



Contribution ID: 141

Type: **not specified**

XDP - Challenges and Future Work

Tuesday, 13 November 2018 09:20 (35 minutes)

XDP already offers rich facilities for high performance packet processing, and has seen deployment in several production systems. However, this does not mean that XDP is a finished system; on the contrary, improvements are being added in every release of Linux, and rough edges are constantly being filed down. The purpose of this talk is to discuss some of these possibilities for future improvements, including how to address some of the known limitations of the system. We are especially interested in soliciting feedback and ideas from the community on the best way forward.

The issues we are planning to discuss include, but are not limited to:

- User experience and debugging tools: How do we make it easier for people who are not familiar with the kernel or XDP to get to grips with the system and be productive when writing XDP programs?
- Driver support: How do we get to full support for XDP in all drivers? Is this even a goal we should be striving for?
- Performance: At high packet rates, every micro-optimisation counts. Things like inlining function calls in drivers are important, but also batching to amortise fixed costs such as DMA mapping. What are the known bottlenecks, and how do we address them?
- QoS and rate transitions: How should we do QoS in XDP? In particular, rate transitions (where a faster link feeds into a slower) are currently hard to deal with from XDP, and would benefit from, e.g., Active Queue Management (AQM). Can we adapt some of the AQM and QoS facilities in the regular networking stack to work with XDP? Or should we do something different?
- Accelerating other parts of the stack: Tom Herbert started the discussion on accelerating transport protocols with XDP back in 2016. How do we make progress on this? Or should we be doing something different? Are there other areas where we can extend XDP's processing model to provide useful accelerations?

Presenters: DANGAARD BROUER, Jesper (Red Hat); HØILAND-JØRGENSEN, Toke (Karlstad University)

Session Classification: Networking Track