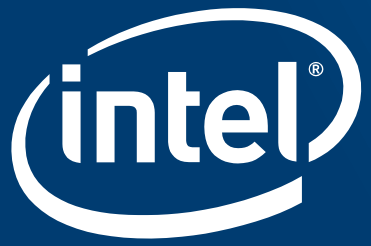

LINUX FULL IPSEC OFFLOAD

Linux IPsec Workshop, March 2018

IPsec Offload Modes

	Crypto Offload	Full IPsec Offload (new)
	Current IPsec offload support in xfrm	Topic for Discussion
SW	IPsec encap/decap Padding insertion/validation Anti-replay Counters updates SPD	
HW	Encrypt/Decrypt/Integrity	IPsec encap/decap Encrypt/Decrypt/Integrity Padding insertion/validation Anti-replay Counter updates SPD

IPsec Full Offload – Discussion Areas

- Handling differing device capabilities
- Exception Handling / SW Fall-back
- Padding
- SA Lifetimes
- Counters
- SPD offload

IPsec Full Offload – Device Capabilities

- How to distinguish between full offload and encrypt/decrypt offload modes?
  - Use existing NETIF_F_HW_ESP & use xfrm_state_offload.flags to indicate different offload modes?
  - Alternative: XFRM user specifies desired offload mode?
  - Alternative: ???
- Device offload capability check as part of xdo_dev_state_add

IPsec Full Offload - Behaviour on Exceptions?

Software Fall-back

- SA miss
- Fragmentation?

Drop Packet

- ICV validation failure
- Anti-replay check failure
- Packet length errors
- Tx packet size exceeds MTU after encap
- Fragmentation

IP Reassembly – Anti-Replay

- Fragments sent to SW for Reassembly + IPsec
- Non-Fragmented packets processed in HW
- Reassembly latencies may cause reassembled packet to fall outside of the anti-replay window.

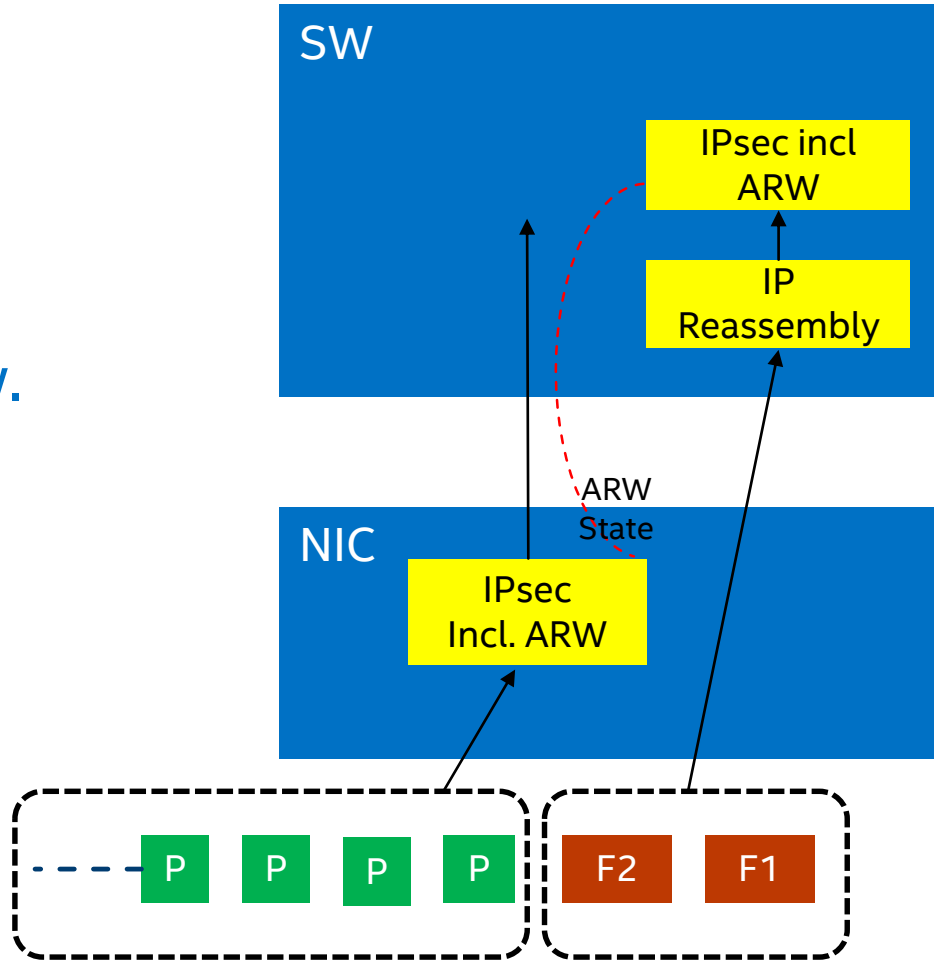

Max Reassembly Time (usec) Before ARW failures

ARW Size ->	128	256	512	1024	4096
10Gbps	10.9	21.7	43.4	86.8	347.3
40Gbps	2.7	5.4	10.9	21.7	86.8
100Gbps	1.1	2.2	4.3	8.7	34.7

Ideal: Reassemble in HW or Drop Fragments

IPsec Full Offload – Padding, Lifetime & Counters

IPsec Padding

- Insert minimum required padding on Tx
- Validate and strip padding on Rx
- TFC?

SA Lifetime Limits

- Byte & Packet count, Elapsed time
- SW vs HW split? Configurable?
- HW generated events on reaching soft/hard limit

Per SA Counters:

- Integrity failures
- Replay failures
- Replay window failures
- Pkts processed?
- Bytes processed?

New xdo callback required to retrieve counters from device.

IPsec Full Offload – SPD Policy

Policy offloaded in all full offload use cases?

Policy selector complexity? Fully compatible with xfrm_selector

IPsec Full Offload – sk_buff

Reuse existing xfrm_offload status & flags to indicate processing performed