

# Hands-on Workshop on Open vSwitch and Software-defined Networking

Ali ALSabeh, Jorge Crichigno  
University of South Carolina  
<http://ce.sc.edu/cyberinfra>  
aalsabeh@email.sc.edu, jcrichigno@cec.sc.edu

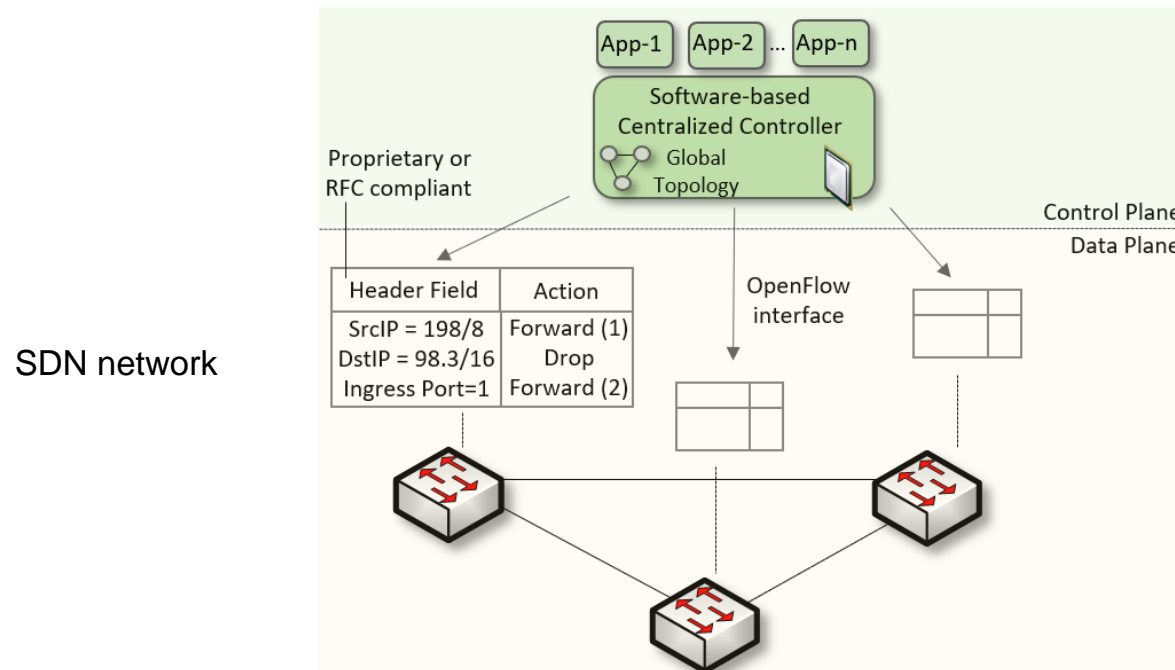
WASTC 2021 virtual Faculty Development Weeks (vFDW)  
June 22, 2021

# Overview of OpenFlow

---

# OpenFlow Overview


- OpenFlow is a protocol specification that describes the communication between OpenFlow switches and an OpenFlow controller
- The consortium responsible for the OpenFlow specification is the Open Networking Foundation (ONF), which was created in 2011



P. Goransson, B. Chuck, C. Timothy, "Software defined networks: a comprehensive approach" Morgan Kaufmann, 2016.


# OpenFlow Switch / Controller

- The core function of a switch is to take packets arriving on one port (path X, port 2) and forward it through another port (port N)
- Potential actions
  - (A) Forward the packet out a local port; (B) Drop the packet; (C) Pass the packet to the controller via a **PKT\_IN** message
- When the controller has a data packet to forward out through the switch, it uses the OpenFlow **PACKET\_OUT** message (e.g., routing advertisements, complex decisions)



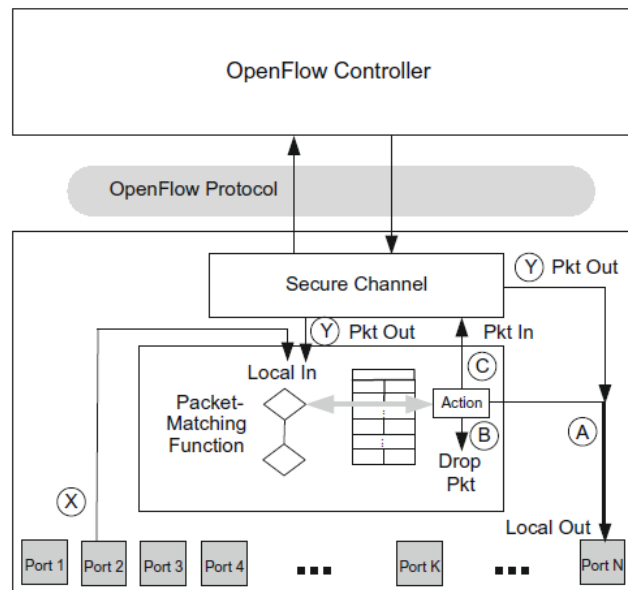
# OpenFlow Protocol

- The protocol consists of a set of messages that are sent from the controller to the switch and a corresponding set of messages that are sent in the opposite direction
- The most basic operations are defining, modifying, and deleting flows
- A flow is a set of packets transferred from one network endpoint to another endpoint



# Controller-switch Secure Channel

- The secure channel is the path used for communications between the OpenFlow controller and the OpenFlow device
- Generally, this communication is secured by TLS-based encryption, though unencrypted TCP connections are allowed
- Connections may be in-band or out-of-band



# Flow Table

- The flow table lies at the core of the definition of an OpenFlow switch
- A flow table consists of flow entries
- A flow entry consists of header fields, counters, and actions associated with that entry

| Flow Entry 0  |  | Flow Entry 1  |                                 |     | Flow Entry F  |                                      |     | Flow Entry M  |                                      |
|---------------|--|---------------|---------------------------------|-----|---------------|--------------------------------------|-----|---------------|--------------------------------------|
| Header Fields | Inport 12<br>192.32.10.1,<br>Port 1012 | Header Fields | Inport *<br>209.*.**,<br>Port * |     | Header Fields | Inport 2<br>192.32.20.1,<br>Port 995 |     | Header Fields | Inport 2<br>192.32.30.1,<br>Port 995 |
| Counters      | val                                    | Counters      | val                             | ■■■ | Counters      | val                                  | ■■■ | Counters      | val                                  |
| Actions       | val                                    | Actions       | val                             |     | Actions       | val                                  |     | Actions       | val                                  |

# Flow Table

- The flow table lies at the core of the definition of an OpenFlow switch
- A flow table consists of flow entries
- A flow entry consists of header fields, counters, and actions associated with that entry
- Example: OvS flow table

Flow table, switch s1

Flow  
entry

```
root@admin:/home/sdn# ovs-ofctl dump-flows s1
cookie=0x10000ea6f4b8e, duration=2519.646s, table=0, n_packets=280, n_bytes=11760, priority=40000,arp actions=CONTROLLER:65535
cookie=0x100009465555a, duration=2519.646s, table=0, n_packets=0, n_bytes=0, priority=40000,dl_type=0x88cc actions=CONTROLLER:65535
cookie=0x100007a585b6f, duration=2519.644s, table=0, n_packets=0, n_bytes=0, priority=40000,dl_type=0x8942 actions=CONTROLLER:65535
cookie=0x10000021b41dc, duration=2123.090s, table=0, n_packets=2, n_bytes=196, priority=5,ip actions=CONTROLLER:65535
cookie=0x5f00002fa2d3c1, duration=2104.661s, table=0, n_packets=2055, n_bytes=201390, priority=10,in_port="s1-eth1",dl_src=ba:03:97:
90:39:4e,dl_dst=e2:8c:0c:de:82:db actions=output:"s1-eth2"
cookie=0x5f000031ebed71, duration=2104.661s, table=0, n_packets=2055, n_bytes=201390, priority=10,in_port="s1-eth2",dl_src=e2:8c:0c:
de:82:db,dl_dst=ba:03:97:90:39:4e actions=output:"s1-eth1"
root@admin:/home/sdn#
```



# Flow Table

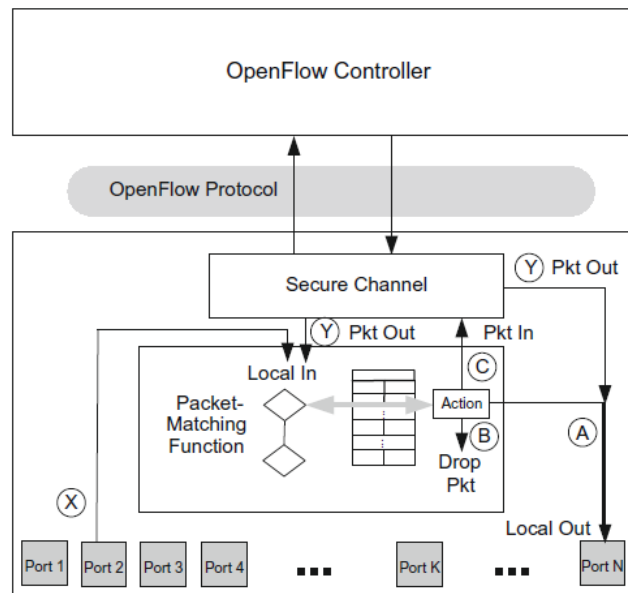
---

- The header fields are used as match criteria to determine whether an incoming packet matches this entry
- The counters are used to track statistics relative to this flow, such as how many packets have been forwarded or dropped for this flow
- The actions fields prescribe what to do with a packet matching this entry

|               |             |
|---------------|-------------|
| Header Fields | Field value |
| Counters      | Field value |
| Actions       | Field value |

# Actions and Packet Forwarding

- The required actions that must be supported by a flow entry are to either forward or drop the matched packet
- The most common case is that the output action specifies a physical port on which the packet should be forwarded



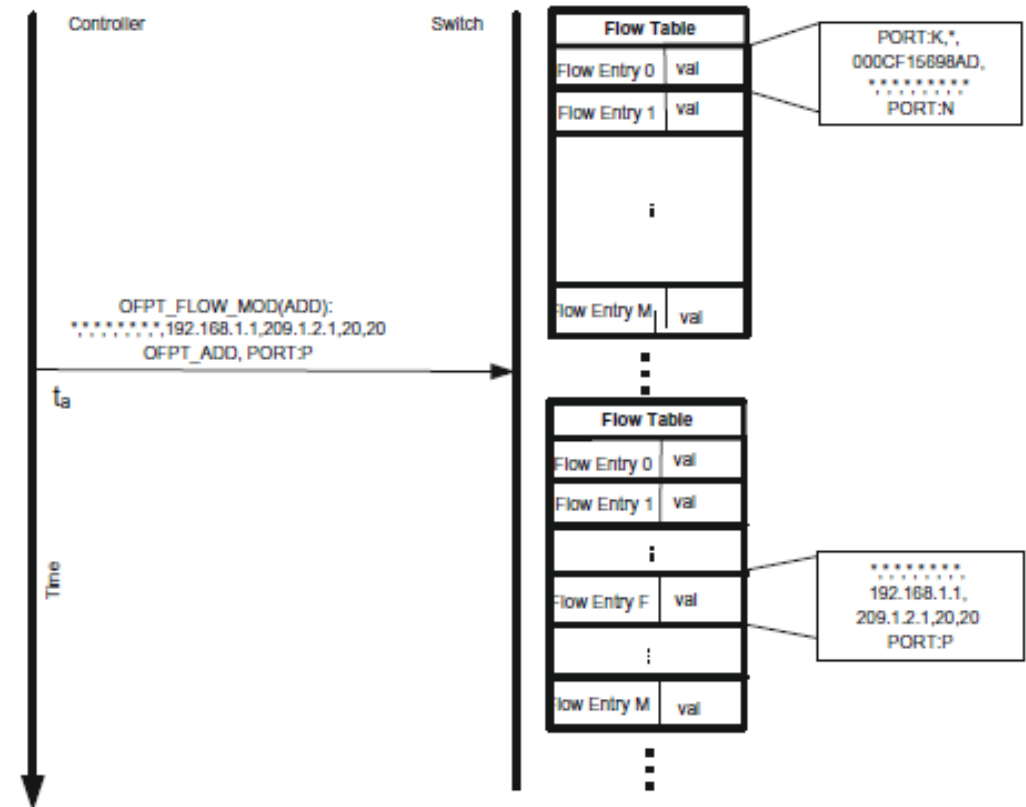
# Messaging between Controller and Switch

- Each message between controller and switch starts with the OpenFlow header
- The header specifies the OpenFlow version, message type, message length, and transaction ID of the message
- Three categories
  - Symmetric: can be sent by controller or switch w/o solicitation
  - Controller-switch: sent by controller to switch
  - Async: can be sent by switch to controller when there is any state change in the system

| Message Type             | Category          | Subcategory          |
|--------------------------|-------------------|----------------------|
| HELLO                    | Symmetric         | Immutable            |
| ECHO_REQUEST             | Symmetric         | Immutable            |
| ECHO_REPLY               | Symmetric         | Immutable            |
| VENDOR                   | Symmetric         | Immutable            |
| FEATURES_REQUEST         | Controller-switch | Switch configuration |
| FEATURES_REPLY           | Controller-switch | Switch configuration |
| GET_CONFIG_REQUEST       | Controller-switch | Switch configuration |
| GET_CONFIG_REPLY         | Controller-switch | Switch configuration |
| SET_CONFIG               | Controller-switch | Switch configuration |
| PACKET_IN                | Async             | NA                   |
| FLOW_REMOVED             | Async             | NA                   |
| PORT_STATUS              | Async             | NA                   |
| ERROR                    | Async             | NA                   |
| PACKET_OUT               | Controller-switch | Cmd from controller  |
| FLOW_MOD                 | Controller-switch | Cmd from controller  |
| PORT_MOD                 | Controller-switch | Cmd from controller  |
| STATS_REQUEST            | Controller-switch | Statistics           |
| STATS_REPLY              | Controller-switch | Statistics           |
| BARRIER_REQUEST          | Controller-switch | Barrier              |
| BARRIER_REPLY            | Controller-switch | Barrier              |
| QUEUE_GET_CONFIG_REQUEST | Controller-switch | Queue configuration  |
| QUEUE_GET_CONFIG_REPLY   | Controller-switch | Queue configuration  |

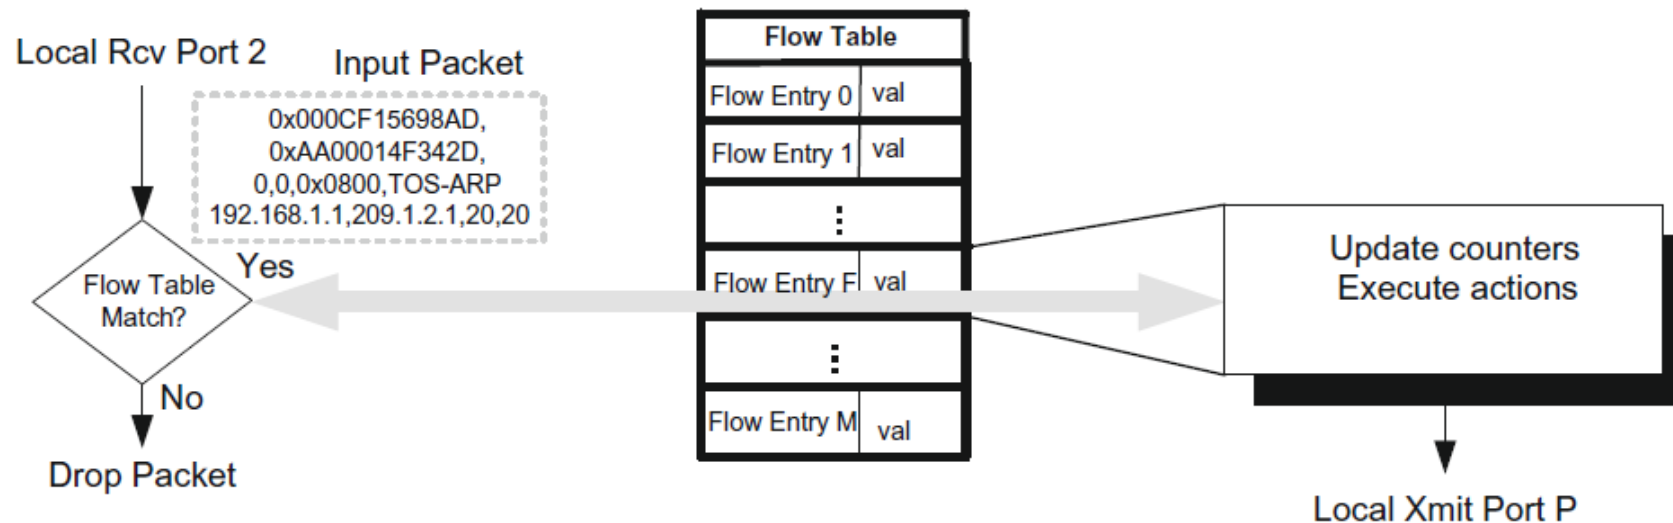
# Example: Controller Programming Flow Table

- At  $t_a$ , the controller sends a FLOW\_MOD (ADD) command
- A flow is added for packets entering the switch on any port
  - Source IP: 192.168.1.1
  - Destination IP: 209.1.2.1
  - Source TCP port: 20
  - Destination TCP port: 20
  - All other match fields have been wildcarded
  - The outport port is specified as P



# Example: Basic Packet Forwarding

- A packet arrives at the switch through port 2 with source IPv4 192.168.1.1 and destination IPv4 209.1.2.1
- The packet-matching function scans the flow table starting at flow entry 0 and finds a match in flow entry F
- Flow entry F stipulates that a matching packet should be forwarded out port P



# OpenFlow Additions

- The OpenFlow interface started simple, with few protocols that could be matched against incoming packets
- Over few years, the specification has been extended with many more header fields and new protocols

| Version      | Date      | Header fields             |
|--------------|-----------|---------------------------|
| OpenFlow 1.0 | Dec. 2009 | 12 (Ethernet, TCP, IPv4)  |
| OpenFlow 1.1 | Feb. 2011 | 15 (MPLS, ...)            |
| OpenFlow 1.2 | Dec. 2011 | 36 (ARP, ICMP, IPv6, ...) |
| OpenFlow 1.3 | Jun. 2012 | 40                        |
| OpenFlow 1.4 | Oct. 2013 | 41                        |
| OpenFlow 1.5 | Mar. 2015 | 44                        |

Bossart et al. "P4: Programming Protocol-Independent Packet Processors"  
OpenFlow Switch Specs v1.5.1. Online <https://tinyurl.com/y4j4a5eh>

# Weakness of SDN / OpenFlow

---

- SDN
  - Fixed number of header fields
  - OpenFlow repeatedly extends the specification
  - Long standardization cycles
  - Fixed protocols / header fields
  - Fixed parser
  - Devices still in control of manufacturers
  - Operators / programmers limited to functionality specified in the OpenFlow specification
  - Match+action stages are in series
- P4 switches (see p4.org)
  - Operators / programmers can define their own protocols and header fields
  - Immediate implementation
  - Customized protocols / header fields
  - Devices in control of operators / programmers
  - Match+action stages are in series or in parallel
  - Actions are composed of protocol-independent primitives (switch is not tight to specific protocols)
  - More future-proof


# Lab 6: Introduction to OpenFlow

---



# Lab 6: Introduction to OpenFlow

- The topology consists of an ONOS controller, an OVS device, and hosts h1 and h2
- The lab shows how to
  - Inspect, add, and remove a flow entry manually in switch s1, using the ovs-ofctl command line utility
  - Use the controller to manage flow entries automatically
  - inspect OpenFlow messages exchanged between the ONOS controller and the OVS switch




# Adding Entries to the Flow Table

Adding flows to forward traffic from host h1 to host h2

```
root@admin: /home/sdn
File Actions Edit View Help
root@admin: /home/sdn
root@admin: /home/sdn# ovs-ofctl add-flow s1 in_port=1,actions=output:2
root@admin: /home/sdn#
```

Adding flows to forward traffic from host h2 to host h1


```
root@admin: /home/sdn
File Actions Edit View Help
root@admin: /home/sdn
root@admin: /home/sdn# ovs-ofctl add-flow s1 in_port=1,actions=output:2
root@admin: /home/sdn# ovs-ofctl add-flow s1 in_port=2,actions=output:1
root@admin: /home/sdn#
```



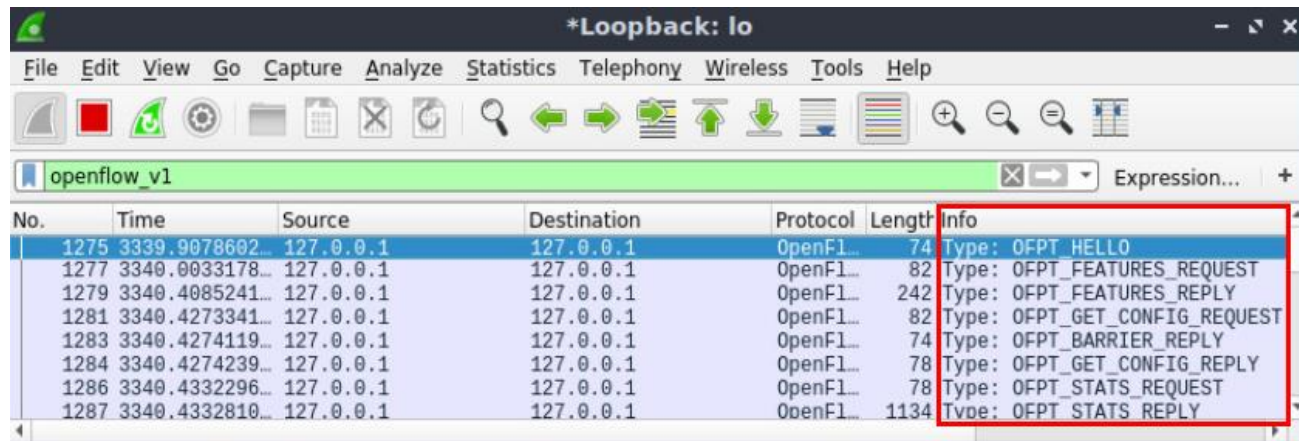
# Activating Application on the Controller

Enabling reactive forwarding application

```
sdn@admin: ~/SDN_Labs/lab6
File Actions Edit View Help
sdn@admin: ~/SDN_Labs/lab6
sdn@root > app activate org.onosproject.fwd
Activated org.onosproject.fwd
sdn@root >
```

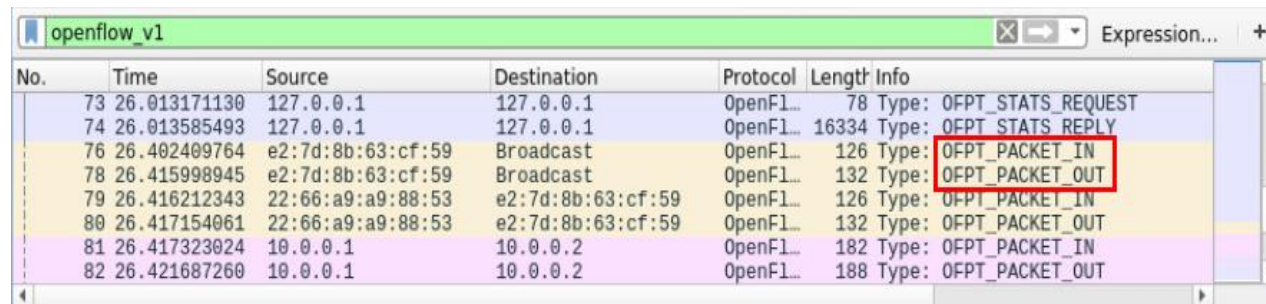


# Capturing OpenFlow Messages



Wireshark capture window titled '\*Loopback: lo' showing a list of OpenFlow messages. The 'Info' column for the last message is highlighted with a red box.

| No.  | Time            | Source    | Destination | Protocol  | Length | Info                          |
|------|-----------------|-----------|-------------|-----------|--------|-------------------------------|
| 1275 | 3339.9078602... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 74     | Type: OFPT_HELLO              |
| 1277 | 3340.0033178... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 82     | Type: OFPT_FEATURES_REQUEST   |
| 1279 | 3340.4085241... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 242    | Type: OFPT_FEATURES_REPLY     |
| 1281 | 3340.4273341... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 82     | Type: OFPT_GET_CONFIG_REQUEST |
| 1283 | 3340.4274119... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 74     | Type: OFPT_BARRIER_REPLY      |
| 1284 | 3340.4274239... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 78     | Type: OFPT_GET_CONFIG_REPLY   |
| 1286 | 3340.4332296... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 78     | Type: OFPT_STATS_REQUEST      |
| 1287 | 3340.4332810... | 127.0.0.1 | 127.0.0.1   | OpenFl... | 1134   | Type: OFPT_STATS_REPLY        |



Wireshark capture window titled 'openflow\_v1' showing a list of OpenFlow messages. The 'Info' column for messages 76 and 78 is highlighted with a red box.

| No. | Time         | Source            | Destination       | Protocol  | Length | Info                     |
|-----|--------------|-------------------|-------------------|-----------|--------|--------------------------|
| 73  | 26.013171130 | 127.0.0.1         | 127.0.0.1         | OpenFl... | 78     | Type: OFPT_STATS_REQUEST |
| 74  | 26.013585493 | 127.0.0.1         | 127.0.0.1         | OpenFl... | 16334  | Type: OFPT_STATS_REPLY   |
| 76  | 26.402409764 | e2:7d:8b:63:cf:59 | Broadcast         | OpenFl... | 126    | Type: OFPT_PACKET_IN     |
| 78  | 26.415998945 | e2:7d:8b:63:cf:59 | Broadcast         | OpenFl... | 132    | Type: OFPT_PACKET_OUT    |
| 79  | 26.416212343 | 22:66:a9:a9:88:53 | e2:7d:8b:63:cf:59 | OpenFl... | 126    | Type: OFPT_PACKET_IN     |
| 80  | 26.417154061 | 22:66:a9:a9:88:53 | e2:7d:8b:63:cf:59 | OpenFl... | 132    | Type: OFPT_PACKET_OUT    |
| 81  | 26.417323024 | 10.0.0.1          | 10.0.0.2          | OpenFl... | 182    | Type: OFPT_PACKET_IN     |
| 82  | 26.421687260 | 10.0.0.1          | 10.0.0.2          | OpenFl... | 188    | Type: OFPT_PACKET_OUT    |

