



Contribution ID: 274

Type: **not specified**

SwitchDev offload optimizations

Monday 9 September 2019 17:45 (45 minutes)

Linux has a nice SW bridge implementation which provides most of the classic Ethernet switching features. DSA and SwitchDev frameworks allow us to represent HW switch devices in Linux and potentially offload the SW forwarding to HW.

But the offloading facilities are not perfect, and there seem to be room for further improvements:

- Limiting the flooding of L2-Multicast traffic. IGMP snooping can limit the flooding of L3 traffic, but L2-Multicast traffic are always flooded.
- Today all bridge slave interfaces are put into promiscuous mode to allow learning/flooding. But if the bridge is offloaded with HW capable of doing learning/learning, then this should not be necessary.
- When not put into promiscuous mode, the struct net_device structure has a list of multicast addresses which should be received by the interface. But when VLAN sub-interfaces are created, the VLAN information is lost when addresses are installed in the mc list.
- The assumption in the bridge code is that all multicast frames goes to the CPU. But what would it actually take only to request the needed multicast frames to the CPU?
- Challenges in adding new redundancy and protection protocols to the kernel, and how to offload such protocols to HW.

The intend with the talk is to present some of the issues we are facing in adding DSA/SwitchDev drivers for existing and near-time future HW. I will have few solutions to present, but will give our thoughts on how it may be solved. Hopefully with will result in good discussions and input from the audience.

Background information: I'm working on a SwitchDev driver for a yet to be released HW Ethernet switch. It will be a TSN switch targeting industrial networks, with HW accelerators to implement redundancy protocols. CPU power are very limited, and latency are extremely important, which is why it is important for us to improve the HW offload facilities.

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary author: Mr NIELSEN, Allan

Presenter: Mr NIELSEN, Allan

Session Classification: Networking Summit Track