P4-TC Offload Kernel Driver

Anjali Singhai Jain
Namrata Limaye

Agenda

1. P4-TC in kernel - Introduction
2. P4-TC Offload Driver Design
3. Current Status
4. Control Plane Vendor Driver supporting P4-TC
5. IPU Kernel Control Plane Driver
6. Ethernet Kernel Control Plane Driver
7. Data Plane Design
8. Summary, useful links and contacts

P4-TC in Kernel

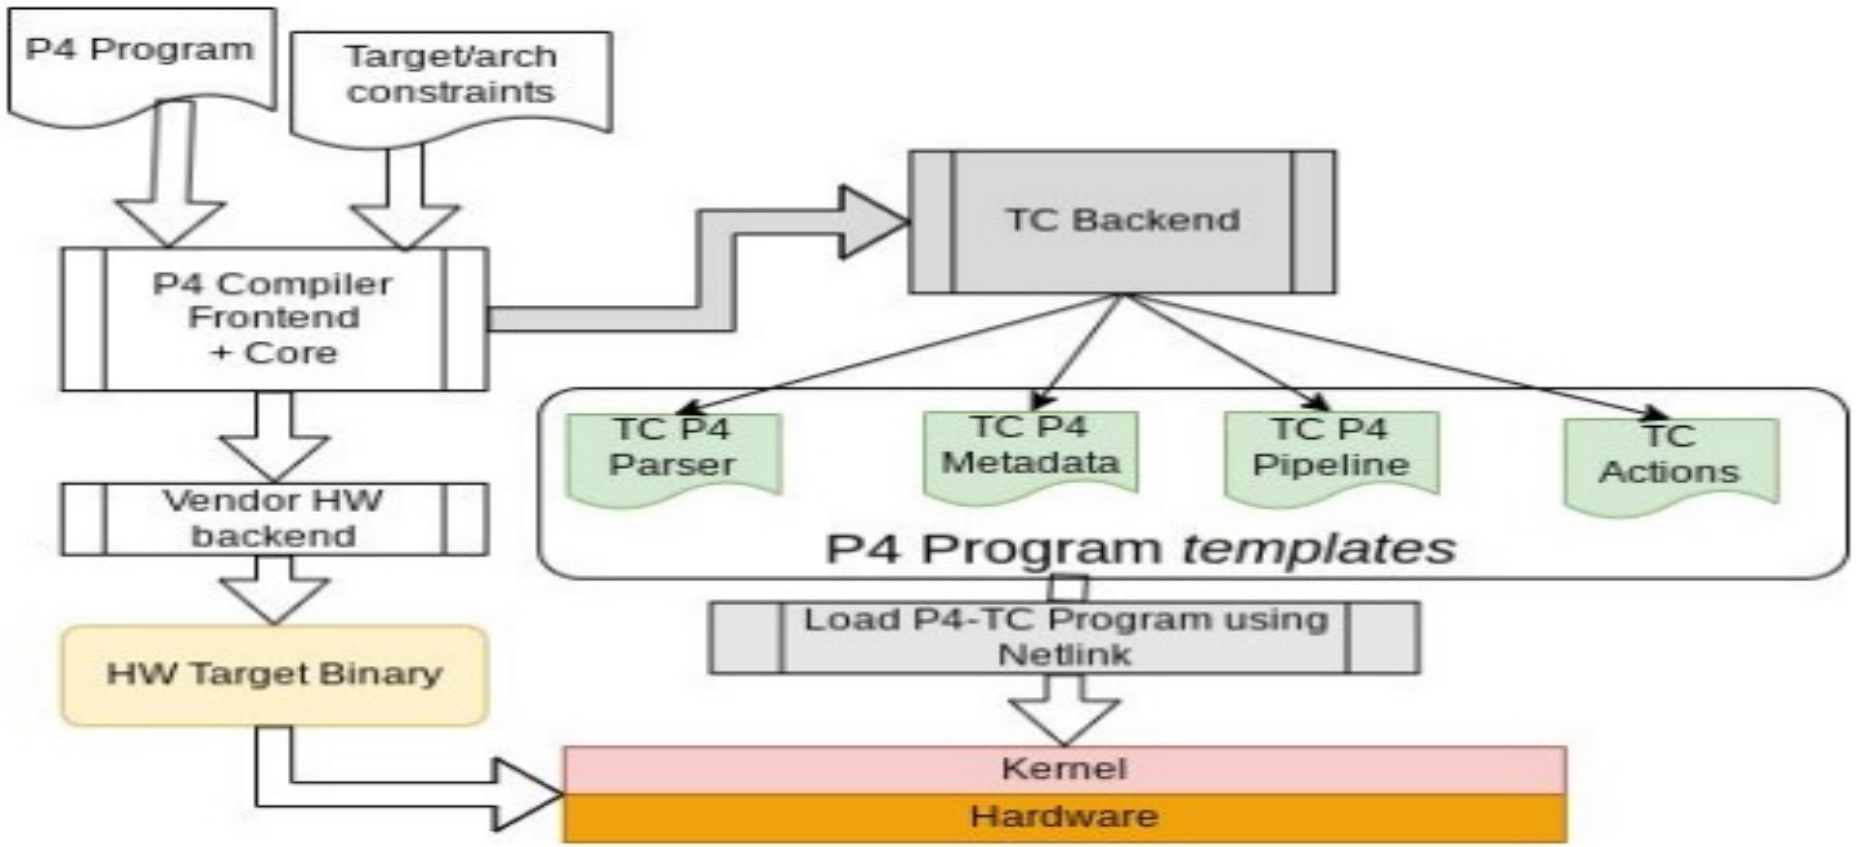

P4-TC Offload Driver Design

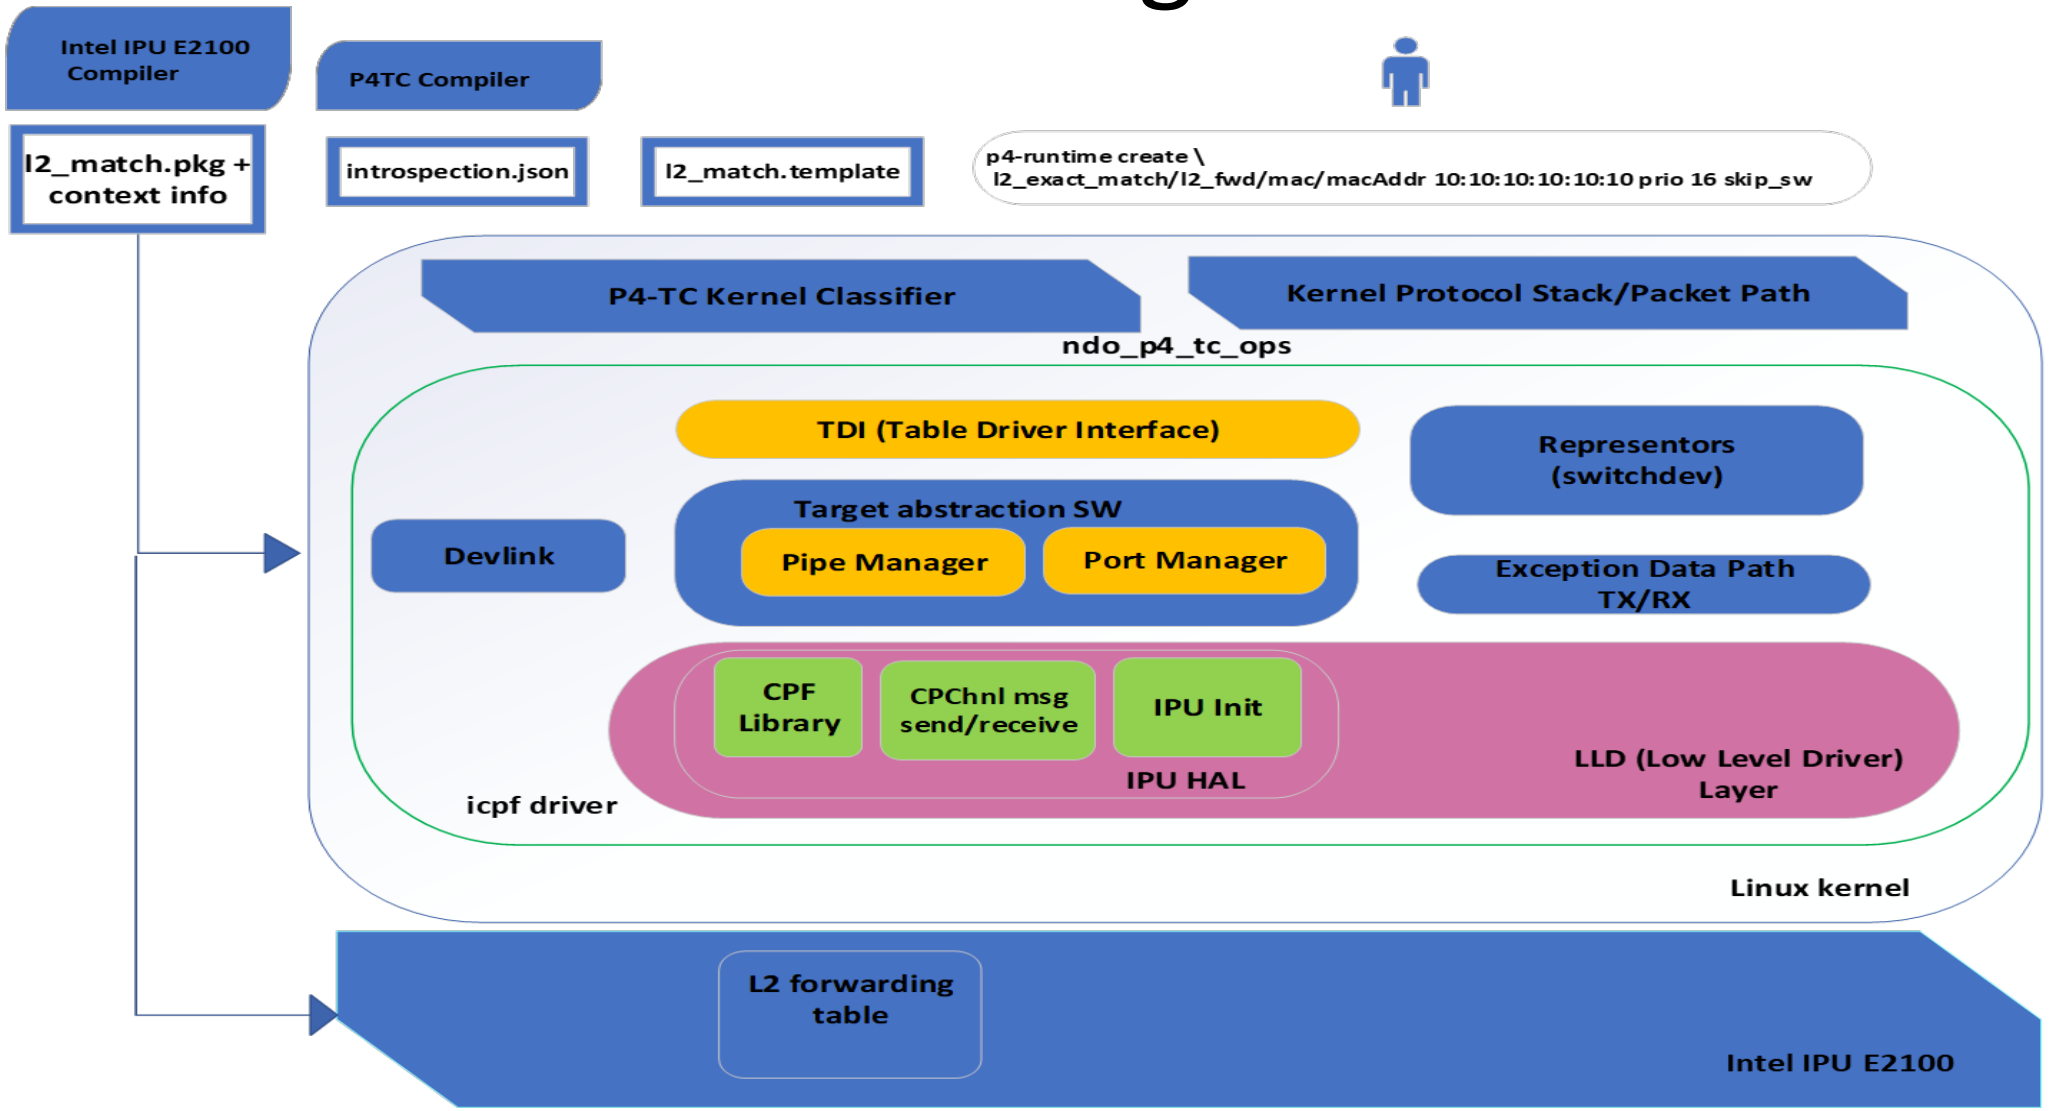

Current Status

- History – User space P4 Control Plane and components open-sourced under ipdk-io and p4lang. TDI used for reference .
- The P4-TC driver interfaces under review in collaboration with multiple vendors
- P4-TC Opensource meetings every 2 weeks
- Driver development in progress
  - Simple L2 offload completed.
  - Work in progress on switchdev/representor support in CPF driver
  - Translation of TDI (p4lang/ipdk-io) for kernel in progress
  - Initialization of P4 dataplane design under review
- Opensource the driver by December 2023
- Upstreaming in Linux kernel right after
- Demo in P4 Conference by Neha Singh

Control Plane Vendor driver supporting p4tc

- Main ingredients:
  - Kernel P4-TC path to software data plane
  - Switchdev and representors for exception packets
  - Initialization component for the active P4 pipeline in target HW
  - Runtime interfaces for Dataplane updates using P4-TC and ndo callbacks
- Single control plane vendor driver with no change works for any future protocol or Data Plane enhancements done via P4-TC
- Not the IDPF Common driver that Intel is upstreaming as of now.
- Although for Dataplane/Exception path, IDPF Spec is used as well.

https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=idpf

Kernel Control Plane: IPU

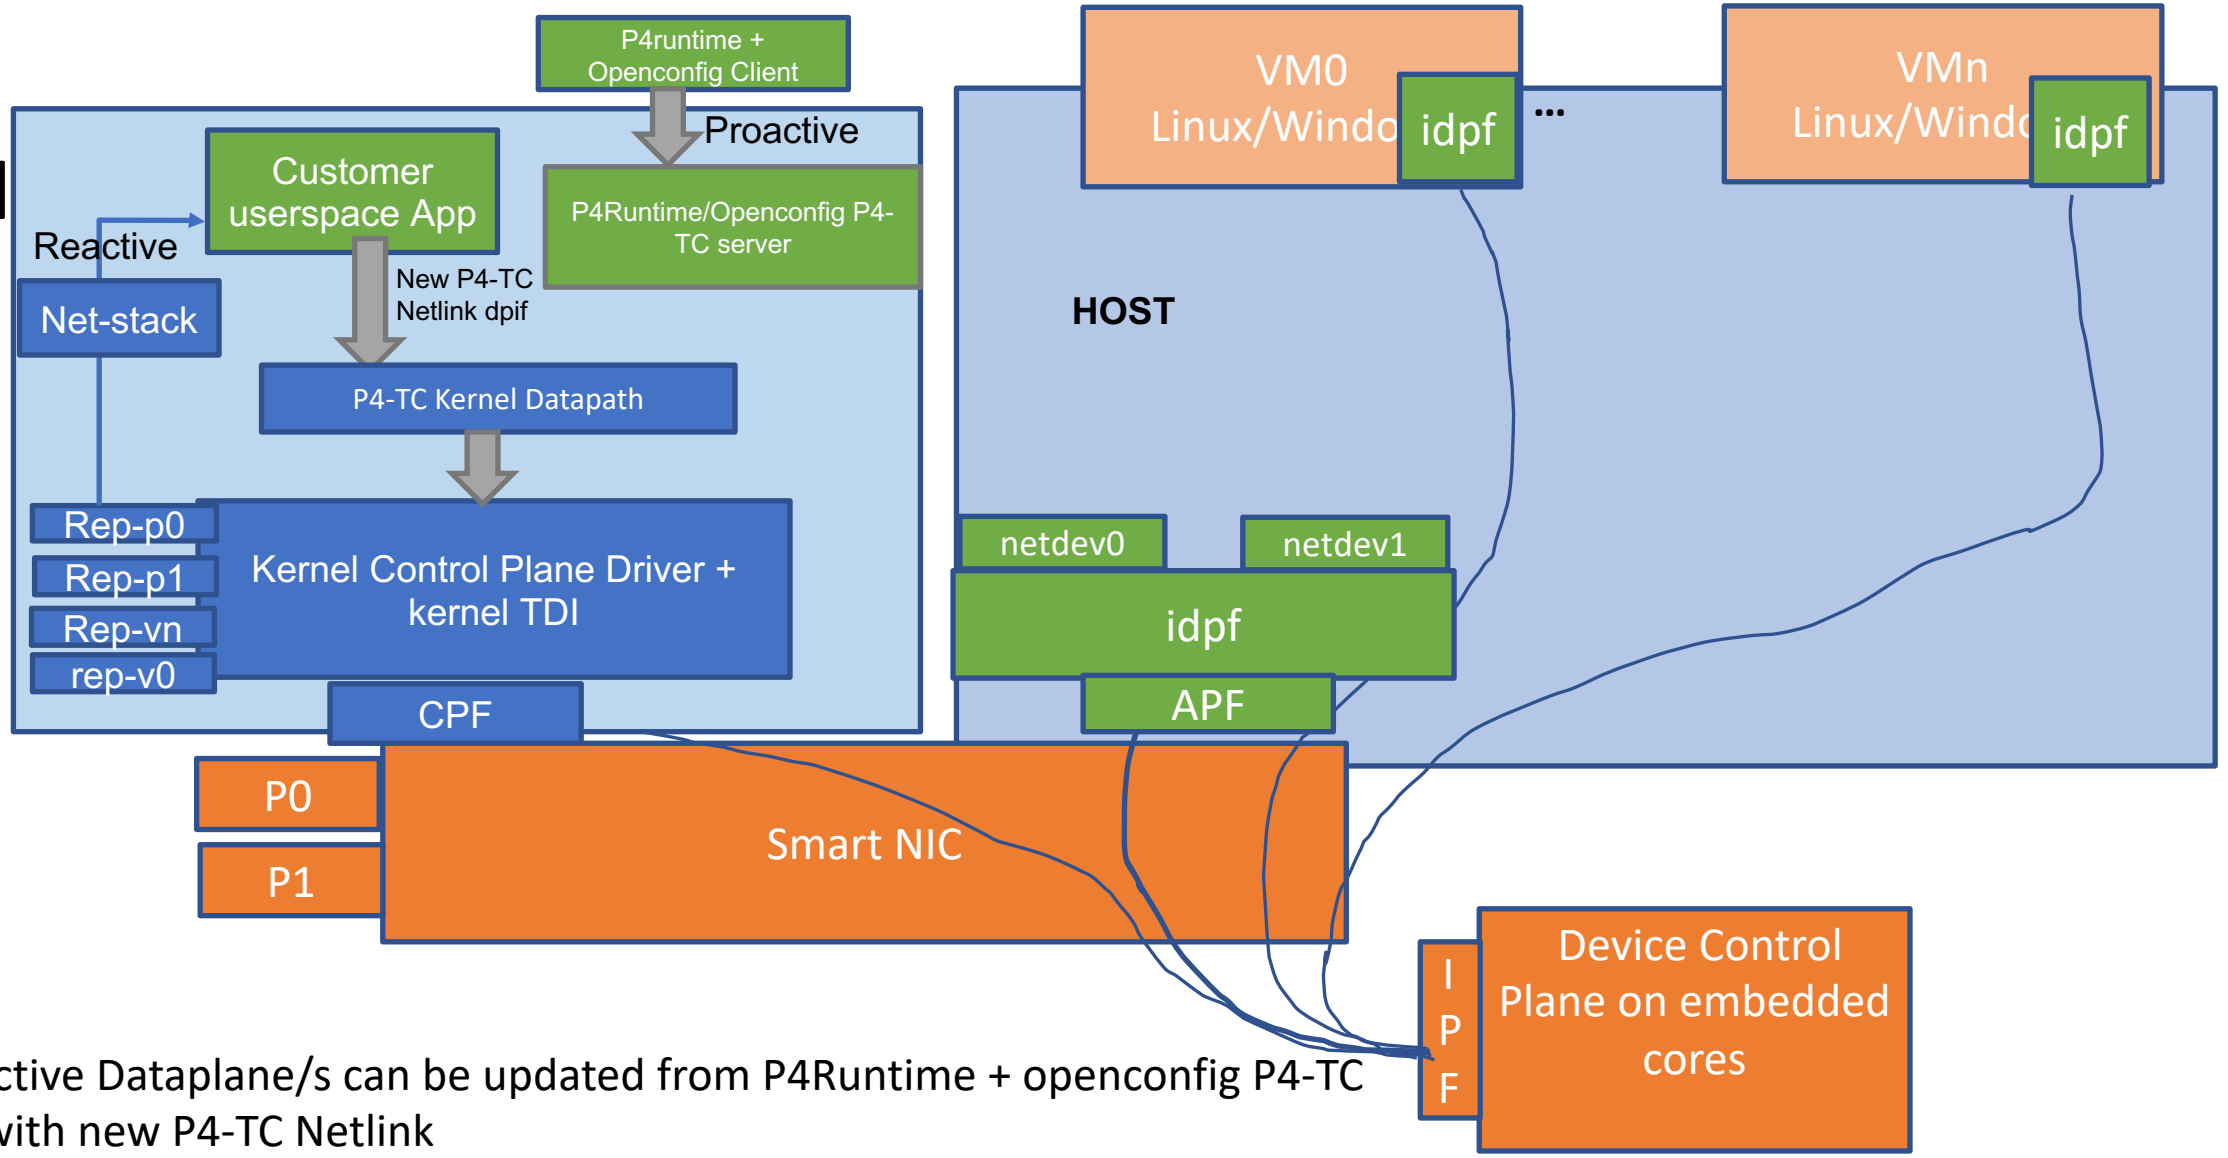

1. P4 Proactive Dataplane/s can be updated from P4Runtime + openconfig P4-TC Server with new P4-TC Netlink
2. Kernel configs and tables get populated via netlink listeners in the KCP Driver
3. OVS Offload via tc-p4 hooks (new dpif in OVS) – Preferred
4. Switchdev/s on CPF driver for exception path
5. IDPF: Base Ethernet offload features, no switchdev etc

Color legend

HW blocks

Pre-existing software

New software

Kernel Control Plane : Smart NIC

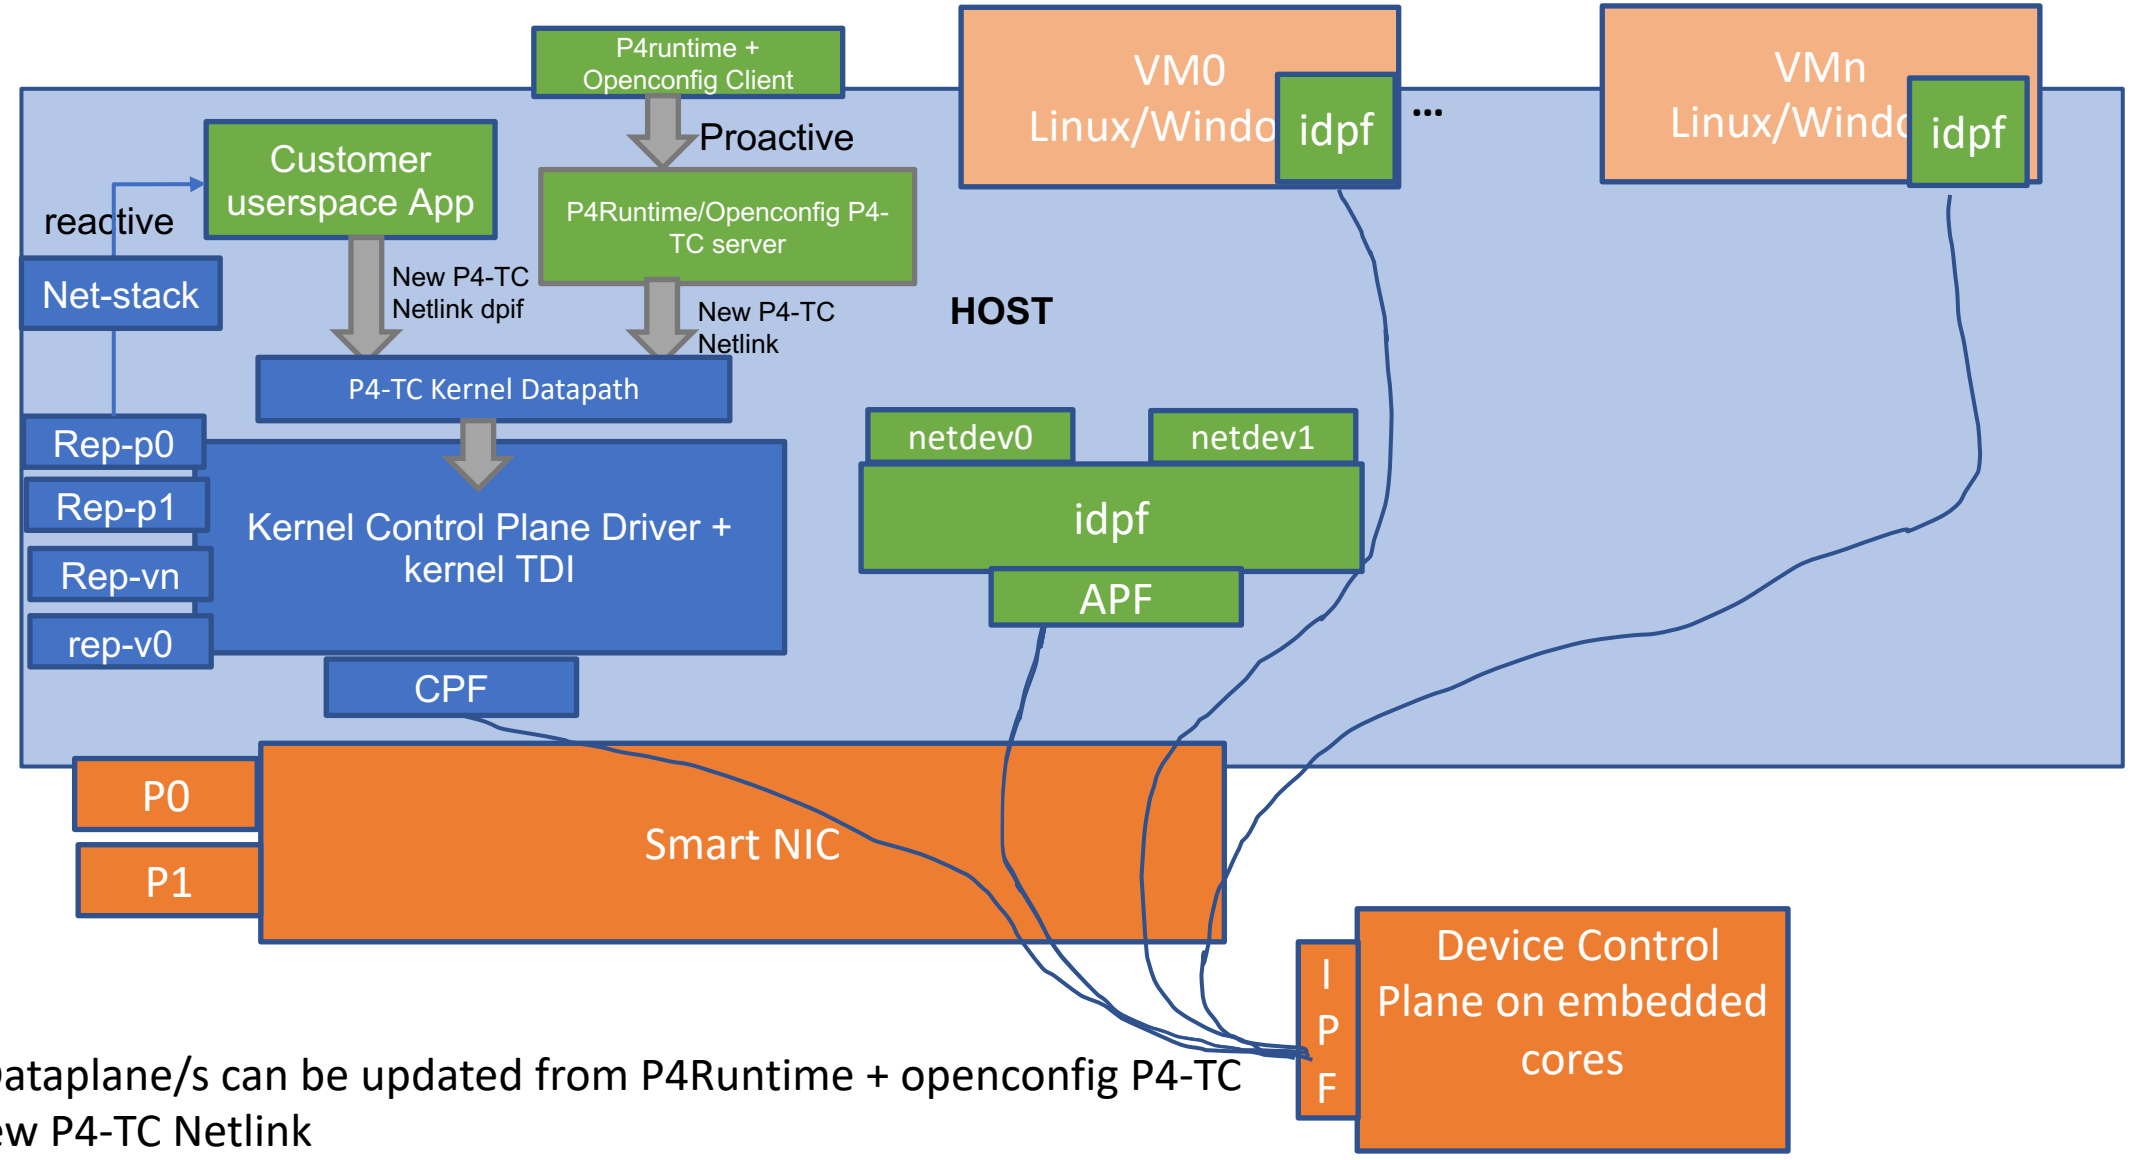

1. P4 Proactive Dataplane/s can be updated from P4Runtime + openconfig P4-TC Server with new P4-TC Netlink
2. Kernel configs and tables get populated via netlink listeners in the KCP Driver
3. OVS Offload via tc-p4 hooks (new dpif in OVS) – Preferred
4. Switchdev/s on CPF driver for exception path
5. IDPF: Base Ethernet offload features, no switchdev etc

Color legend

- HW blocks
- Pre-existing software
- New software

Data Plane Design

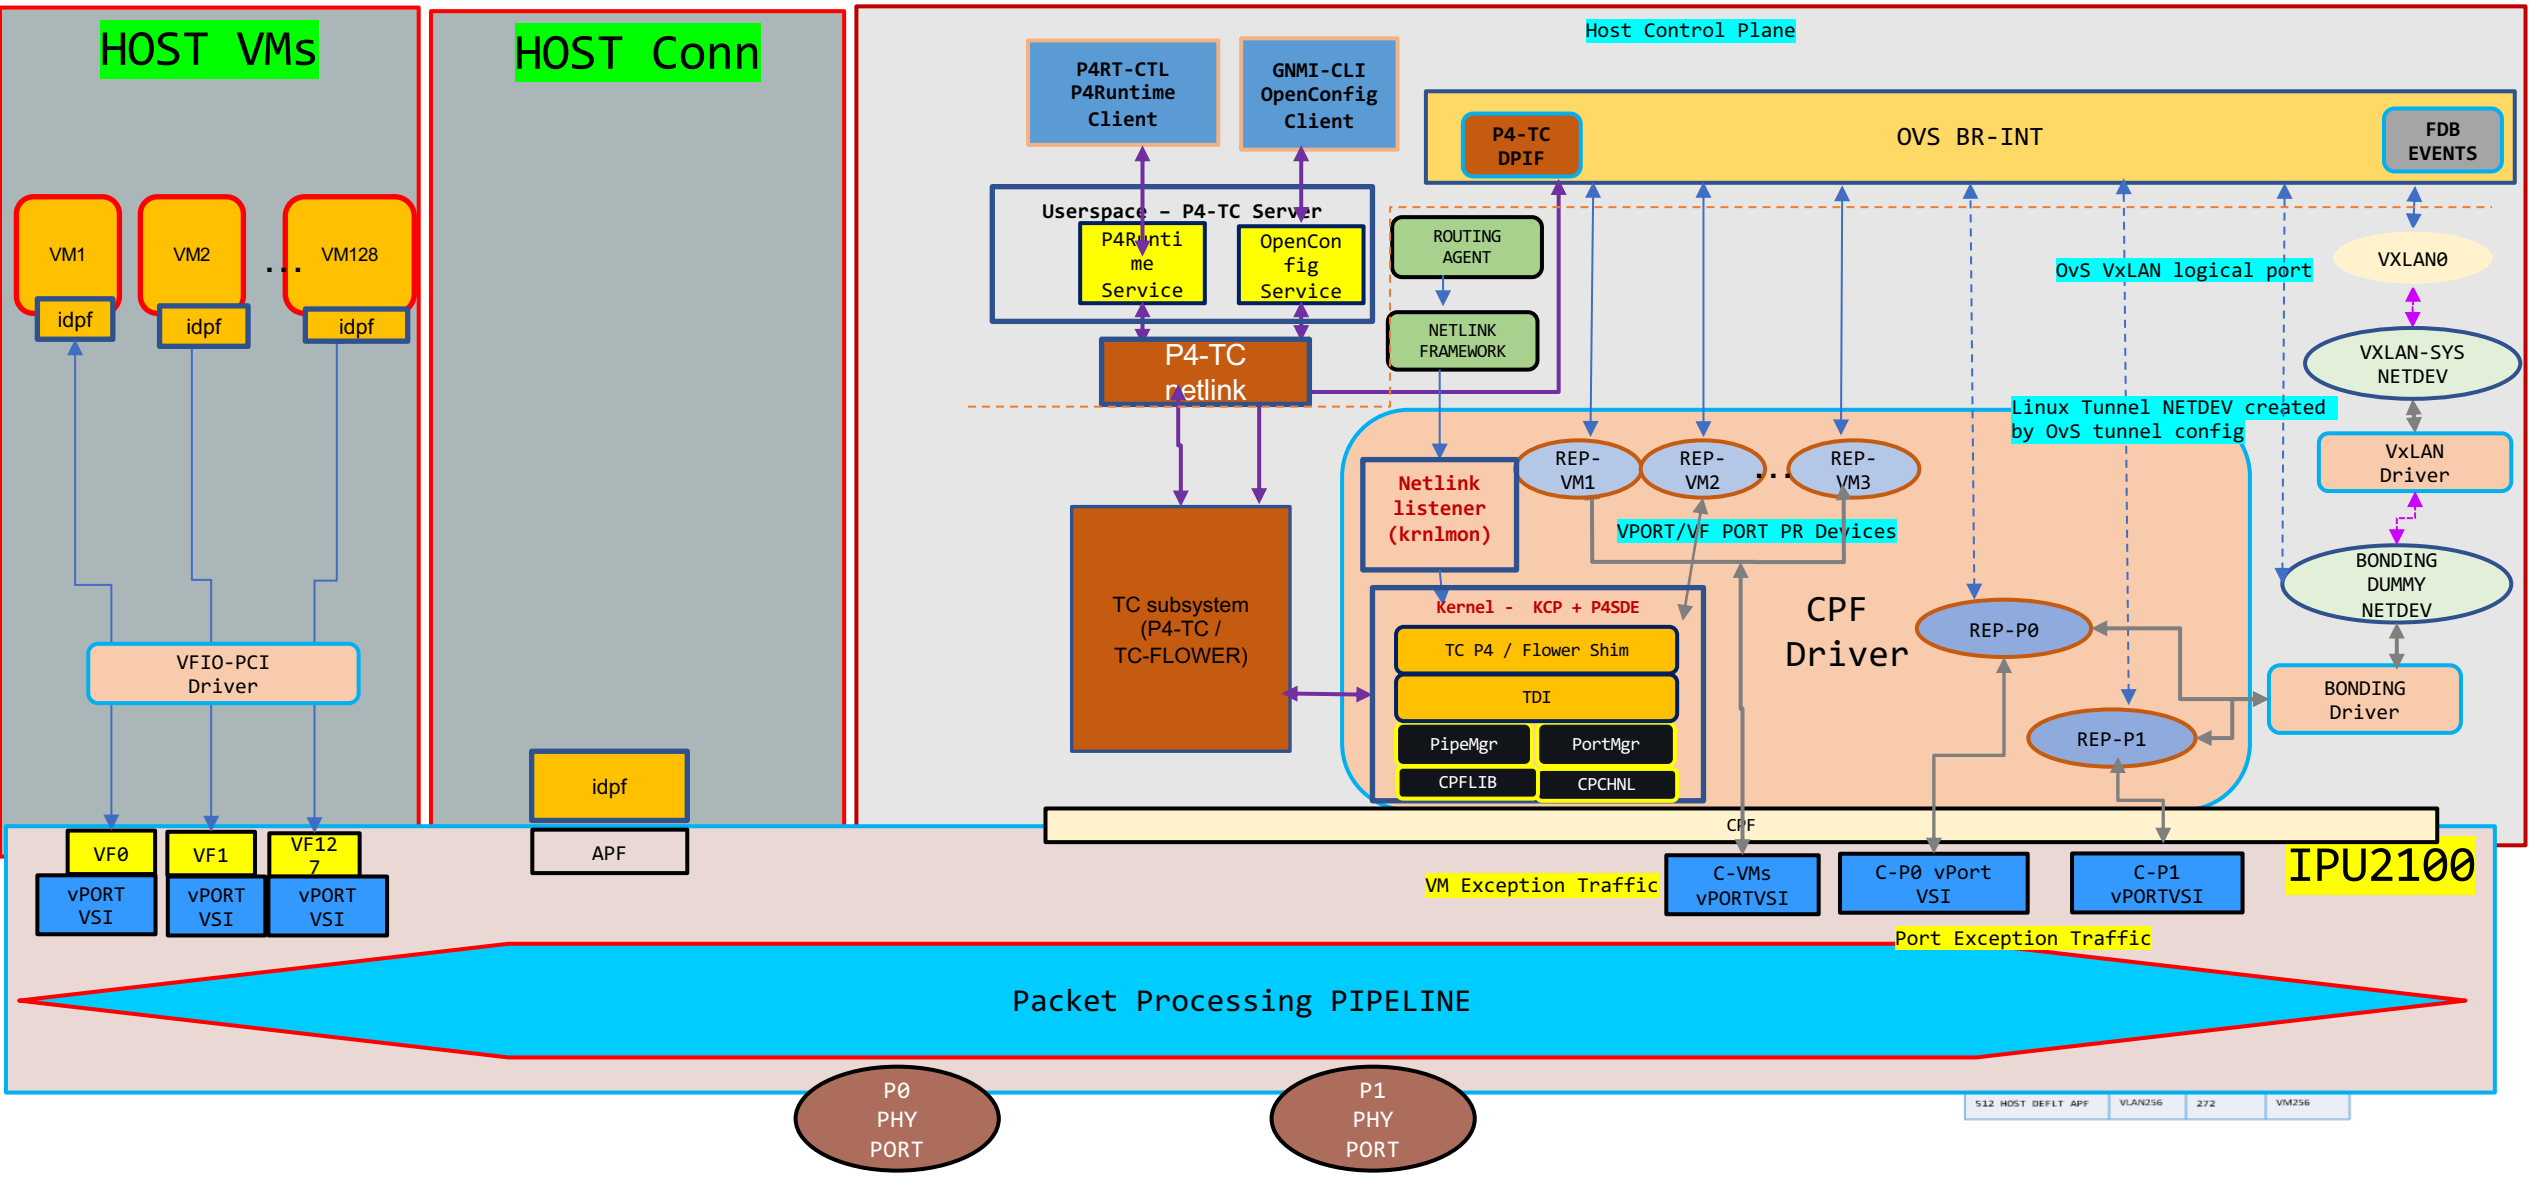

Summary

- Useful Links
  - https://github.com/p4lang/tdi - TDI
  - https://github.com/ipdk-io - IPDK (Demo by Nupur Uttarwar for networking recipe and Nupur Jain on K8s recipe)
  - https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=idpf- - IDPF
  - https://www.p4tc.dev/
- Contacts
  - Anjali.Singhai@intel.com
  - Namrata.Limaye@intel.com
  - James.choi@intel.com

Thank You!

Anjali Singhai
Namrata Limaye