much larger group of testers. Even if you have contributed a driver for hardware which is not yet available, you will be surprised by how many people will build your code into their kernels. And, of course, where there are testers, there will be bug reports.

The worst sort of bug reports are regressions. If your patch causes a regression, you'll find an uncomfortable number of eyes upon you; regressions need to be fixed as soon as possible. If you are unwilling or unable to fix the regression (and nobody else does it for you), your patch will almost certainly be removed during the stabilization period. Beyond negating all of the work you have done to get your patch into the mainline, having a patch pulled as the result of a failure to fix a regression could well make it harder for you to get work merged in the future.

After any regressions have been dealt with, there may be other, ordinary bugs to deal with. The stabilization period is your best opportunity to fix these bugs and ensure that your code's debut in a mainline kernel release is as solid as possible. So, please, answer bug reports, and fix the problems if at all possible. That's what the stabilization period is for; you can start creating cool new patches once any problems with the old ones have been taken care of.

And don't forget that there are other milestones which may also create bug reports: the next mainline stable release, when prominent distributors pick up a version of the kernel containing your patch, etc. Continuing to respond to these reports is a matter of basic pride in your work. If that is insufficient motivation, though, it's also worth considering that the development community remembers developers who lose interest in their code after it's merged. The next time you post a patch, they will be evaluating it with the assumption that you will not be around to maintain it afterward.

6.3: Other Things That Can Happen

One day, you may open your mail client and see that somebody has mailed you a patch to your code. That is one of the advantages of having your code out there in the open, after all. If you agree with the patch, you can either forward it on to the subsystem maintainer (be sure to include a proper From: line so that the attribution is correct, and add a signoff of your own), or send an Acked-by: response back and let the original poster send it upward.

If you disagree with the patch, send a polite response explaining why. If possible, tell the author what changes need to be made to make the patch acceptable to you. There is a certain resistance to merging patches which are opposed by the author and maintainer of the code, but it only goes so far. If you are seen as needlessly blocking good work, those patches will eventually flow around you and get into the mainline anyway. In the Linux kernel, nobody has absolute veto power over any code. Except maybe Linus.

On very rare occasion, you may see something completely different: another developer posts a different solution to your problem. At that point, chances are that one of the two patches will not be merged, and “mine was here first” is not considered to be a compelling technical argument. If somebody else’s patch displaces yours and gets into the mainline, there is really only one way to respond: be pleased that your problem got solved and get on with your work. Having one’s work shoved aside in this manner can be hurtful and discouraging, but the community will remember your reaction long after they have forgotten whose patch actually got merged.

7: Advanced Topics

At this point, hopefully, you have a handle on how the development process works. There is still more to learn, however! This section will cover a number of topics which can be helpful for developers wanting to become a regular part of the Linux kernel development process.

7.1: Managing Patches With Git

The use of distributed version control for the kernel began in early 2002, when Linus first started playing with the proprietary BitKeeper application. While BitKeeper was controversial, the approach to software version management it embodied most certainly was not. Distributed version control enabled an immediate acceleration of the kernel development project. In current times, there are several free alternatives to BitKeeper. For better or for worse, the kernel project has settled on git as its tool of choice.

Managing patches with git can make life much easier for the developer, especially as the volume of those patches grows. Git also has its rough edges and poses certain hazards; it is a young and powerful tool which is still being civilized by its developers. This document will not attempt to teach the reader how to use git; that would be sufficient material for a long document in its own right. Instead, the focus here will be on how git fits into the kernel development process in particular. Developers who wish to come up to speed with git will find more information at:

<http://git.or.cz/>

<http://www.kernel.org/pub/software/scm/git/docs/user-manual.html>

and on various tutorials found on the web.

The first order of business is to read the above sites and get a solid understanding of how git works before trying to use it to make patches available to others. A git-using developer should be able to obtain a copy of the mainline repository, explore the revision history, commit changes to the tree, use branches, etc. An understanding of git’s tools for the rewriting of history (such as rebase) is also useful. Git comes with its own terminology and concepts; a new user of git should know about refs, remote branches, the index, fast-forward merges, pushes and pulls, detached heads, etc. It can all be a little intimidating at the outset, but the concepts are not that hard to grasp with a bit of study.

Using git to generate patches for submission by email can be a good exercise while coming up to speed. When you are ready to start putting up git trees for others to look at, you will, of course, need a server that can be pulled from. Setting up such a server with git-daemon is relatively straightforward if you have a system which is accessible to the Internet. Otherwise, free, public hosting sites (Github, for example) are starting to appear on the net. Established developers can get an account on kernel.org, but those are not easy to come by; see <http://kernel.org/faq/> for more information.

The normal git workflow involves the use of a lot of branches. Each line of development can be separated into a separate “topic branch” and maintained independently. Branches in git are cheap, there is no reason to not make free use of them. And, in any case, you should not do your development in any branch which you intend to ask others to pull from. Publicly-available branches should be created with care; merge in patches from development branches when they are in complete form and ready to go – not before.

Git provides some powerful tools which can allow you to rewrite your development history. An inconvenient patch (one which breaks bisection, say, or which has some other sort of obvious bug) can be fixed in place or made to disappear from the history entirely. A patch series can be rewritten as if it had been written on top of today's mainline, even though you have been working on it for months. Changes can be transparently shifted from one branch to another. And so on. Judicious use of git's ability to revise history can help in the creation of clean patch sets with fewer problems.

Excessive use of this capability can lead to other problems, though, beyond a simple obsession for the creation of the perfect project history. Rewriting history will rewrite the changes contained in that history, turning a tested (hopefully) kernel tree into an untested one. But, beyond that, developers cannot easily collaborate if they do not have a shared view of the project history; if you rewrite history which other developers have pulled into their repositories, you will make life much more difficult for those developers. So a simple rule of thumb applies here: history which has been exported to others should generally be seen as immutable thereafter.

So, once you push a set of changes to your publicly-available server, those changes should not be rewritten. Git will attempt to enforce this rule if you try to push changes which do not result in a fast-forward merge (i.e. changes which do not share the same history). It is possible to override this check, and there may be times when it is necessary to rewrite an exported tree. Moving changesets between trees to avoid conflicts in linux-next is one example. But such actions should be rare. This is one of the reasons why development should be done in private branches (which can be rewritten if necessary) and only moved into public branches when it's in a reasonably advanced state.

As the mainline (or other tree upon which a set of changes is based) advances, it is tempting to merge with that tree to stay on the leading edge. For a private branch, rebasing can be an easy way to keep up with another tree, but rebasing is not an option once a tree is exported to the world. Once that happens, a full merge must be done. Merging occasionally makes good sense, but overly frequent merges can clutter the history needlessly.

Suggested technique in this case is to merge infrequently, and generally only at specific release points (such as a mainline – rc release). If you are nervous about specific changes, you can always perform test merges in a private branch. The git “rerere” tool can be useful in such situations; it remembers how merge conflicts were resolved so that you don't have to do the same work twice.

One of the biggest recurring complaints about tools like git is this: the mass movement of patches from one repository to another makes it easy to slip in ill-advised changes which go into the mainline below the review radar. Kernel developers tend to get unhappy when they see that kind of thing happening; putting up a git tree with unreviewed or off-topic patches can affect your ability to get trees pulled in the future. Quoting Linus:

You can send me patches, but for me to pull a git patch from you, I need to know that you know what you're doing, and I need to be able to trust things *without* then having to go and check every individual change by hand.
(<http://lwn.net/Articles/224135/>)

To avoid this kind of situation, ensure that all patches within a given branch stick closely to the associated topic; a “driver fixes” branch should not be making changes to the core memory management code. And, most importantly, do not use a git tree to bypass the review process. Post an occasional summary of the tree to the relevant list, and, when the time is right, request that the tree be included in linux-next.

If and when others start to send patches for inclusion into your tree, don't forget to review them. Also ensure that you maintain the correct authorship information; the git “am” tool does its best in this regard, but you may have to add a “From:” line to the patch if it has been relayed to you via a third party.

When requesting a pull, be sure to give all the relevant information: where your tree is, what branch to pull, and what changes will result from the pull. The git request-pull command can be helpful in this regard; it will format the request as other developers expect, and will also check to be sure that you have remembered to push those changes to the public server.

7.2: Reviewing Patches

Some readers will certainly object to putting this section with “advanced topics” on the grounds that even beginning kernel developers should be reviewing patches. It is certainly true that there is no better way to learn how to program in the kernel environment than by looking at code posted by others. In addition, reviewers are forever in short supply; by looking at code you can make a significant contribution to the process as a whole.

Reviewing code can be an intimidating prospect, especially for a new kernel developer who may well feel nervous about questioning code – in public – which has been posted by those with more experience. Even code written by the most experienced developers can be improved, though. Perhaps the best piece of advice for reviewers (all reviewers) is this: phrase review comments as questions rather than criticisms. Asking “how does the lock get released in this path?” will always work better than stating “the locking here is wrong.”

Different developers will review code from different points of view. Some are mostly concerned with coding style and whether code lines have trailing white space. Others will focus primarily on whether the change implemented by the patch as a whole is a good thing for the kernel or not. Yet others will check for problematic locking, excessive stack usage, possible security issues, duplication of code found elsewhere, adequate documentation, adverse effects on performance, user-space ABI changes, etc. All types of review, if they lead to better code going into the kernel, are welcome and worthwhile.

8: For More Information

There are numerous sources of information on Linux kernel development and related topics. First among those will always be the Documentation directory found in the kernel source distribution. The top-level HOWTO file is an important starting point; SubmittingPatches and SubmittingDrivers are also something which all kernel developers should read. Many internal kernel APIs are documented using the kerneldoc mechanism; “make htmldocs” or “make pdfdocs” can be used to generate those documents in HTML or PDF format (though the version of TeX shipped by some distributions runs into internal limits and fails to process the documents properly).

Various web sites discuss kernel development at all levels of detail. Your author would like to humbly suggest <http://lwn.net/> as a source; information on many specific kernel topics can be found via the LWN kernel index at:

<http://lwn.net/Kernel/Index/>

Beyond that, a valuable resource for kernel developers is:

<http://kernelnewbies.org/>

Information about the linux-next tree gathers at:

<http://linux.f-seidel.de/linux-next/pmwiki/>

And, of course, one should not forget <http://kernel.org/>, the definitive location for kernel release information. There are a number of books on kernel development:

- Linux Device Drivers, 3rd Edition (Jonathan Corbet, Alessandro Rubini, and Greg Kroah-Hartman). Online at <http://lwn.net/Kernel/LDD3/>.
- Linux Kernel Development (Robert Love).
- Understanding the Linux Kernel (Daniel Bovet and Marco Cesati).

All of these books suffer from a common fault, though: they tend to be somewhat obsolete by the time they hit the shelves, and they have been on the shelves for a while now. Still, there is quite a bit of good information to be found there. Documentation for git can be found at:

<http://www.kernel.org/pub/software/scm/git/docs/>

<http://www.kernel.org/pub/software/scm/git/docs/user-manual.html>

9: Conclusion

Congratulations to anybody who has made it through this long-winded document. Hopefully it has provided a helpful understanding of how the Linux kernel is developed and how you can participate in that process.

In the end, it's the participation that matters. Any open source software project is no more than the sum of what its contributors put into it. The Linux kernel has progressed as quickly and as well as it has because it has been helped by an impressively large group of developers, all of whom are working to make it better. The kernel is a premier example of what can be done when thousands of people work together toward a common goal.

The kernel can always benefit from a larger developer base, though. There is always more work to do. But, just as importantly, most other participants in the Linux ecosystem can benefit through contributing to the kernel. Getting code into the mainline is the key to higher code quality, lower maintenance and distribution costs, a higher level of influence over the direction of kernel development, and more. It is a situation where everybody involved wins. Fire up your editor and come join us; you will be more than welcome.