



↔ vs

Open vSwitch

December 10th-11th, 2019 | Westford MA

How to dump your miniflow bits

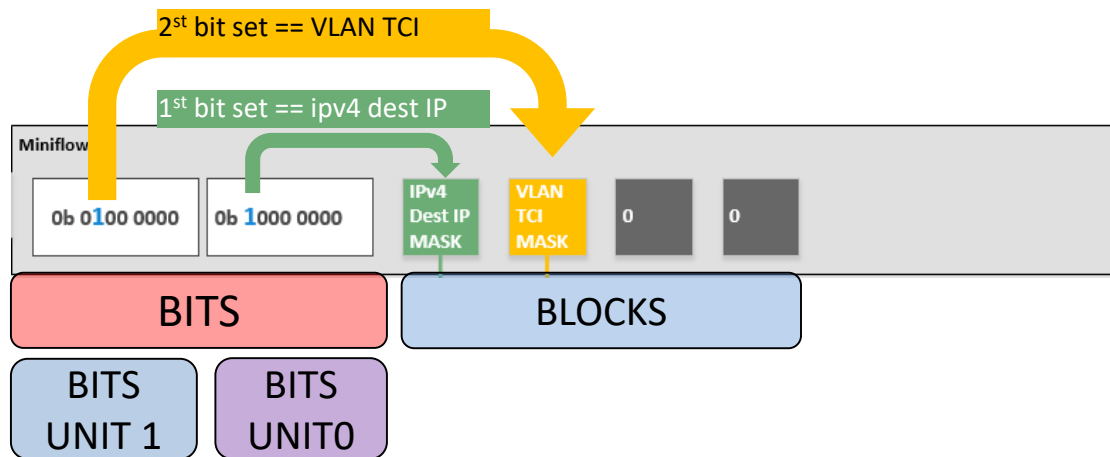
Ian Stokes
Intel

Content

- Miniflow Refresh
- DPCLS Optimization : 2.11 vs 2.12
- Are you benefitting from DPCLS Optimization?
- DPCLS “subtable-show” tool

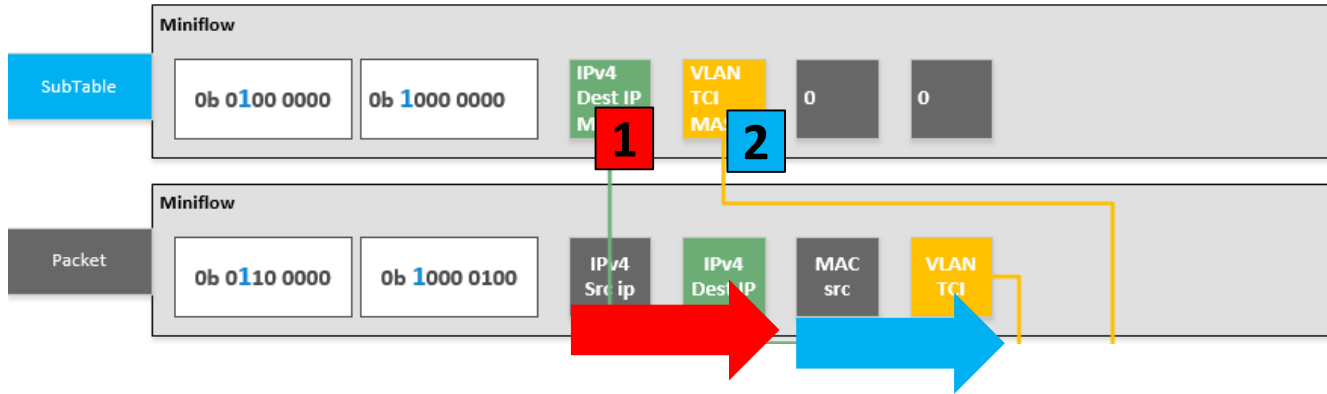
Miniflow Refresh

- **What is a miniflow?**
 - Structure to represent metadata
 - **Packets** and **Subtables** have miniflows



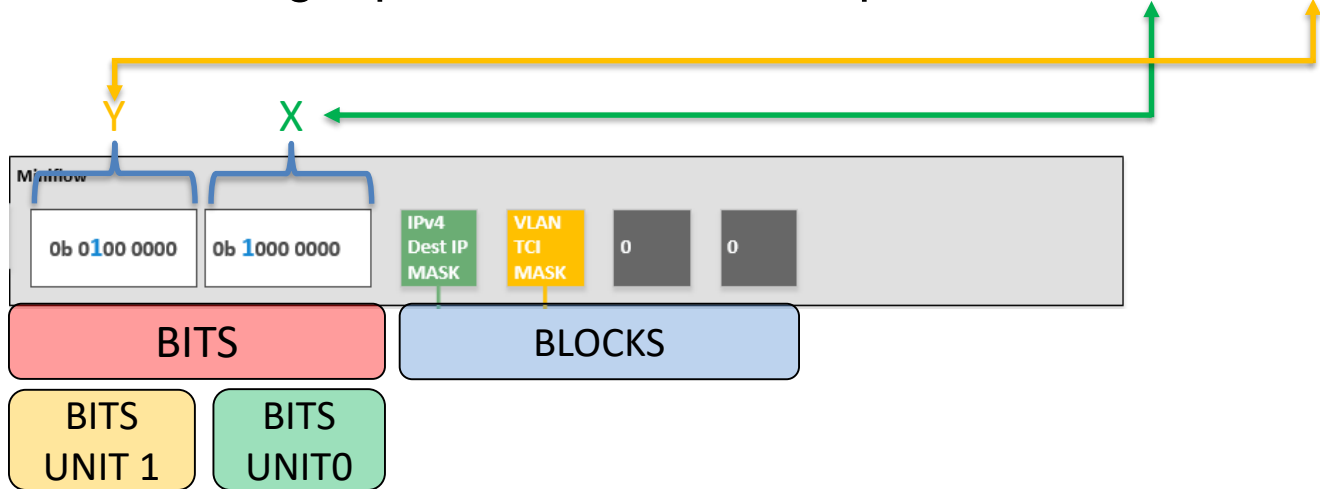
DPCLS Behaviour: 2.11 vs 2.12

- OVS 2.11 - Packet to subtable search



DPCLS Behaviour: 2.11 vs 2.12

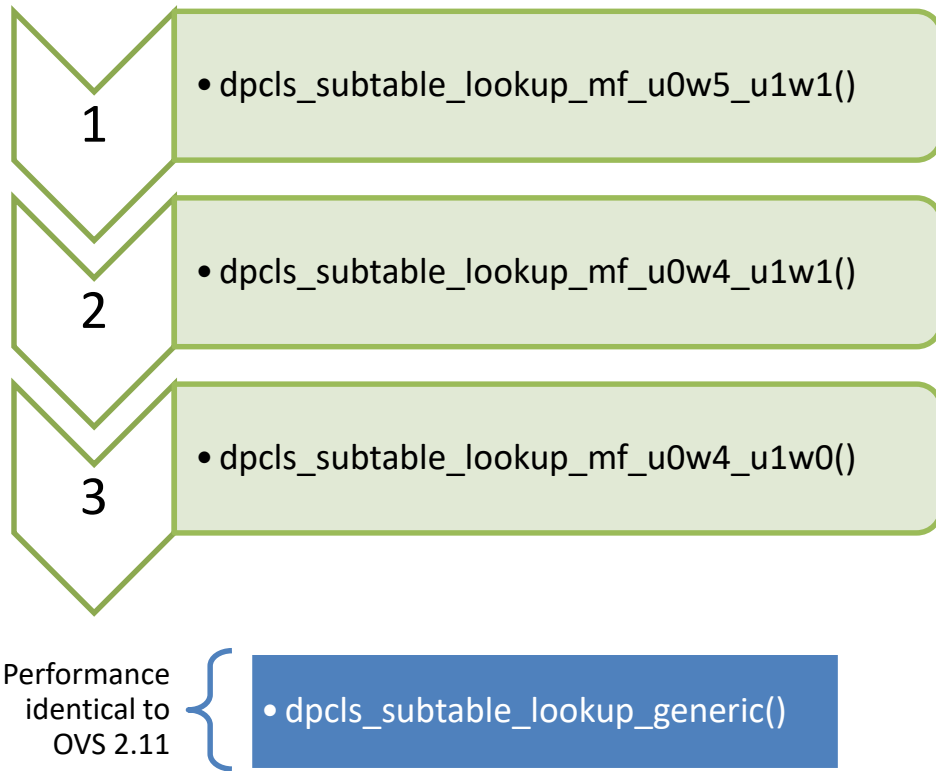
- OVS 2.12
 - DPCLS refactored to support **generic scalar functions**.
 - Subtables can have a specialized miniflow lookup function.
 - E.g. `dpcls_subtable_lookup_mf__u0wX_u1wY`



Result

Community reported
10 - 15%
performance
improvement.

Are you benefitting from DPCLS Optimization?



- Benefit if traffic matches subtable using one of 3 functions.
- Otherwise generic look up is used i.e. no performance gain

- Easy to add new function e.g.

`dpcls_subtable_lookup_mf__u0wX_u1wY()`

- Question:
- What should **X** and **Y** be set to?

DPCLS “subtable-show” tool

- `ovs-appctl dpif-netdev/subtable-show`
 - Iterate through DPCLS subtables.
 - Outputs miniflow bits in hex.

```
-bash-4.4$ sudo ./ovs-appctl dpif-netdev/subtable-show
pmd thread numa_id 0 core_id 2:
  dpcls port 3:
    subtable:
      unit_0: 4 (0x4)
      unit_1: 1 (0x1)
```

- <https://patchwork.ozlabs.org/patch/1191979/>

DPCLS “subtable-show” tool

- `dpcls_subtable_lookup_mf__u0wX_u1wY()`
- Question: What should X and Y be set to?

```
-bash-4.4$ sudo ./ovs-appctl dpif-netdev/subtable-show
pmd thread numa_id 0 core_id 2:
  dpcls port 3:
    subtable:
      unit 0: 4 (0x4) X
      unit 1: 1 (0x1) Y
pmd thread numa_id 0 core_id 3:
  dpcls port 2:
    subtable:
      unit_0: 4 (0x4)
      unit_1: 1 (0x1)
-bash-4.4$
```


Questions

- Email: ian.stokes@intel.com
- Email: emma.finn@intel.com