

Hands-on Open vSwitch and Software-defined Networking

Jorge Crichigno, Shahrin Sharif, Elie Kfoury

University of South Carolina

<http://ce.sc.edu/cyberinfra>

jcrichigno@cec.sc.edu, ssharif@email.sc.edu, ekfoury@email.sc.edu

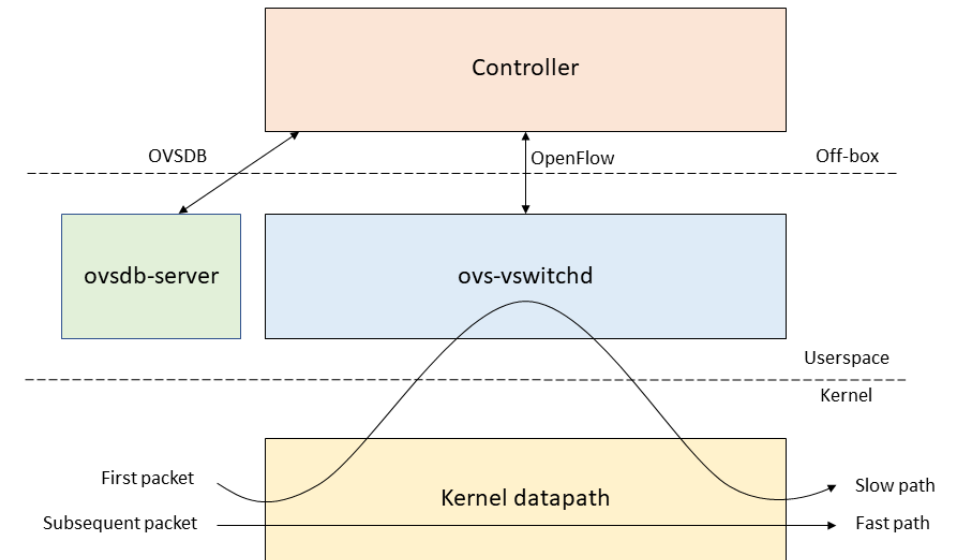
WASTC 2021 virtual Faculty Development Weeks (vFDW)

June 24, 2021


Lab 9: Open vSwitch Kernel Datapath

Open vSwitch Kernel Module

- Kernel module directs the packet to the userspace. The userspace makes the decisions about the actions to be taken against the packet according to OpenFlow entries (slow path)
- The action entry is stored in the kernel and used to forward subsequent packets which makes the forwarding faster (fast path)




Lab Topology



Slow path

- Initially, kernel table is empty


```
ovs@admin: ~  
File Actions Edit View Help  
ovs@admin: ~  
ovs@admin:~$ sudo ovs-dpctl dump-flows  
ovs@admin:~$
```



Fast path

- The userspace makes the decisions about the actions to be taken

```
"Host: h1"  
root@admin:/home/ovs# ping 10.0.0.2  
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.  
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=0.933 ms  
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=0.096 ms  
64 bytes from 10.0.0.2: icmp_seq=3 ttl=64 time=0.071 ms  
64 bytes from 10.0.0.2: icmp_seq=4 ttl=64 time=0.065 ms  
64 bytes from 10.0.0.2: icmp_seq=5 ttl=64 time=0.120 ms  
64 bytes from 10.0.0.2: icmp_seq=6 ttl=64 time=0.077 ms  
64 bytes from 10.0.0.2: icmp_seq=7 ttl=64 time=0.067 ms
```



Fast path

- Action entries are installed in the kernel

```
ovs@admin: ~  
File Actions Edit View Help  
ovs@admin: ~  
ovs@admin:~$ sudo ovs-dpctl dump-flows  
recirc_id(0),in_port(5),eth(src=00:00:00:00:00:02,dst=00:00:00:00:00:01),eth_type  
(0x0800),ipv4(frag=no), packets:10, bytes:980, used:0.316s, actions:3  
recirc_id(0),in_port(6),eth(src=00:00:00:00:00:01,dst=00:00:00:00:00:02),eth_type  
(0x0800),ipv4(frag=no), packets:10, bytes:980, used:0.316s, actions:4  
recirc_id(0),in_port(3),eth(src=00:00:00:00:00:01,dst=00:00:00:00:00:02),eth_type  
(0x0800),ipv4(frag=no), packets:10, bytes:980, used:0.316s, actions:5  
recirc_id(0),in_port(4),eth(src=00:00:00:00:00:02,dst=00:00:00:00:00:01),eth_type  
(0x0800),ipv4(frag=no), packets:10, bytes:980, used:0.316s, actions:6  
recirc_id(0),in_port(3),eth(src=00:00:00:00:00:01,dst=00:00:00:00:00:02),eth_type  
(0x0806), packets:1, bytes:42, used:5.437s, actions:5  
recirc_id(0),in_port(6),eth(src=00:00:00:00:00:01,dst=00:00:00:00:00:02),eth_type  
(0x0806), packets:0, bytes:0, used:never, actions:4  
recirc_id(0),in_port(5),eth(src=00:00:00:00:00:02,dst=00:00:00:00:00:01),eth_type  
(0x0806), packets:1, bytes:42, used:5.437s, actions:3  
recirc_id(0),in_port(4),eth(src=00:00:00:00:00:02,dst=00:00:00:00:00:01),eth_type  
(0x0806), packets:1, bytes:42, used:5.437s, actions:6  
ovs@admin:~$
```

