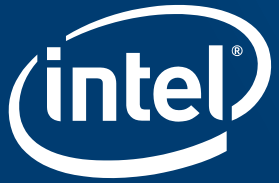

FORTVILLE LINUX IPSEC OFFLOAD

Josh Hay, ND
Don Skidmore, ND
Mar 2018

Network Division

Agenda

- Overview
- Status
- Performance Preview

POC Overview

- No separate control plane for Configuration and Metadata
  - Use one L2 tag/Ethertype to denote Configuration packets
  - Different L2 tag/Ethertype to denote Metadata in a packet

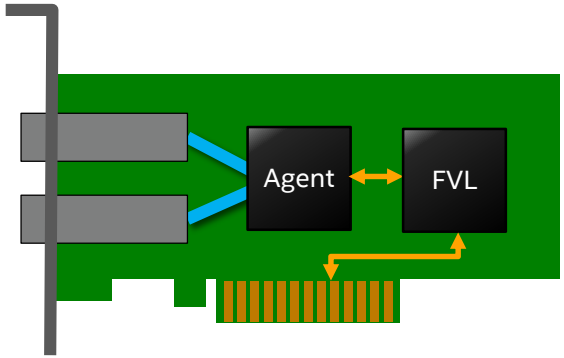

Status

- POC is code complete
  - IPv4, Transport mode, TSO
  - Both control plane and data plane fully integrated
- Validation and debug in progress
- Still no virtualization support

First Glance Numbers

- Formal testing and validation in progress; these are NOT formal results (“dev testing”)
  - Systems are not tuned, not symmetric, but everything run in the same conditions

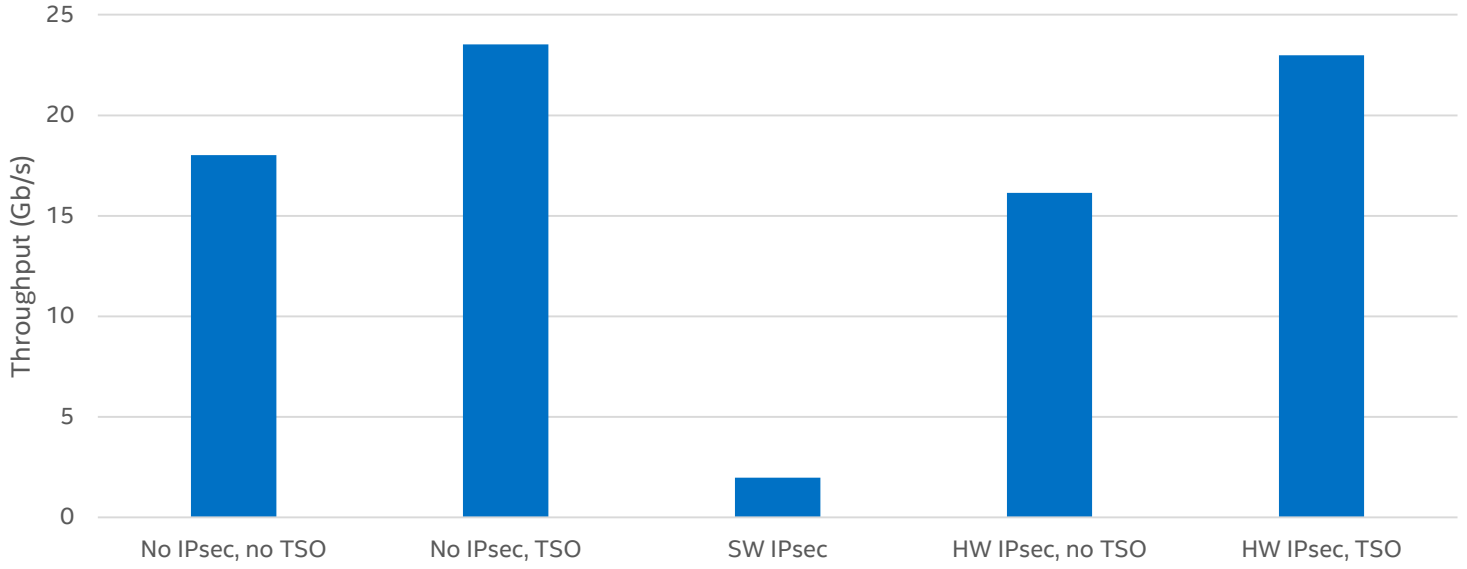

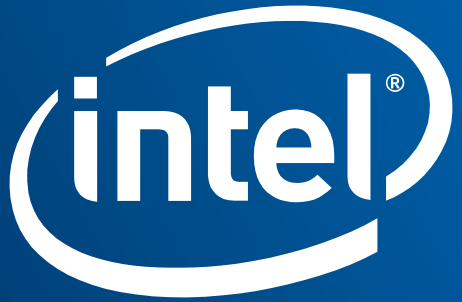
intel

IPsec Control Packet flow

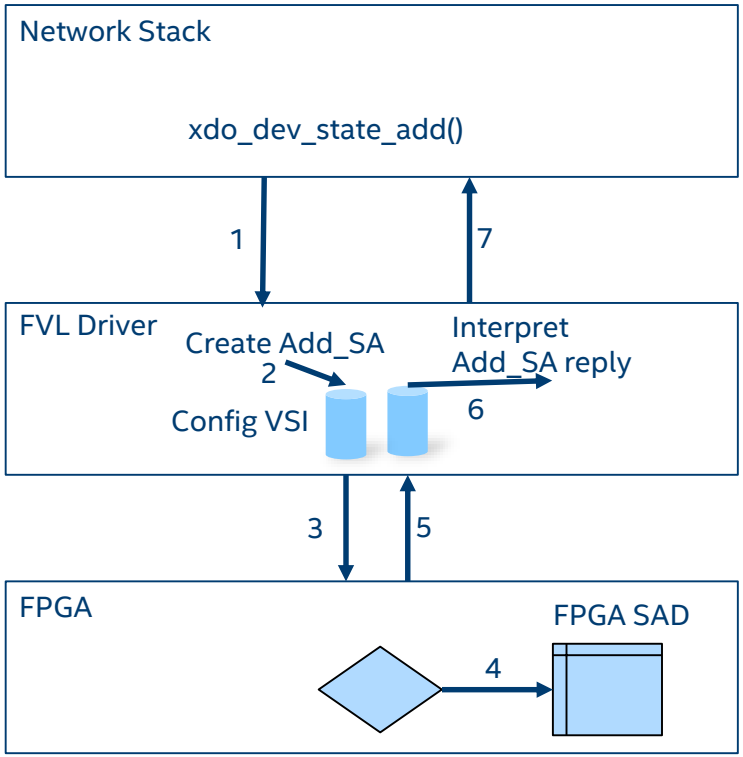

1. Stack calls xdo_dev_state_add to add an SA.
2. Driver creates an Add_SA control packet
3. The Add_SA packet is sent to the FPGA via “control” VSI
4. The FPGA adds the SA to its SAD if possible.
5. The FPGA sends a Add_SA reply to the driver
6. The driver receives the reply and interprets it.
7. The return value to xdo_dev_state_add reflects what we received in the Add_SA reply

Simplified Packet Format

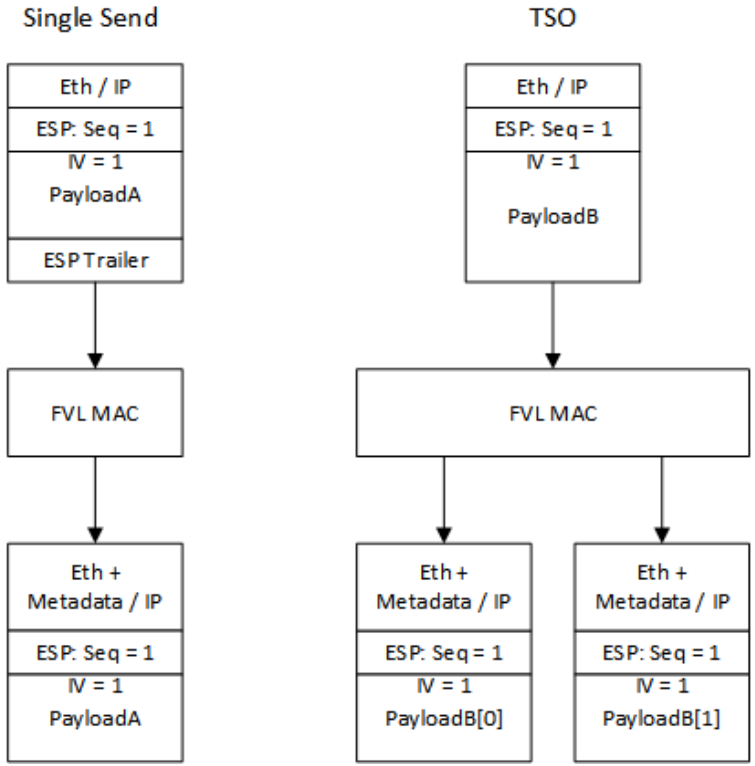

TSO Sequence Number Solution

- Problem: The header is replicated exactly for each segment, but parts of it need to be changed per segment
- Solution: Update RTL to track Sequence Number/IV to the SA entry in the SAD and replace these fields in the packet segments on the fly
  - Also reduces metadata consumption

```
If (IVDB[SA].IV <= packet.IV)
    IVDB[SA].IV = packet.IV + 1
Else if (packet.IV < IVDB[SA].IV)
    packet.IV = IVDB[SA].IV
    IVDB[SA].IV++
```