

Published on *Tux Machines* (<http://www.tuxmachines.org>)

[Home](#) > [content](#) > Kernel and Graphics: Another Attack on the GPL, Power Management and Thermal Control Microconference, Intel and AMD

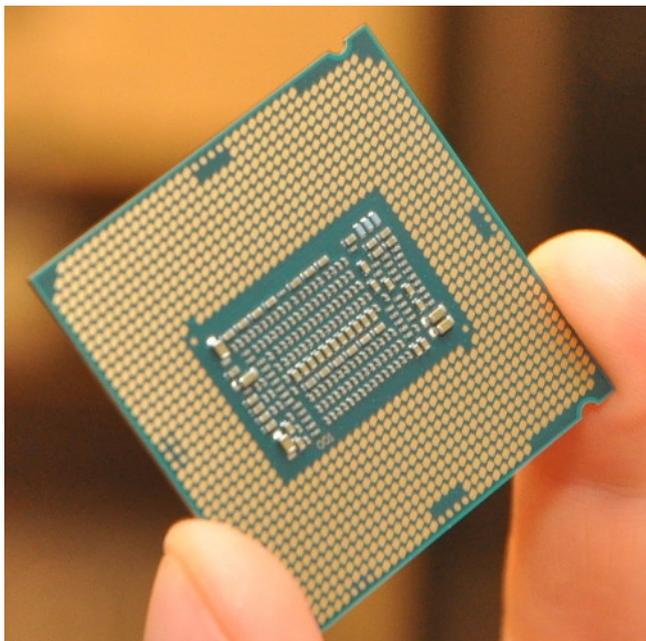
Kernel and Graphics: Another Attack on the GPL, Power Management and Thermal Control Microconference, Intel and AMD

By *Roy Schestowitz*

Created 05/08/2020 - 9:54am

Submitted by Roy Schestowitz on Wednesday 5th of August 2020 09:54:39 AM Filed under [Graphics/Benchmarks](#) [1]

[Linux](#) [2]



- [Wrap it before you tap it? No, say Linux developers: 'GPL condom' for Nvidia driver is laughed out of the kernel](#) [3]

Linux devs have dismissed a proposed patch to the kernel that would only work with a Nvidia driver, motivating a second patch that will prevent disguised use of proprietary code in GPL modules.

The Linux Kernel licensing rules make provision for proprietary third-party modules but state that they must be tagged as such.

This "cannot be used for modules with source code in the kernel tree. Modules tagged that way are tainting the kernel with the 'P' flag when loaded and the kernel module loader refuses to link such modules against symbols which are exported with EXPORT_SYMBOL_GPL()."

Facebook developer Jonathan Lemon put forward an RFC (Request for Comments) on a proposal to implement DMA (Direct Memory Access) zero-copy between a network card and a GPU to enhance network performance, while keeping the protocol processing on the CPU. The use case is for "GPUs used for machine learning, which are located near the NICs, and have a high bandwidth PCI connection between the GPU/NIC," states the RFC.

The code relies on Nvidia's proprietary driver for Linux, noticed by kernel maintainer Greg Kroah-Hartman, who observed: "OK, now you are just trolling us. Nice job, I shouldn't have read the previous patches. Please, go get a lawyer to sign-off on this patch, with their corporate email address on it. That's the only way we could possibly consider something like this."

- [Power Management and Thermal Control Microconference Accepted into 2020 Linux Plumbers Conference](#)[4]

We are pleased to announce that the Power Management and Thermal Control Microconference has been accepted into the 2020 Linux Plumbers Conference!

Power management and thermal control is an important area in the Linux ecosystem to help with the global environment. Optimizing the amount of work that is achieved while having long battery life and keeping the box from overheating is critical in today's world. This meeting will focus on continuing to have Linux be an efficient operating system while still lowering the cost of running a data center.

Last year's meetup at Linux Plumbers resulted in the introduction of thermal pressure support into the CPU scheduler as well as several improvements to the thermal framework, such as a netlink implementation of thermal notification and improvements to CPU cooling. Discussions from last year also helped to improve systems-wide suspend testing tools.

- [Intel Tiger Lake OpenCL Support On Linux Now Considered Production Ready](#)[5]

With all the recent work on Intel's open-source compute stack around the vector back-end and GPU code generation with their ISPC compiler there was another significant milestone achieved that went unnoticed until spotting the change a few days ago.

The open-source Intel Compute Runtime in the past two weeks now has "production" ready OpenCL support for the forthcoming Gen12 Tiger Lake graphics. That's good news with Tiger Lake laptops expected to market soon.

- [RADV ACO Back-End Begins Tackling Navi 2 / GFX10.3 Support](#) [6]

With the "Sienna Cichlid" and "Navy Flounder" open-source driver support as what appear to be the first "Navi 2" GPUs and the first of the "GFX10.3" generation on the graphics engine side there is the initial kernel support with Linux 5.9 and the initial Mesa support for 20.2. That Mesa support has been focused on RadeonSI as the official OpenGL driver as well as Mesa's RADV driver as the Radeon Vulkan driver in-tree but not officially supported by AMD. That RADV support is currently un-tested. Both drivers currently depend upon the "AMDGPU" back-end found in the forthcoming LLVM 11.0 with its initial GFX10.3 support. But now on the RADV driver side there is preliminary GFX10.3 bits landing for the popular "ACO" back-end.

ACO is the back-end worked on by Valve and other stakeholders like open-source graphics driver engineers from Google and Red Hat. But as ACO isn't officially supported by AMD, there hasn't been any patches from them in wiring up the Navi 2 / GFX10.3 support for this AMDGPU LLVM alternative. Rhys Perry as part of Valve's Linux driver efforts though has worked out what should be the initial changes needed for this yet-to-be-released hardware with ACO.

[Graphics/Benchmarks Linux](#)

Source URL: <http://www.tuxmachines.org/node/140654>

Links:

[1] <http://www.tuxmachines.org/taxonomy/term/148>

[2] <http://www.tuxmachines.org/taxonomy/term/63>

[3] https://www.theregister.com/2020/08/04/gpl_condom_nvidia_linux_kernel/

[4] <https://www.linuxplumbersconf.org/blog/2020/power-management-and-thermal-control-microconference-accepted-into-2020-linux-plumbers-conference/>

[5] https://www.phoronix.com/scan.php?page=news_item&px=Tiger-Lake-OpenCL-Ready

[6] https://www.phoronix.com/scan.php?page=news_item&px=RADV-ACO-For-Navi-2