

Kernel: Wayland, NVIDIA and Linux Development (LWN)

By *Roy Schestowitz*

Created 23/05/2019 - 6:43am

Submitted by Roy Schestowitz on Thursday 23rd of May 2019 06:43:58 AM Filed under [Graphics/Benchmarks](#) [1]
[Linux](#) [2]

- [Problems Being Investigated Under Wayland Itches Program, Including Gaming Performance](#)[3]

Last week we wrote about a "Wayland Itches" program being devised by prolific open-source contributor Hans de Goede of Red Hat. The goal of this program is to address itches/paper-cuts/problems in using GNOME Shell atop Wayland. He's received a fair amount of feedback so far and has some early indications to share.

Hans de Goede wrote two blog posts today outlining the early feedback to his Wayland Itches project. Two items he is going to look into initially are middle-click on title/header bar to lower the Window not working for native applications and sudo/pfexec not working on Wayland. For the sudo/pfexec support, Hans is planning to optionally support the ability for GUI apps to connect when running as root. That was rejected upstream before but his plan is for this to be an optional feature for enabling the xauth file for allowing XWayland as root by GNOME-Shell/Mutter.

- [NVIDIA 418.52.07 Linux Driver Wires In Two More Extensions](#) [4]

NVIDIA today released the 418.52.07 Linux driver as an updated build intended for Vulkan developers with it introducing support for two more extensions.

- [BPF: what's good, what's coming, and what's needed](#) [5]

The 2019 Linux Storage, Filesystem, and Memory-Management Summit differed somewhat from its predecessors in that it contained a fourth track dedicated to the BPF virtual machine. LWN was unable to attend most of those sessions, but a couple of BPF-related talks were a part of the broader program. Among those was a plenary talk by Dave Miller, described as "a wholistic view" of why BPF is successful, its current state, and where things are going.

Years ago, Miller began, Alexei Starovoitov showed up at a netfilter conference promoting his ideas for extending BPF. He described how it could be used to efficiently implement various types of switching fabric ? any type, in fact. Miller said that he didn't understand the power of this idea until quite a bit later.

- [The first half of the 5.2 merge window](#) [6]

When he released the 5.1 kernel, Linus Torvalds noted that he had a family event happening in the middle of the 5.2 merge window and that he would be offline for a few days in the middle. He appears to be trying to make up for lost time before it happens: over 8,300 non-merge changesets have found their way into the mainline in the first four days. As always, there is a wide variety of work happening all over the kernel tree.

- [DAX semantics](#) [7]

In the filesystems track at the 2019 Linux Storage, Filesystem, and Memory-Management Summit, Ted Ts'o led a discussion about an inode flag to indicate DAX files, which is meant to be applied to files that should be directly accessed without going through the page cache. XFS has such a flag, but ext4 and other filesystems do not. The semantics of what the flag would mean are not clear to Ts'o (and probably others), so the intent of the discussion was to try to nail those down.

Dan Williams said that the XFS DAX flag is silently ignored if the device is not DAX capable. Otherwise, the file must be accessed with DAX. Ts'o said there are lots of questions about what turning on or off a DAX flag might mean; does it matter whether there are already pages in the page cache, for example. He said that he did not have any strong preference but thought that all filesystems should stick with one interpretation.

While Christoph Hellwig described things as "all broken", Ts'o was hoping that some agreement could be reached among the disparate ideas of what a DAX flag would mean. A few people think there should be no flag and that it should all be determined automatically, but most think the flag is useful. He suggested starting with something "super conservative", such as only being able to set the flag for zero-length files or only empty directories where the files

in it would inherit the flag. Those constraints could be relaxed later if there was a need.

•

[A filesystem for virtualization](#) [8]

A new filesystem aimed at sharing host filesystems with KVM guests, virtio-fs, was the topic of a session led by Miklos Szeredi at the 2019 Linux Storage, Filesystem, and Memory-Management Summit. The existing solution, which is based on the 9P filesystem from Plan 9, has some shortcomings, he said. Virtio-fs is a prototype that uses the Filesystem in Userspace (FUSE) interface.

The existing 9P-based filesystem does not provide local filesystem semantics and is "pretty slow", Szeredi said. The FUSE-based virtio-fs (RFC patches) is performing "much better". One of the ideas behind the new filesystem is to share the page cache between the host and guests, so there would be no data duplication for multiple guests accessing the same files from the host filesystem.

There are still some areas that need work, however. Metadata and the directory entry cache (dcache) cannot be shared, because data structures cannot be shared between the host and guests. There are two ways to handle that. Either there can be a round trip from the guest to the host for each operation to ensure the coherence of the metadata cache and dcache, or the guest can cache that information and somehow revalidate the cache on each operation without going to the host kernel.

•

[Common needs for Samba and NFS](#) [9]

Amir Goldstein led a discussion on things that the two major network filesystems for Linux, Samba and NFS, could cooperate on at the end of day one of the 2019 Linux Storage, Filesystem, and Memory-Management Summit. In particular, are there needs that both filesystems have that the kernel is not currently providing? He had some ideas of areas that might be tackled, but was looking for feedback from the assembled filesystem developers.

He has recently just started looking at the kernel NFS daemon (knfsd) as it is a lesser use case for the customers of his company's NAS device. Most use Samba (i.e. SMB). He would like to see both interoperate better with other operating systems, though.

•

[NFS topics](#) [10]

Trond Myklebust and Bruce Fields led a session on some topics of interest in the NFS world at the 2019 Linux Storage, Filesystem, and Memory-Management Summit. Myklebust discussed the intersection of NFS and containers, as well adding TLS support to NFS. Fields also had

some container changes to discuss, along with a grab bag of other areas that need attention.

Myklebust began with TLS support for the RPC layer that underlies NFS. One of the main issues is how to do the upcall from the RPC layer to a user-space daemon that would handle the TLS handshake. There is kernel support for doing TLS once the handshake is complete; hardware acceleration of TLS was added in the last year based on code from Intel and Mellanox, he said. RPC will use that code, but there is still the question of handling the handshake.

[Graphics/Benchmarks Linux](#)

Source URL: <http://www.tuxmachines.org/node/124166>

Links:

- [1] <http://www.tuxmachines.org/taxonomy/term/148>
- [2] <http://www.tuxmachines.org/taxonomy/term/63>
- [3] https://www.phoronix.com/scan.php?page=news_item&px=Wayland-Itches-Problems
- [4] https://www.phoronix.com/scan.php?page=news_item&px=NVIDIA-Linux-418.52.07
- [5] <https://lwn.net/Articles/787856/>
- [6] <https://lwn.net/Articles/787963/>
- [7] <https://lwn.net/Articles/787973/>
- [8] <https://lwn.net/Articles/788333/>
- [9] <https://lwn.net/Articles/788335/>
- [10] <https://lwn.net/Articles/788292/>