

envoycon

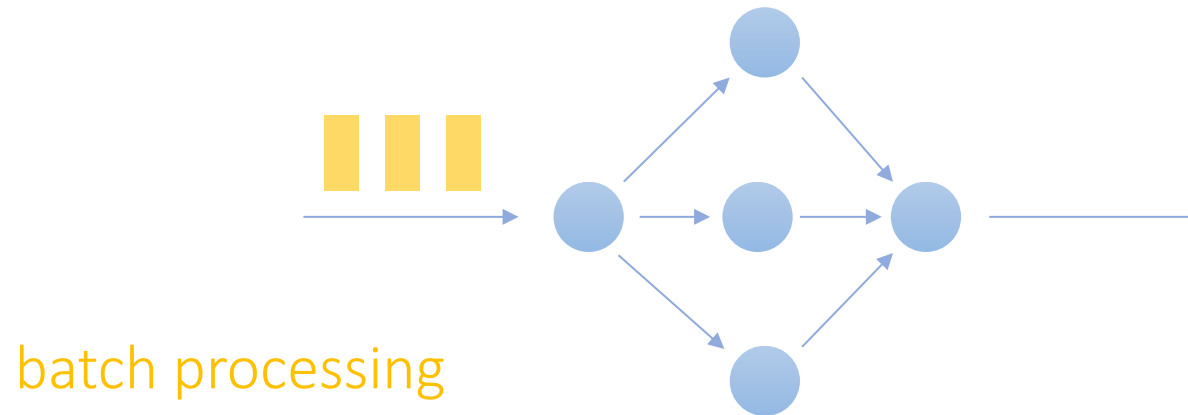
October 15, 2020  
Virtual Experience

# Using VPP as Envoy's Network Stack

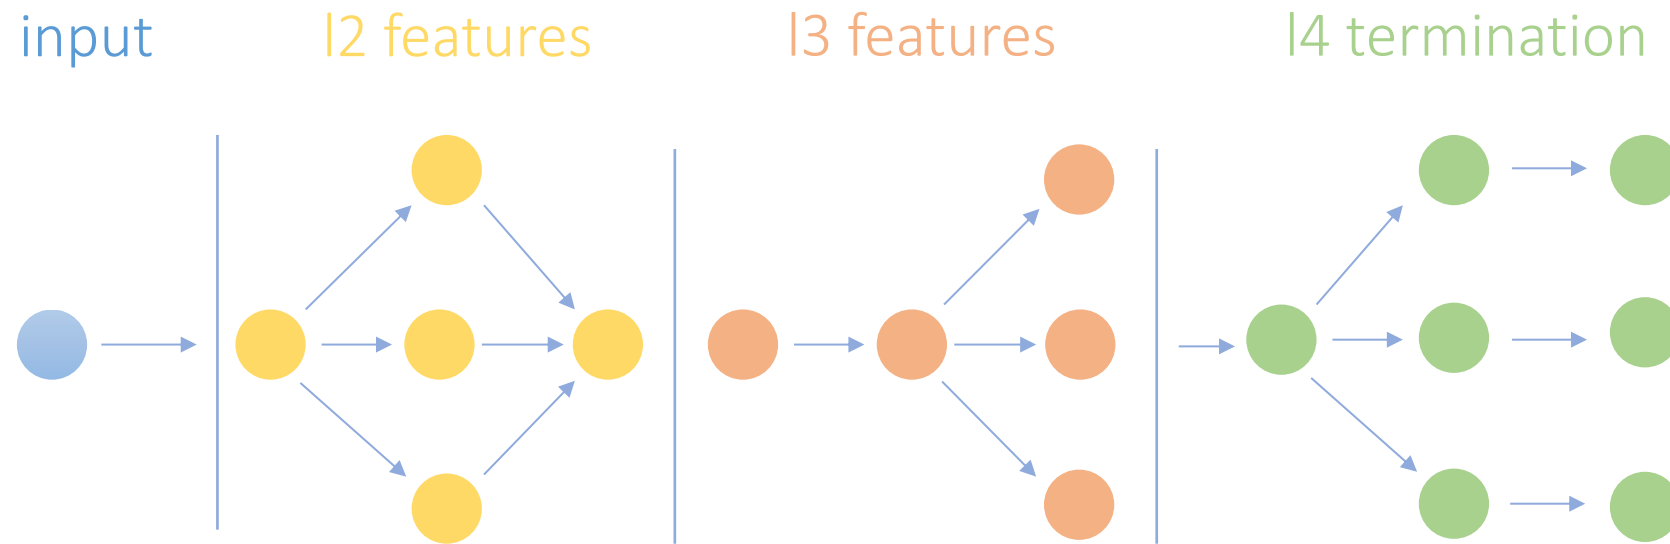
*Florin Coras, Cisco*

# Vector Packet Processing (VPP)

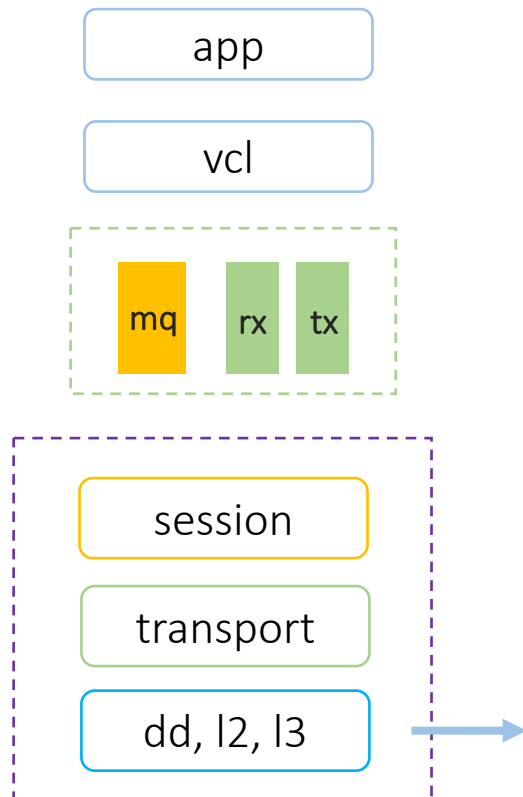
forwarding model as a graph



# Vector Packet Processing (VPP)

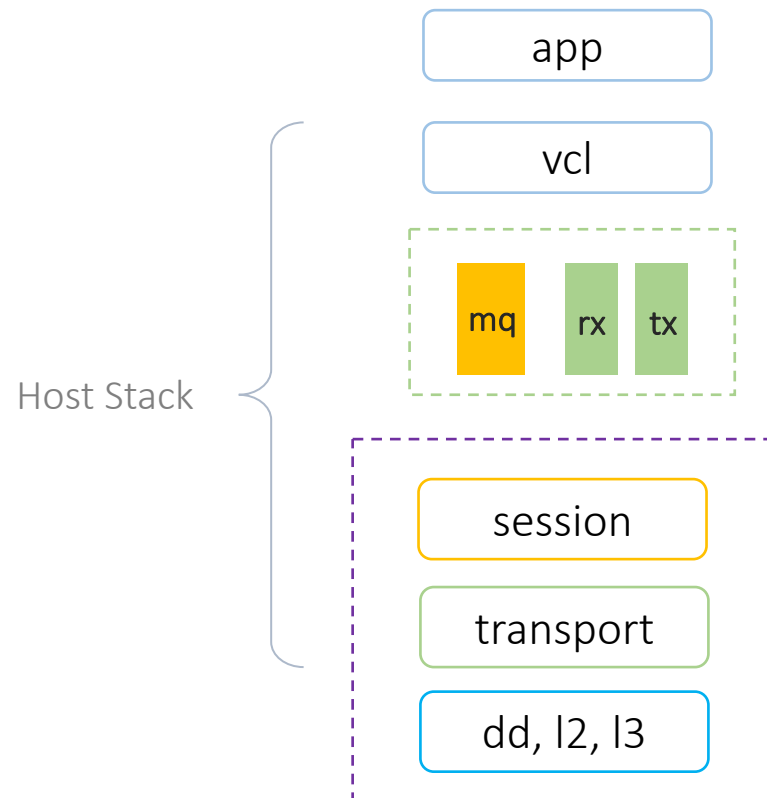


# Network Stack



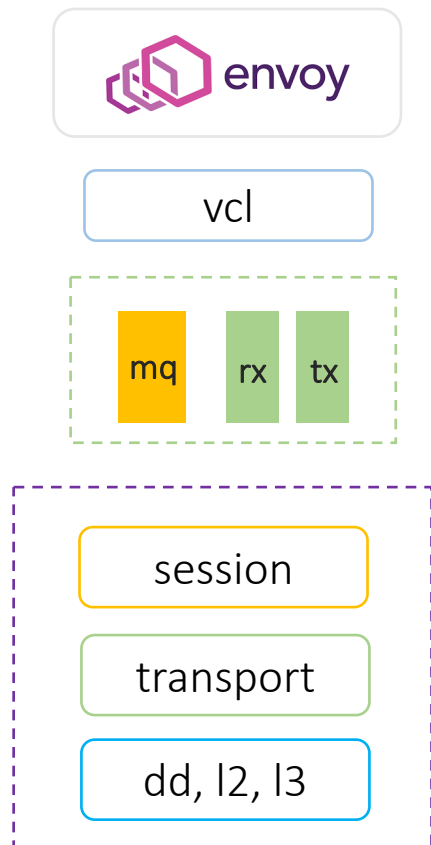
- DPDK and native device drivers
- Switching, Routing
- IPsec, ACL, NAT, MPLS, SR, tunneling ...

# Host Stack



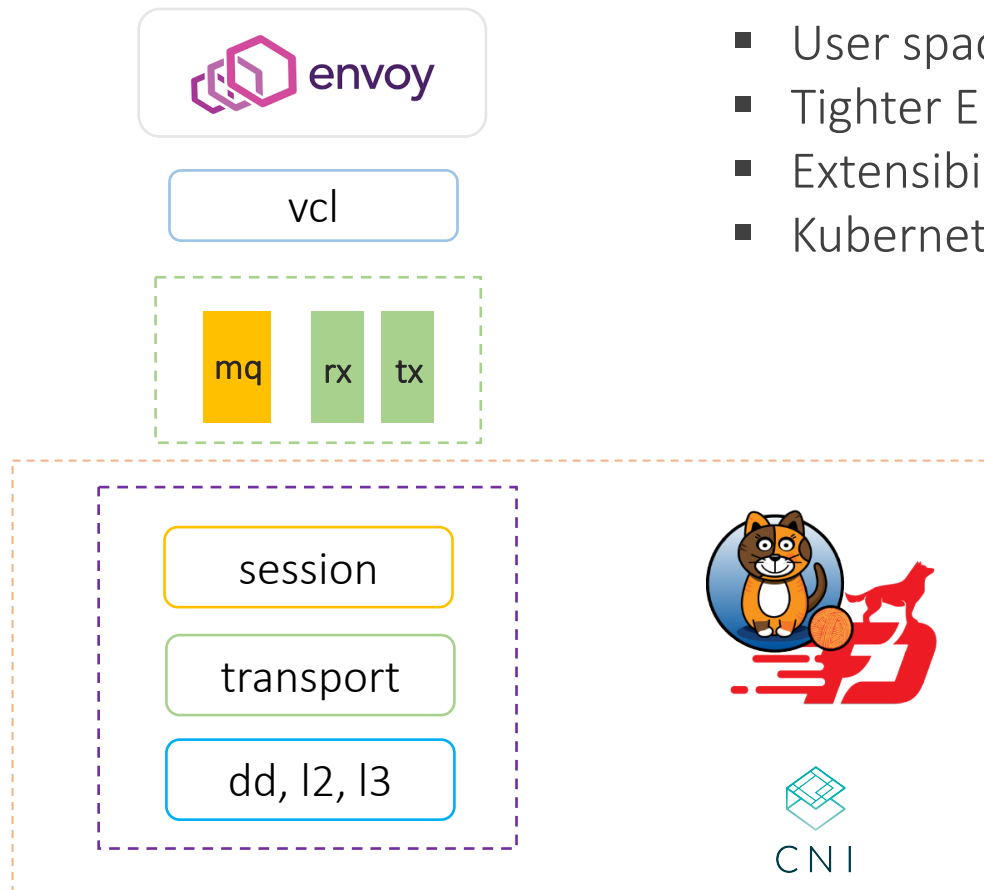
- Transports: TCP, UDP, QUIC, TLS
- Session scheduling, on demand tx pacing, builtin app support
- Shared memory infra for io/ctrl events and data exchanges with external apps
- POSIX-like apis through VCL

# Why?



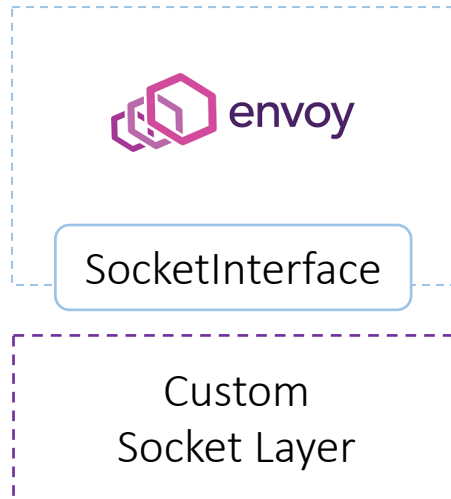
- User space performance
- Tighter Envoy integration with network stack
- Extensibility

# Why?



- User space performance
- Tighter Envoy integration with network stack
- Extensibility
- Kubernetes integration via Calico-VPP ([see talk](#))

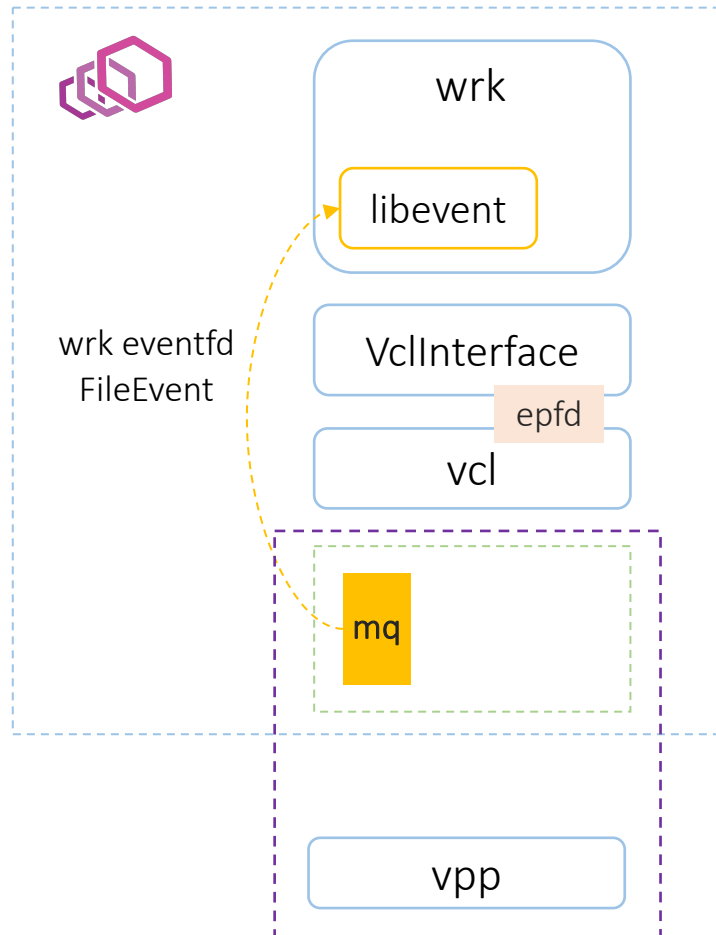
# Envoy VCL Interface



- Prerequisite: Socket and IoHandle refactor
  - Avoid direct use of socket fds
  - Pluggable IoHandle factories (SocketInterfaces)
  - Delegate FileEvent creation to IoHandle
  - Custom BIOs for TLS

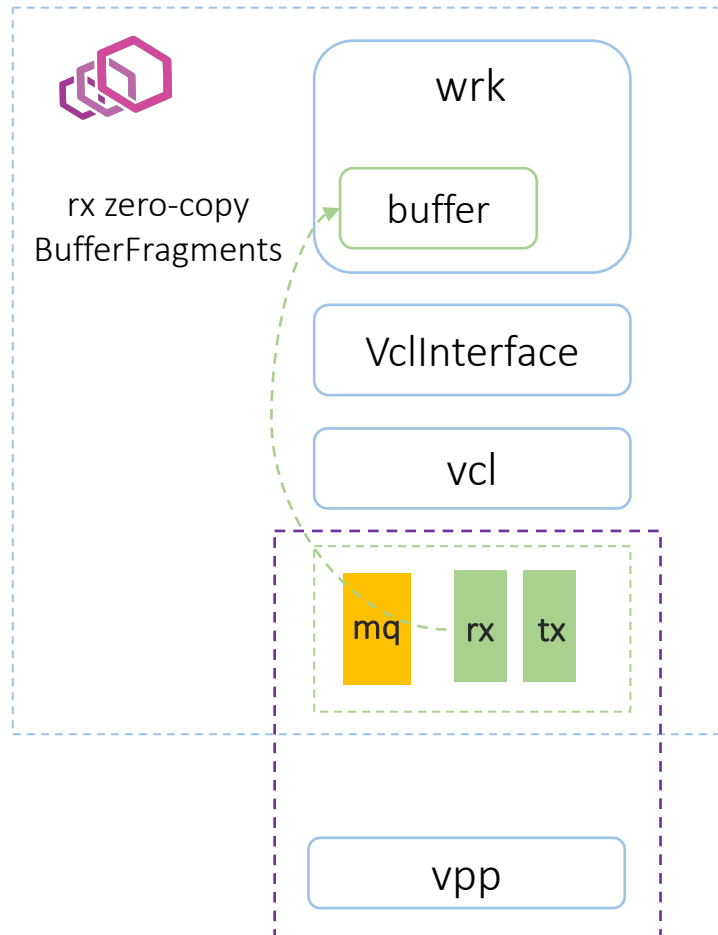


# Envoy VCL Interface



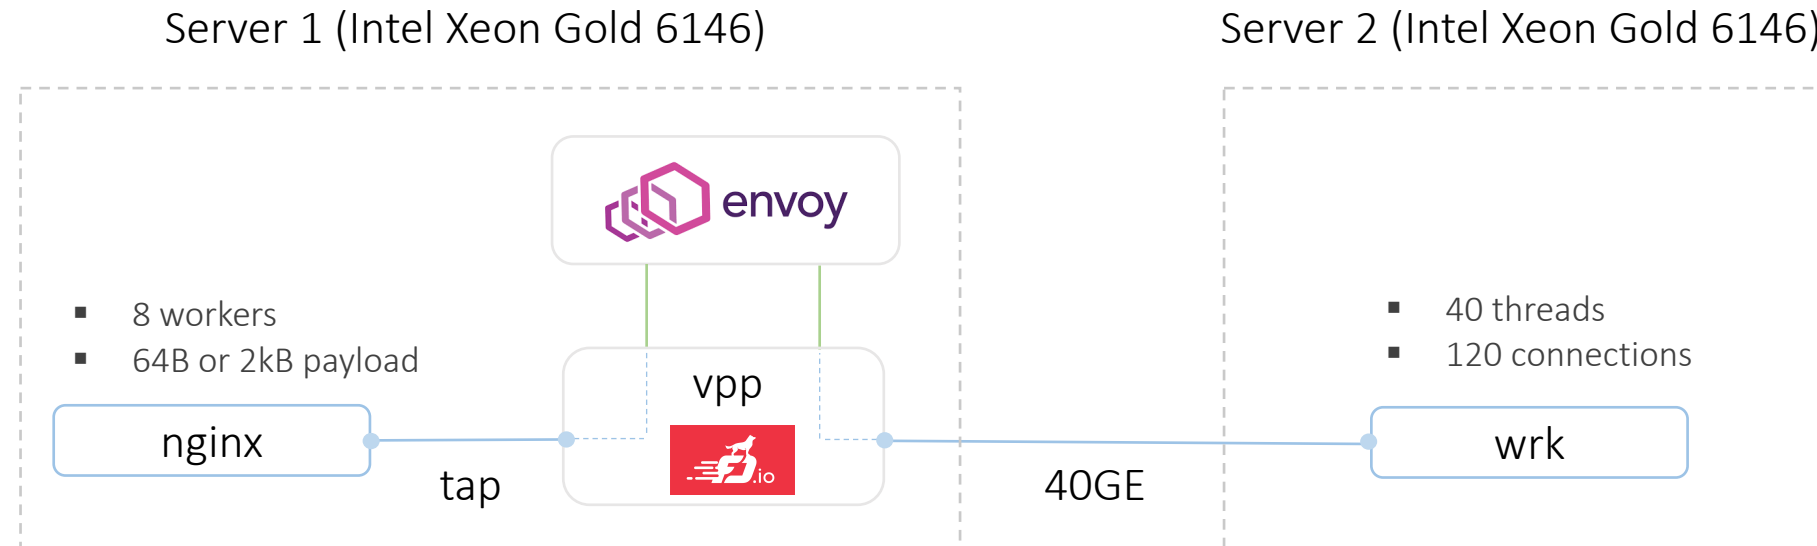
- Prerequisite: Socket and IoHandle refactor
  - Avoid direct use of socket fds
  - Pluggable IoHandle factories (SocketInterfaces)
  - Delegate FileEvent creation to IoHandle
  - Custom BIOs for TLS
- VCL interface
  - Retrieves async VPP session events via VCL worker eventfd FileEvent
  - Maintains per worker epoll fd that tracks all IoHandles
  - Custom VCL IoHandle FileEvents are created and dispatched independent of libevent

# Envoy VCL Interface




- Prerequisite: Socket and IoHandle refactor
  - Avoid direct use of socket fds
  - Pluggable IoHandle factories (SocketInterfaces)
  - Delegate FileEvent creation to IoHandle
  - Custom BIOs for TLS
- VCL interface
  - Retrieves async VPP session events via VCL worker eventfd FileEvent
  - Maintains per worker epoll fd that tracks all IoHandles
  - Custom VCL IoHandle FileEvents are created and dispatched independent of libevent
- Next steps VCL IoHandle Optimizations
  - Optimize socket read operations
  - Investigate minimizing data copying

# Performance




# Performance

### Envoy HTTP Routing (64B)



### Envoy HTTP Routing (2kB)



# Thank You

- Early Envoy-VPP integration code
  - Github: <https://github.com/florincoras/envoy-vpp>
  - Slack: <https://envoyproxy.slack.com/>
- VPP Project:
  - Parent LFN project: <https://fd.io/>
  - Mailing lists: <https://lists.fd.io/g/main>
  - Host Stack documentation: <https://wiki.fd.io/view/VPP/HostStack>