

svn UPDATE

by Steven Kreuzer

Since the topic of this issue is FreeBSD vs. Linux, I thought it would be interesting to take a closer look at some of the changes made to the LinuxKPI subsystem over the past few months. The project was originally announced in May 2016 as a way to no longer have to port changes from the Linux KMS and DRM drivers over to FreeBSD. The idea is to add a set of shims that can act as a com-

patibility layer to allow for these drivers to work with minimal changes, making it easier to follow upstream development and greatly reducing the diff between FreeBSD code and the original code from Linux. In addition, it also speeds up the integration of new changes in FreeBSD, which I am sure is a welcome addition to anyone running FreeBSD on bleeding-edge hardware.

Properly implement `poll_wait()` in the LinuxKPI. This prevents direct use of the `linux_poll_wakeup()` function from unsafe contexts, which can lead to use-after-free issues.

<https://svnweb.freebsd.org/changeset/base/323349>

Resolve IPv6 scope ID issues when using `ip6_find_dev()` in the LinuxKPI.

<https://svnweb.freebsd.org/changeset/base/323351>

Add more sanity checks to `linux_fget()` in the LinuxKPI. This prevents returning pointers to file descriptors that were not created by the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/323347>

Remove unsafe access to the LinuxKPI file structure from `ibcore`. `selwakeup()` is now done by the `wake_up()` family of functions. <https://svnweb.freebsd.org/changeset/base/323350>

Add some miscellaneous definitions to support the DRM drivers.

<https://svnweb.freebsd.org/changeset/base/322795>

Fix for deadlock situation in the LinuxKPI's RCU synchronize API.

<https://svnweb.freebsd.org/changeset/base/322746>

Use integer type to pass around jiffies and/or ticks values in the LinuxKPI because in FreeBSD ticks are 32-bit.

<https://svnweb.freebsd.org/changeset/base/322357>

Implement parts of the `hrtimer` API in the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/320364>

Add `u64_to_user_ptr()` to the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/320337>

Add `ns_to_ktime()` to the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/320336>

Add `noop_lseek()` to the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/320333>

Allow the VM fault handler to be `NULL` in the LinuxKPI when handling a memory map request. When the VM fault handler is `NULL`, a return code of `VM_PAGER_BAD` is returned from the character device's pager populate handler. <https://svnweb.freebsd.org/changeset/base/320189>

svn **UPDATE** continued

- Add `kthread` parking support to the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/320078>
- Add generic `kqueue()` and `kevent()` support to the LinuxKPI character devices. The implementation allows read and write filters to be created and piggybacks on the `poll()` file operation to determine when a filter should trigger. The piggyback mechanism is simply to check for the `EWOULDBLOCK` or `EAGAIN` return code from `read()`, `write()`, or `ioctl()` system calls and then update the `kqueue()` polling state bits. <https://svnweb.freebsd.org/changeset/base/319409>
- Improve `kqueue()` support in the LinuxKPI. Some applications using `kqueue()` do not set non-blocking I/O mode for event-driven read of file descriptors. This results in the LinuxKPI internal `kqueue` read and write event flags having to be updated before the next read **and/or** write system call, otherwise the read **and/or** write system call may block. <https://svnweb.freebsd.org/changeset/base/319501>
- Implement 64-bit atomic operations for the LinuxKPI. <https://svnweb.freebsd.org/changeset/base/294521>

STEVEN KREUZER is a FreeBSD Developer and Unix Systems Administrator with an interest in retro-computing and air-cooled Volkswagens. He lives in Queens, New York, with his wife, daughter, and dog.

Thank you!

The FreeBSD Foundation would like to acknowledge the following companies for their continued support of the Project. Because of generous donations such as these we are able to continue moving the Project forward.



Are you a fan of FreeBSD? Help us give back to the Project and donate today! freebsdfoundation.org/donate/

Please check out the full list of generous community investors at freebsdfoundation.org/donate/sponsors

Uranium



Iridium



Gold



Silver



vmware



STORMSHIELD

