



OVS

Open vSwitch

December 2021

# Recirculating the Journey to Higher Performance in the OVS SW Datapath

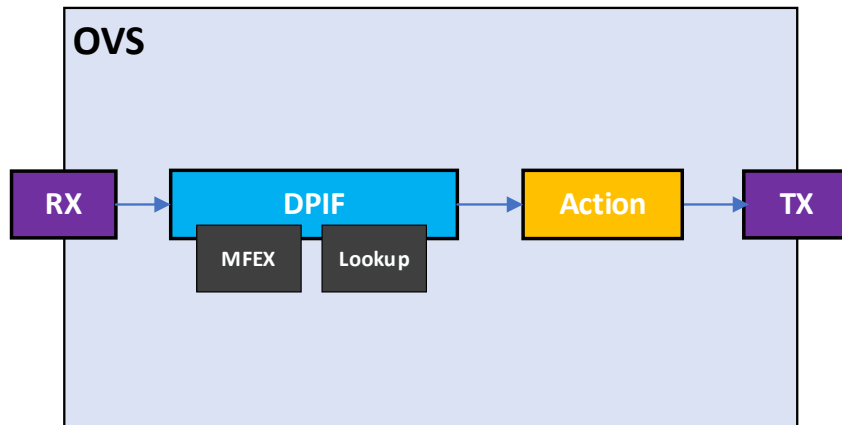
Cian Ferriter, Harry van Haaren, Kumar Amber  
Intel

# OVERVIEW

- **Introduction**
  - DPIF
  - Miniflow Extract
  - Recirculation
- **Enabling AVX-512 Recirculation**
  - For DPIF
  - For MFEX
  - Benefits of AVX-512 Recirculation
- **Validation of AVX-512 Recirculation**

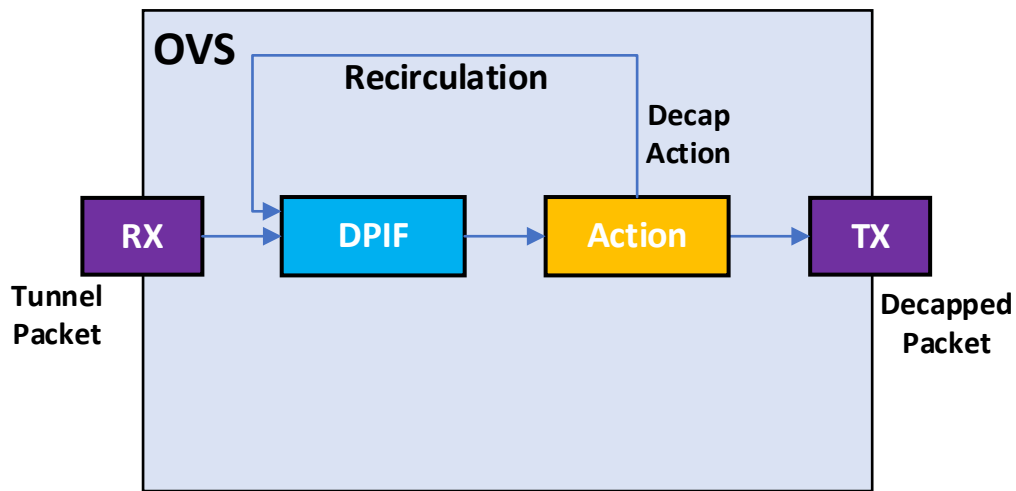
# DPIF AND MFEX

- **DPIF - DataPath InterFace**
  - Connects MFEX, lookup and action blocks
- **MFEX – MiniFlow EXtract**
  - Parses the packet



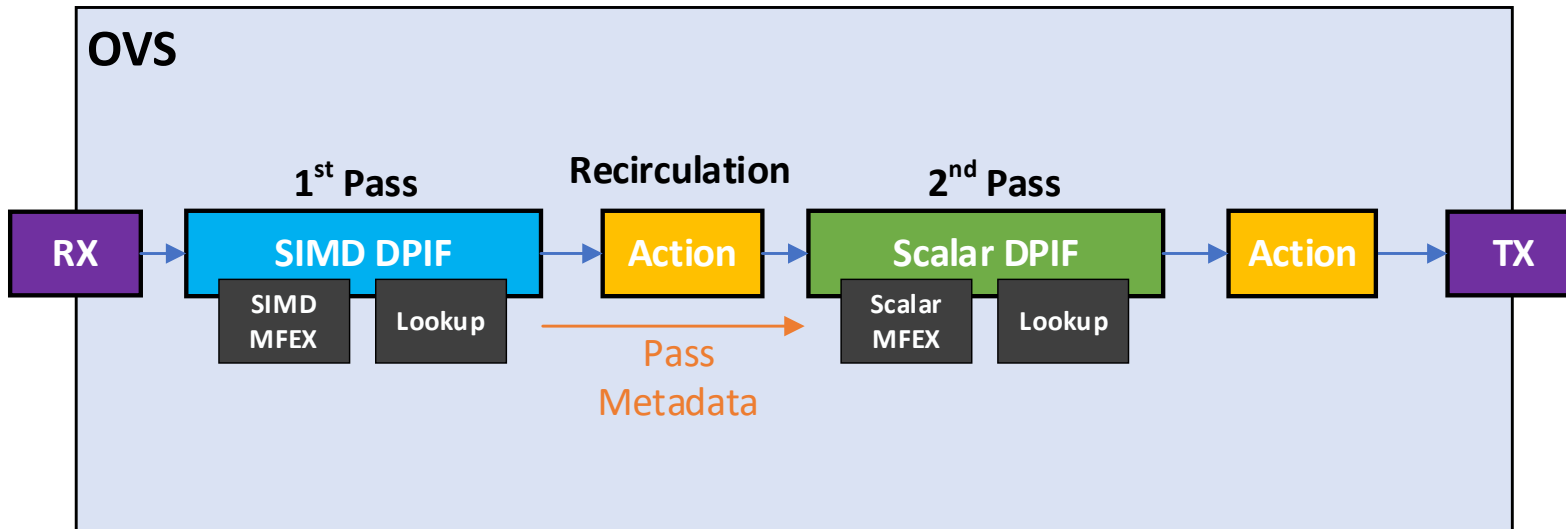
# RECIRCULATION

- **What is recirculation?**
  - Packets traverse OVS datapath multiple times
  - Used in tunnel decap and conntrack



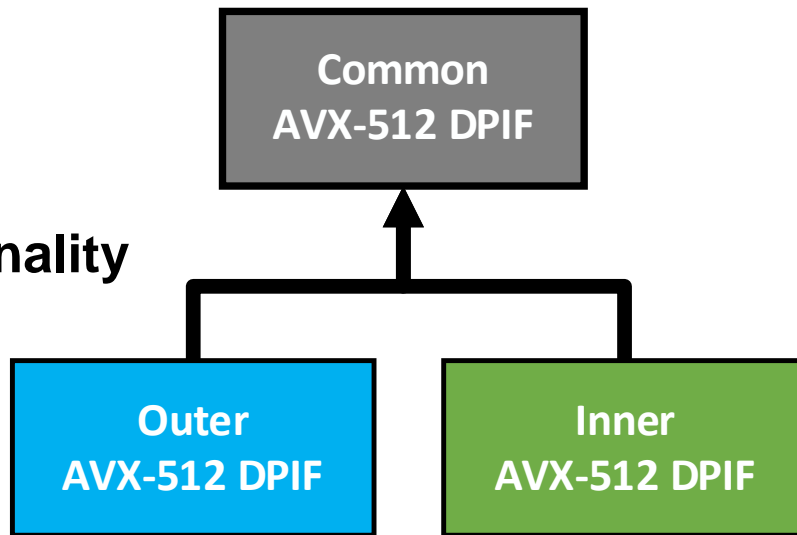
# RECIRCULATION

- Tunnel decap carries extra metadata
  - Makes 2<sup>nd</sup> pass DPIF different



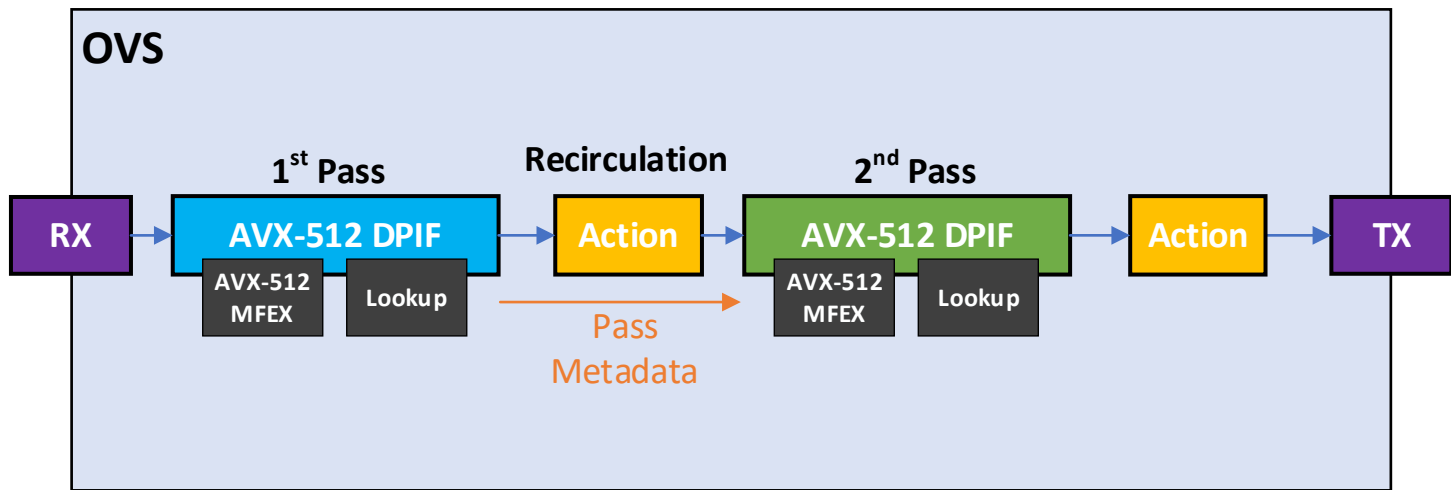
# ENABLING AVX-512 INNER DPIF

- **DPIF changes**
  - Skip metadata initialization
  - Add inner DPIF stats
- **Maintainable DPIFs**
  - Common AVX-512 DPIF functionality
  - Avoid code duplication



# ENABLING AVX-512 INNER MFEX

- Add support for extra data in Miniflow
  - Packet metadata from outer packet
  - Recirculation ID



# BENEFITS OF AVX-512 RECIRCULATION




- **Faster packet batching with AVX-512**
  - Even faster in the worst case (multi-flow)
- **Faster packet parsing**
  - Recirculated packets can have larger miniflows
  - Faster with wider registers



# VALIDATION OF AVX-512 INNER RECIRC

- **Both are validated like outer AVX-512 impls**
- **Use AutoValidator approach for MFEX**
- **Use OVS unit tests for DPIF**
  - **Set AVX-512 inner DPIF on at compile time**

# Code for 2.18!



**Patches will be submitted to the mailing list for review and merge**

**Reviews & Comments welcomed**



**! Thank You !**  
**? Questions ?**

Cian Ferriter [cian.ferriter@intel.com](mailto:cian.ferriter@intel.com)

Harry van Haaren [harry.van.haaren@intel.com](mailto:harry.van.haaren@intel.com)

Kumar Amber [kumar.amber@intel.com](mailto:kumar.amber@intel.com)