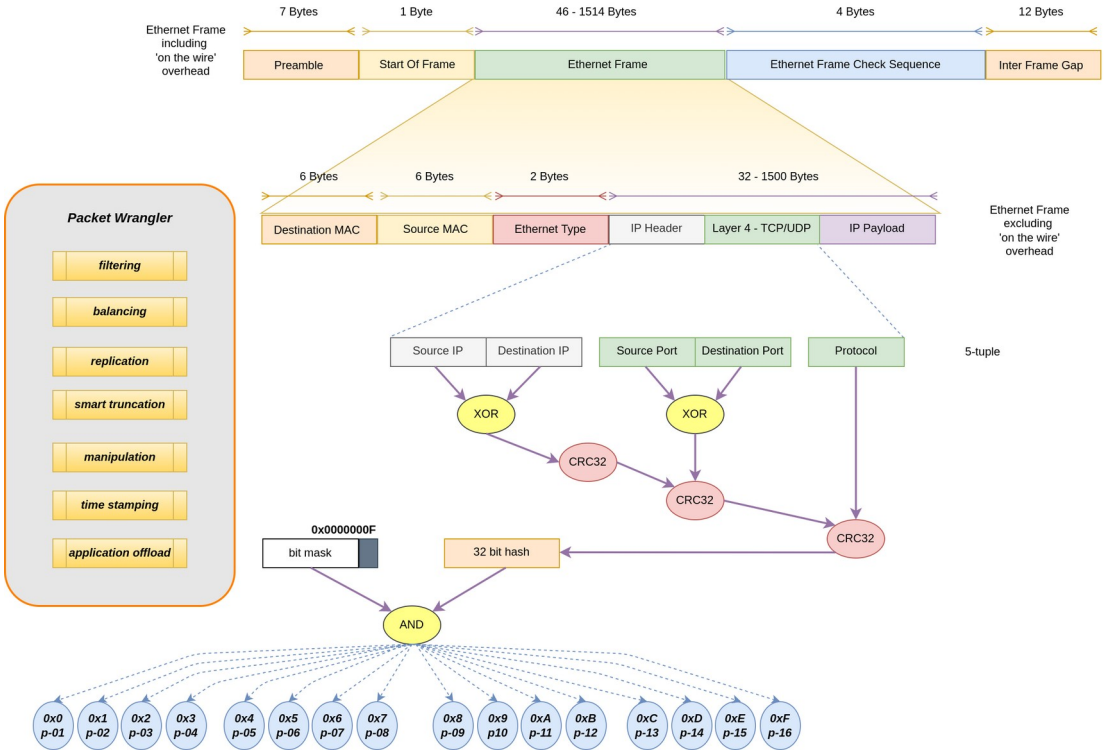

The pne-pwr is a flexible platform that performs a selection of wrangling techniques on Ethernet packets at Layers 2 – 4 on links at speeds from 10G to 400G. The primary platform operates at line rate for non-context techniques such as tuple-hashing load-balancing. The pne-pwr platform can be extended with contextual processors to perform such techniques such as Sticky load-balancing. The capacity of the primary platform units are 32 x 100G and 32 x 400G links. On both platforms the links can be broken down into smaller capacity links - 100, 50, 40, 25, 10G links.

The Horse Wranglers of yesteryear brought multifarious skills and techniques to their work – shoeing, lassoing, branding, breaking, riding, animal psychology and were part veterinarians as well. In Packet Handling applications, several single technique use-cases are achieved by our other pne products such as:

pne-ff	Packet Filtering
pne-bal	Load Balancing
pne-rep	Link Replicator

However some use-cases require a mixed bag of techniques and variations within those techniques – the Packet Wrangler flexibly implements the selection of techniques and custom variations where needed to meet the desired requirements. As the platform is programmable, changes can be made to the feature set to adapt to changing requirements. Flexible link handling allows for ongoing adaption to capacity growth within the limits of a single platform. Platform chaining can allow from growth beyond these dimensions.

The techniques possible for the Packet Wrangler are limited only by the capabilities of the underlying platform and imagination – the following are some techniques that we have implemented:

filtering: line rate filtering of high speed links using IPv4 and IPv6 single and range addresses – forwarding the matching packets to a collecting application or destination. Essentially any layer 2 – 4 entities can be used as filter criteria. Filtering of clear-text tunnels such as GTP can be done using known-port criteria.

balancing: 2-, 3-, and 5-tuple hash balancing of multiple heavy simplex or duplex links over a set of egress ports to lower capacity applications. Most layer 2 – 4 entities can be used as sole criterion or in conjunction with other layer entities and platform entities such as ingress port to form tuple balancing criteria. Fair balancing uses an equal share of the hash space to deliver fair portions of the hashed traffic to the balanced ports. Weighted balancing can be used to deliver ‘weighted’ portions of the hash space to egress ports.

replication: entire links or filtered portions of the packet traffic can be replicated to one or more egress ports. This is very useful when there are several applications that need to look at the same packet traffic.

smart truncation: packet truncation of passive copies of packet traffic is often desirable to reduce load on downstream analyzers and sometimes necessary to allay privacy purposes when performing statistical analysis on public data. Both fixed length and parsed stack truncation is supported.

manipulation: the protocol stack of packets can be manipulated for most elements the Layer 2 – 4, e.g. standard or proprietary headers

time stamping: the pne-pwr platform runs primarily in real-time with only minor asynchronous effect at relatively small egress buffers and very short cross platform latency. When handing off packets to other skewed-time platforms it can be very important to embed a timestamp in the packet to act as a time reference downstream.

application offload: the pne-pwr can off-load downstream applications by performing tasks that would consume resources in the applications. Examples of this are x-tuple hash calculations, origin port tagging and time stamping which can be embedded in additional VLAN tags, unused header options fields, proprietary headers or even the MAC layer.

integration: A REST API makes it a flexible option for any integration into downstream collectors and management systems.

standalone: It can be used as a standalone system, configured by the pne-cli command line interface and delivering its filtered frames to an independent collector port.

footprint: The Tofino: 32 x 100G and Tofino 2: 32 x 400G pne platforms are 1 HU rack components - a very small footprint in terms of space and thermals compared to the port density.

multi-speed links: The Tofino hardware offers 100G ports with breakout DAC options to 4 x 10G and 4 x 25G. The Tofino 2 offers 400G, 200G and 100G ports with breakout DAC options to 4x 25G, 8 x 50G and 2 x 200G. An active DAC offer breakout of 400G to 4 x 100G.

flexible linking: up to 4 * 10 GbE and can be aggregated into a single 40/100 GbE – up to 4 * 25 GbE can be aggregated into a single 100 GbE link using breakout cables.

statistics: can be integrated into the customers application or presented by standalone applications, e.g. InfluxDB, Grafana.

Intel – Barefoot: the pne-platform is based on the Tofino ASIC which is owned and promoted by Intel for Data Plane processing

platforms

- Tofino 1
  - 32 x 100G interfaces – breakouts to 10, 25, 40G
  - 64 x 100G interfaces – breakouts to 10, 25, 40G
- Tofino 2
  - 64 x 100G interfaces – breakouts to 25, 50, 100, 200G
- CLI interface to control configuration and filter tables
- REST API option for direct integration into management systems
- statistics for ports, ingress and egress traffic
- dual power supply, resilient restart

connectivity

- QSFP28 SR, LR – 100 Gbps
- QSFP-DD SR, FR, LR – 400 Gbps
- DAC Breakouts
  - 4 x 25 G to 100 G
  - 2 x 200 G to 400 G
  - 8 x 50 G to 400 G
- DAC
  - 100 G
  - 400 G
- RJ45 - 1 GbE – Management Port