

Implementing Ultra Ethernet Transport with XDP2

Tom Herbert

XDPnet

Overview

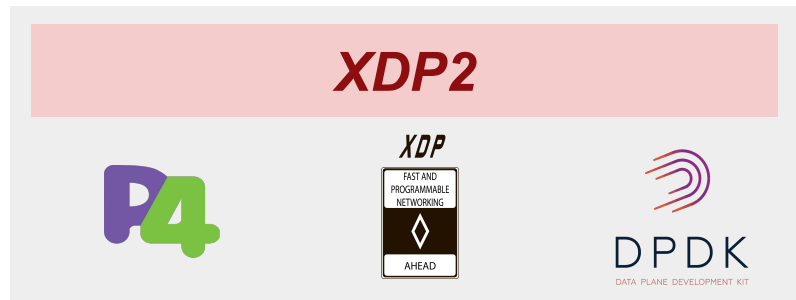
- XDP2 is a new open source programming model for the high performance SW and HW datapath
- The Ultra Ethernet Transport, or UET, is a new transport protocol designed expressly for AI/ML

XDP2 + **Ultra***Ethernet* =
Consortium

Super fast and flexible open source UET implementation

XDP2 is the programming model for the datapath

- User writes their datapath in a language that's convenient for them
- Their code compiles into a variety of target
- It's a type of convergence layer for things like P4, XDP/eBPF, DPDK/VPP



The five pillars of the high performance datapath

*Flexible and high performance
communications infrastructure*

Parsing

Lookup Tables

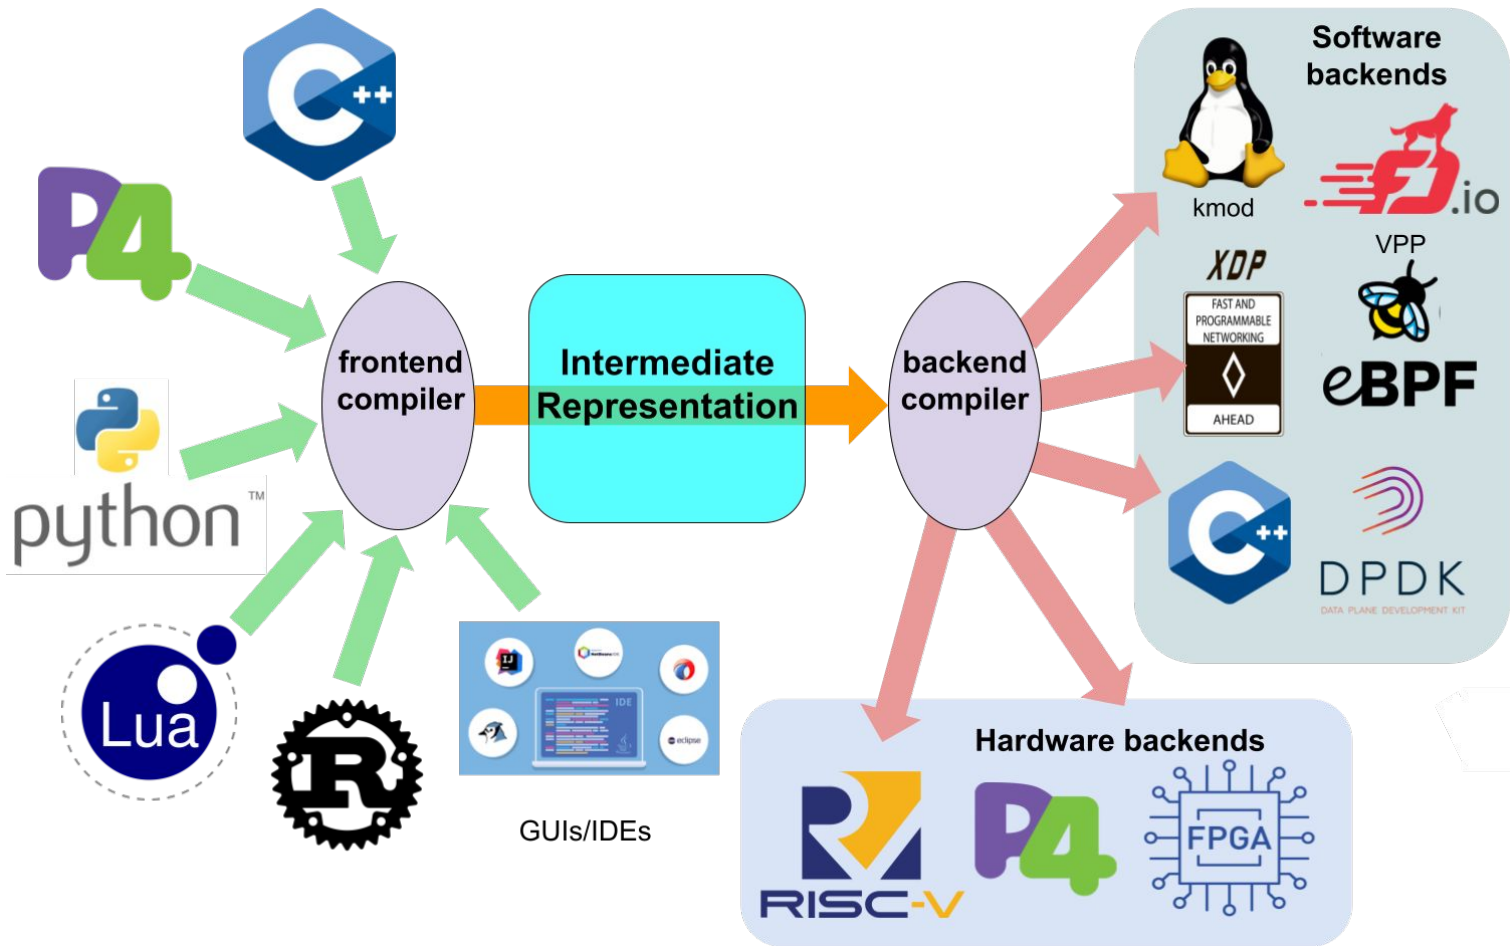
Accelerators

Flow Logic

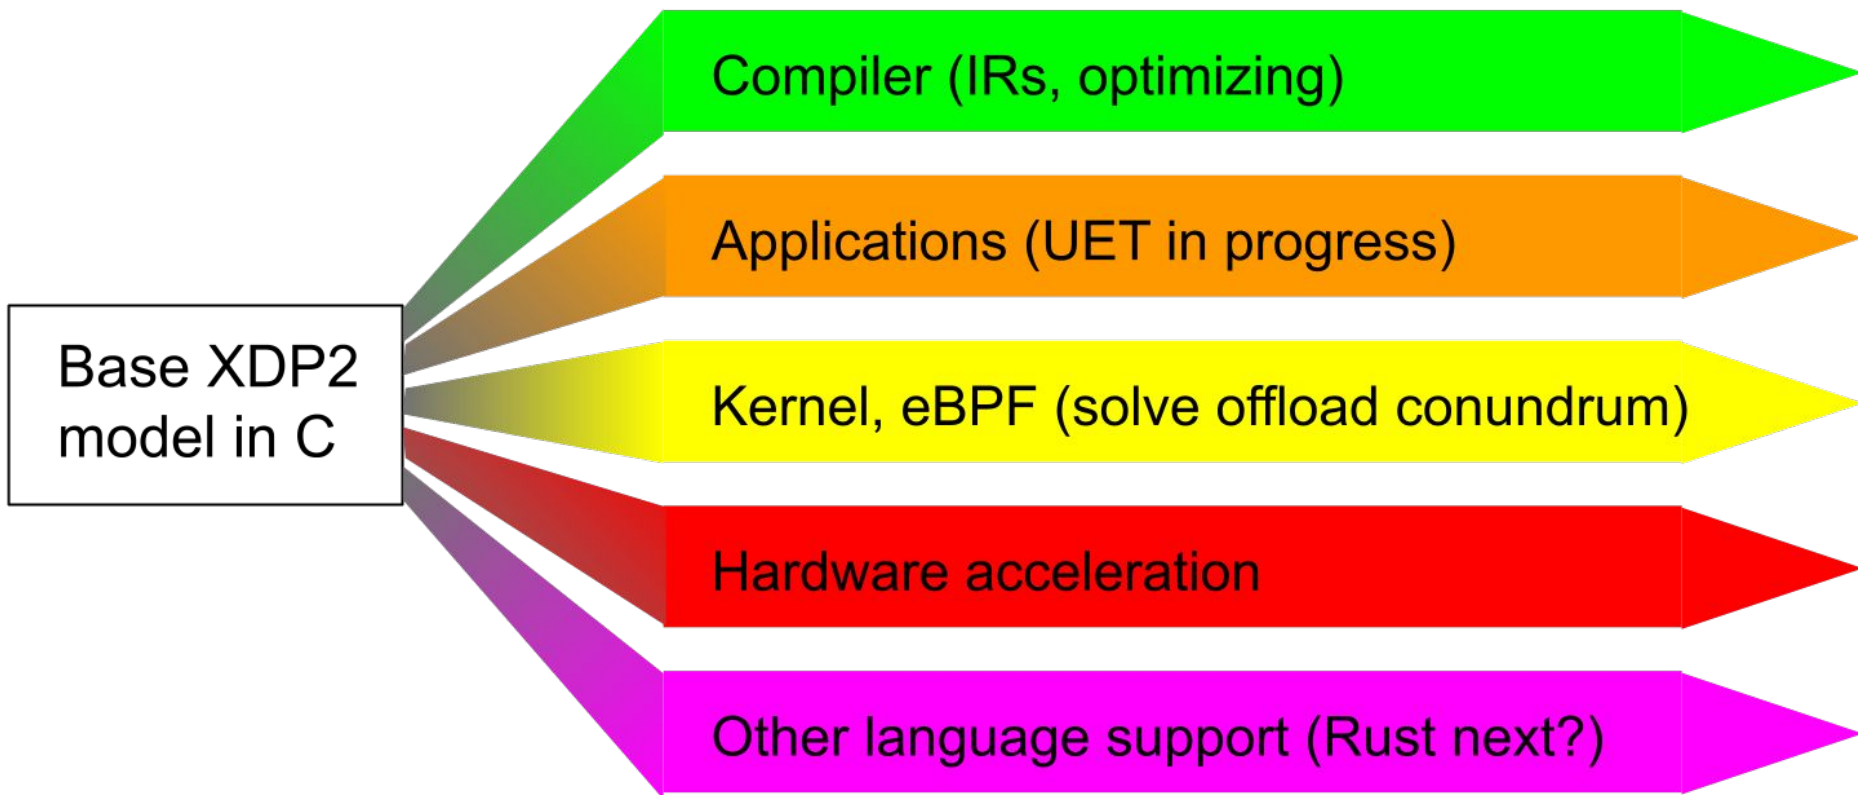
Parallelism

Ubiquitous programming model and
Domain Specific Architecture

Write once, run anywhere,
run well!



XDP2 development vectors



Ultra Ethernet Transport (UET)

The **Ultra Ethernet Consortium** is developing a full-communications stack architecture to meet the growing network demands of AI & HPC at scale

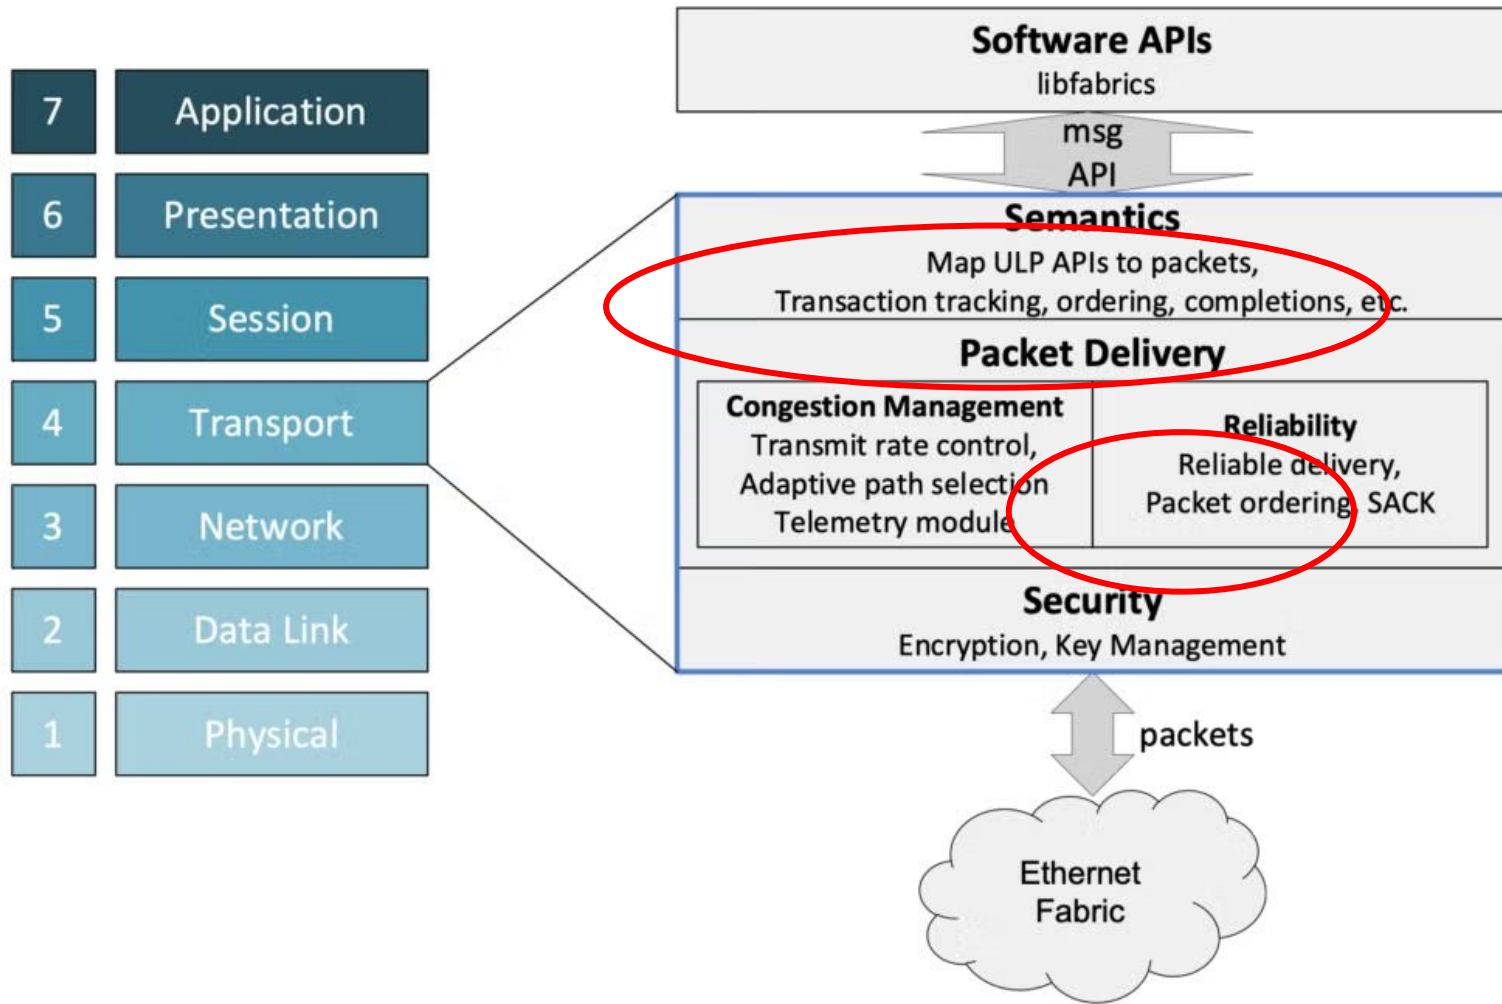
Ultra Ethernet Transport Layer, or **UET**, is the protocol layer that provides reliable packet delivery and messaging semantics for Upper Layer Protocols

UET features



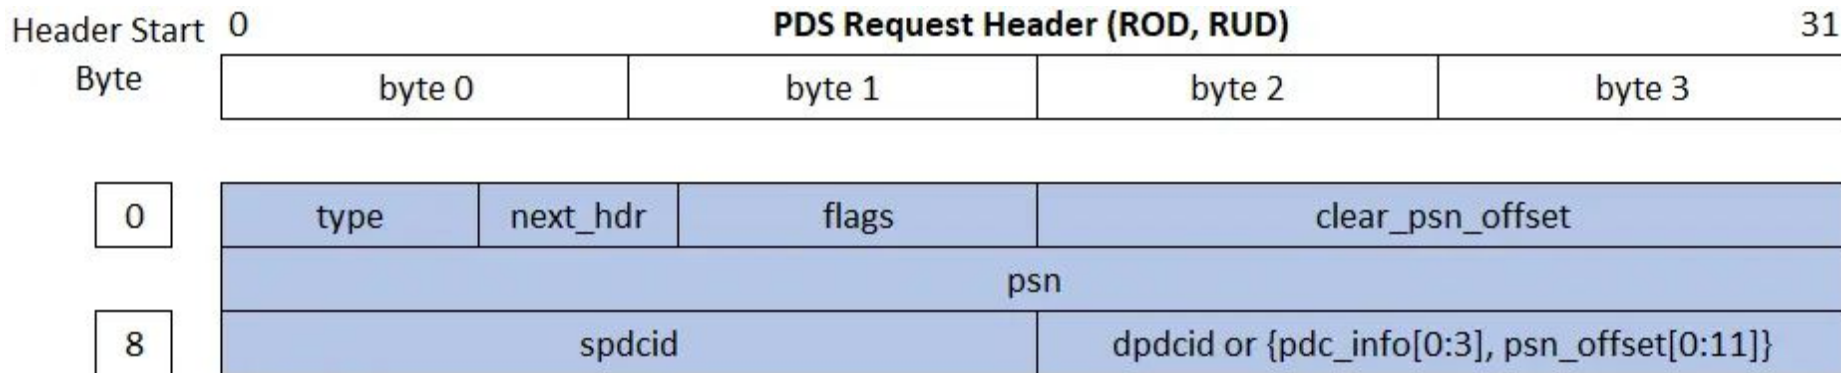
- Low latency, high throughput, simple implementation in HW and SW
- Lightweight connections (PDC)
- Message and initiator/target based comm
- Congestion control with rapid response
- Unordered and ordered packet delivery
- Multi-path and packet spraying
- RDMA and collective operations

UET stack



Packet Delivery Sublayer (PDS)

- Implements reliable delivery of protocol messages
- Request packets, response packets (ACK/NACK), and control packets
- Four *packet delivery* modes

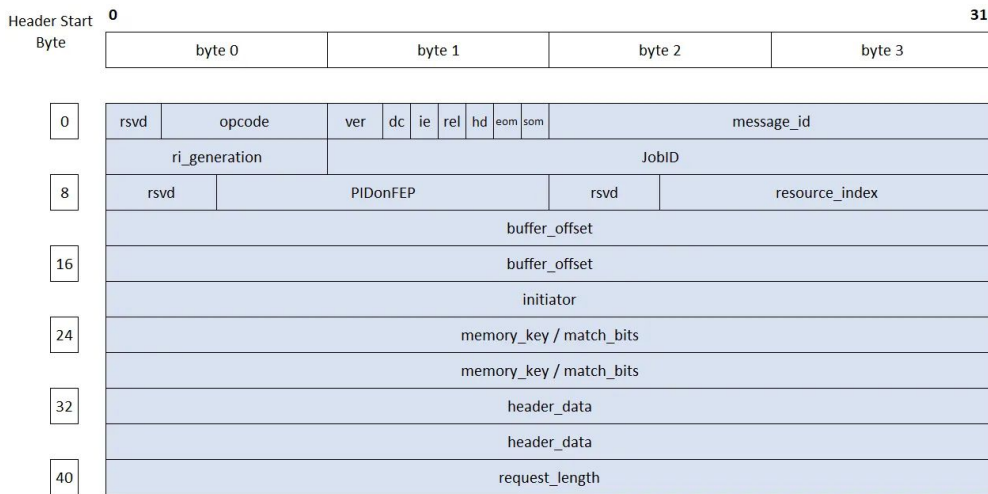


Semantic Sublayer (SES)

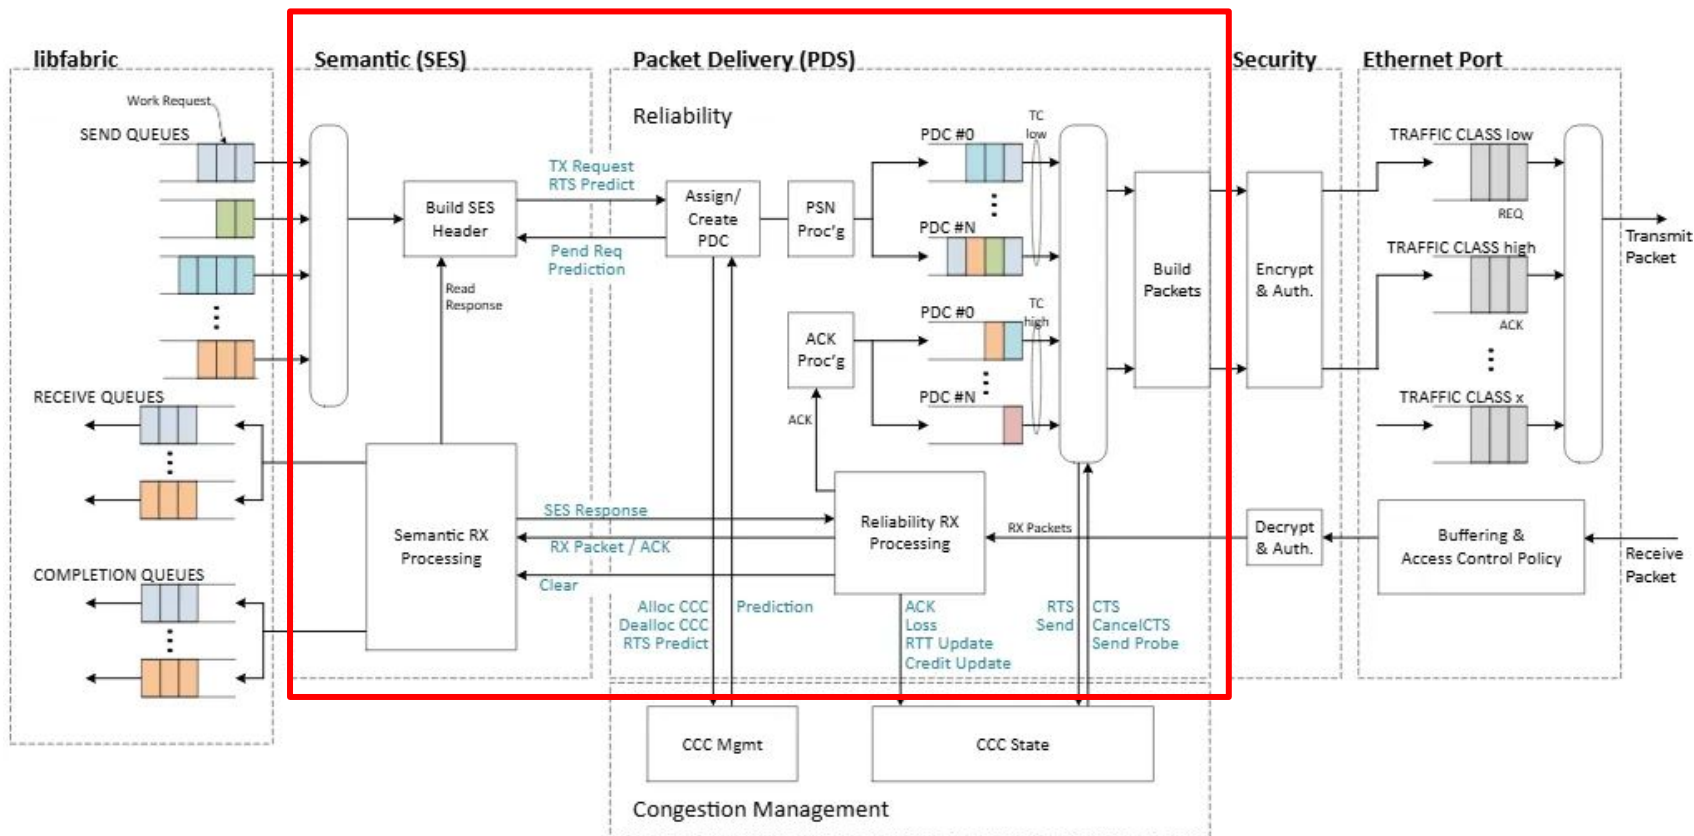
Defines protocol for a initiator-target data transactions including:

- **SEND, WRITE, READ**
- **ATOMICS**
- **RENDEZVOUS** and **DEFERRABLE SEND**
- **RESPONSE** (from target to initiator)

Standard Request Header

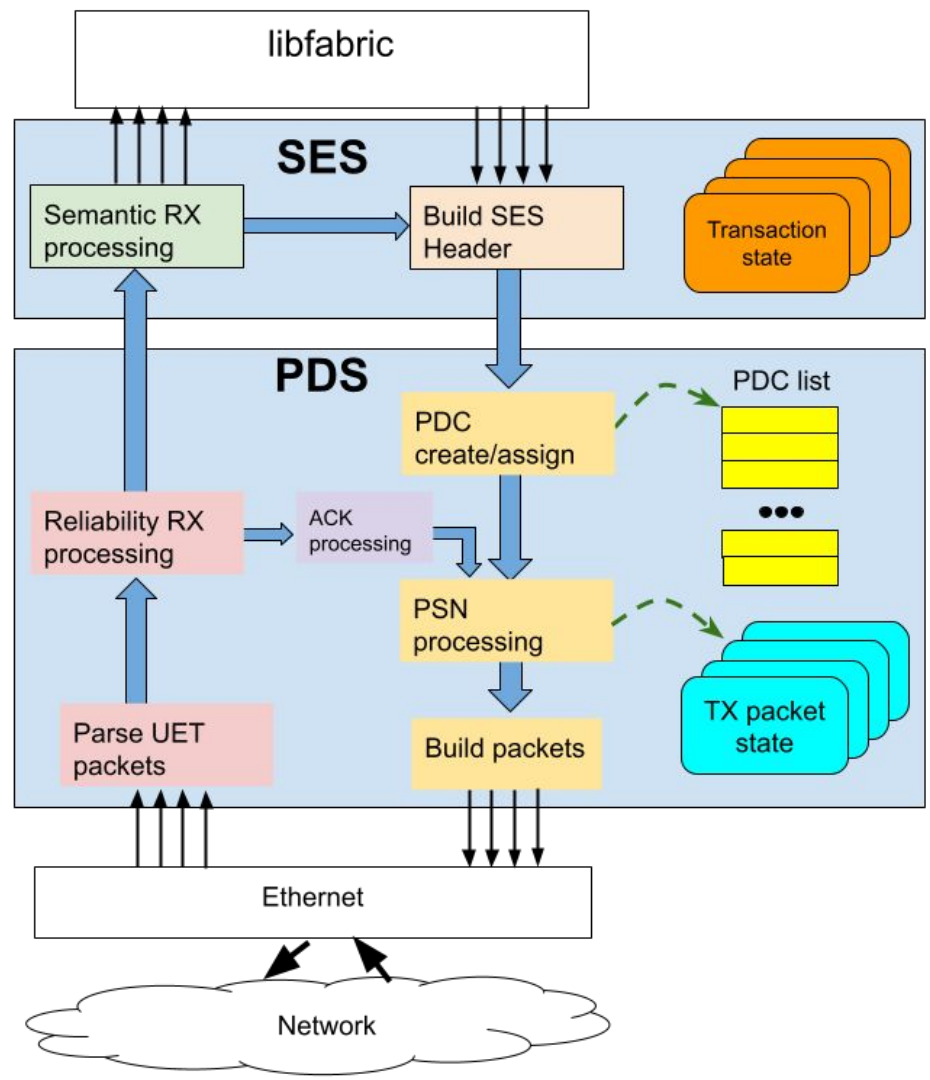


Implementing UET in XDP2



Program components

- Receive path
- Transmit path
- Timeouts
- Interface to ULP (libfabric)
- PDC (connection) managements





XDP2 Services needed

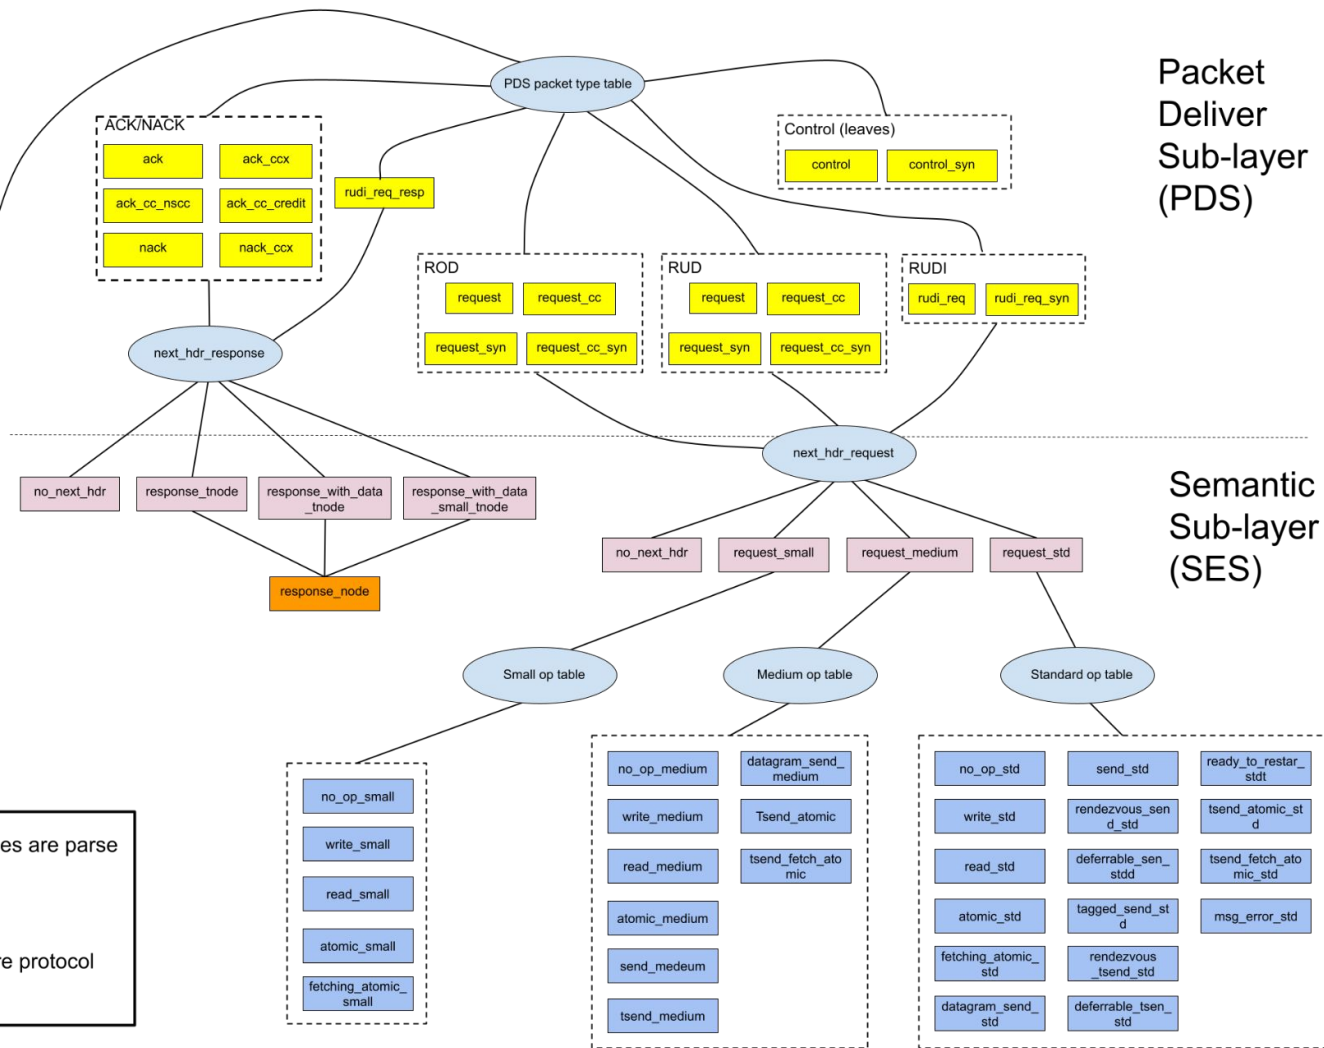
- Protocol parser
- Tables
- Timers
- PVbufs (to hold scatter gather data)
- Bitmaps
- parse_dump to dump UET packets
- Configuration facility
- CLI

```
graph TD; Ethernet[Ethernet] --> Ethernertype([Ethernertype]); Ethernertype --> IPv4[IPv4]; Ethernertype --> IPv6[IPv6]; IPv4 --> IPproto([IP proto]); IPv6 --> IPproto; IPproto --> UDP[UDP]; UDP --> UDPports([UDP ports]); UDPports --> PDSbase[PDS base node];
```

The diagram illustrates the network stack layers. It starts with 'Ethernet' (yellow rectangle) at the top, connected to 'Ethernertype' (light blue oval). 'Ethernertype' branches into 'IPv4' (green rectangle) and 'IPv6' (green rectangle). Both 'IPv4' and 'IPv6' connect to 'IP proto' (light blue oval). 'IP proto' connects to 'UDP' (teal rectangle), which then connects to 'UDP ports' (light blue oval). Finally, 'UDP ports' connects to the 'PDS base node' (magenta rectangle) at the bottom.

 Rectangles are parse nodes

 Ovals are protocol tables



The code on GitHub: **xdp2-dev/xdp2**

UET directories and files

- `src/include/uet`
- `src/lib/uet`
- `src/test/uet`
- `src/tool/packets/uet`
- `documentation/protocol/uet.md`

Thank you!