



Open vSwitch with P4

Cian Ferriter

DPDK Summit Userspace - Dublin- 2016



Legal Disclaimer



General Disclaimer:

© Copyright 2016 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Inside, the Intel Inside logo, Intel Experience What's Inside are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

Technology Disclaimer:

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com].

Performance Disclaimers:

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modelling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

- ▶ P4
- ▶ Benefits of P4 with OVS
- ▶ Packet Processing in OVS
- ▶ OVS vs OVS + P4 Comparison
- ▶ Questions

- ▶ Programming Protocol-Independent Packet Processors
- ▶ High-level Domain Specific Language
- ▶ Describes packet processing applications
- ▶ Open source

Protocol Independent/Reconfigurable

```
table routing {
  reads {
    ipv4.dstAddr : lpm;
  }
  actions {
    do_drop;
    route_ipv4;
  }
  size: 2048;
}

control ingress {
  apply(routing);
}
```



Target Independent



General
Purpose CPUs



FPGAs



NICS



GPUs

Benefits of P₄ with OVS




- ▶ Easy addition of new protocols
- ▶ Performance
 - ▶ Match and parse only on the fields relevant to our switch. E.g. an L2 switch doesn't need to parse L3 headers
- ▶ Separation of concerns
- ▶ Bring OVS-DPDK Userspace Datapath on par with Kernel Datapath
 - ▶ Datapath is implementation-agnostically defined in P4

Change	OVS LOC	P4 LOC
TCP Flags	370	4
Connection Label	411	5

Packet Processing in OVS



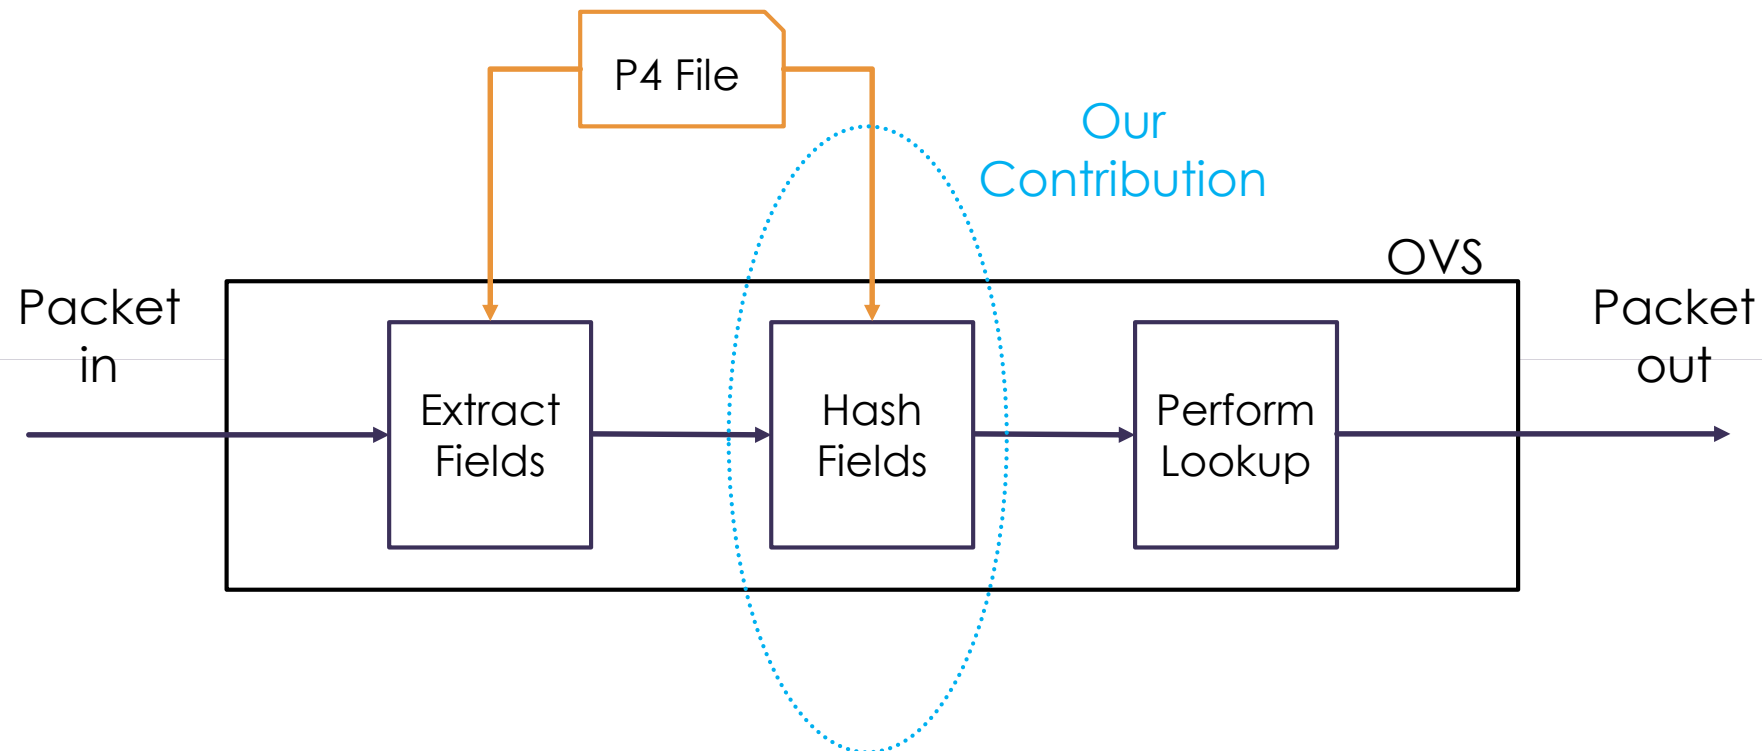
- ▶ In current OVS, the set of fields extracted and hashed is hardcoded
- ▶ Implementation has to support a lot different protocols



Packet Processing in OVS




- ▶ Extract and Hash stages dynamically generated
- ▶ OVS tuned to specific use case



OVS vs OVS + P4 Comparison



Phy-Phy, L3 router case, 1 Core, 64Byte, 1 flow.



Test and System Configurations: Estimates are based on internal Intel analysis using Intel® Server Board S2600WT, Intel(R) Xeon(R) CPU E5-2695 v3 @ 2.30GHz, Intel® 82599ES 10 Gigabit Ethernet Controller

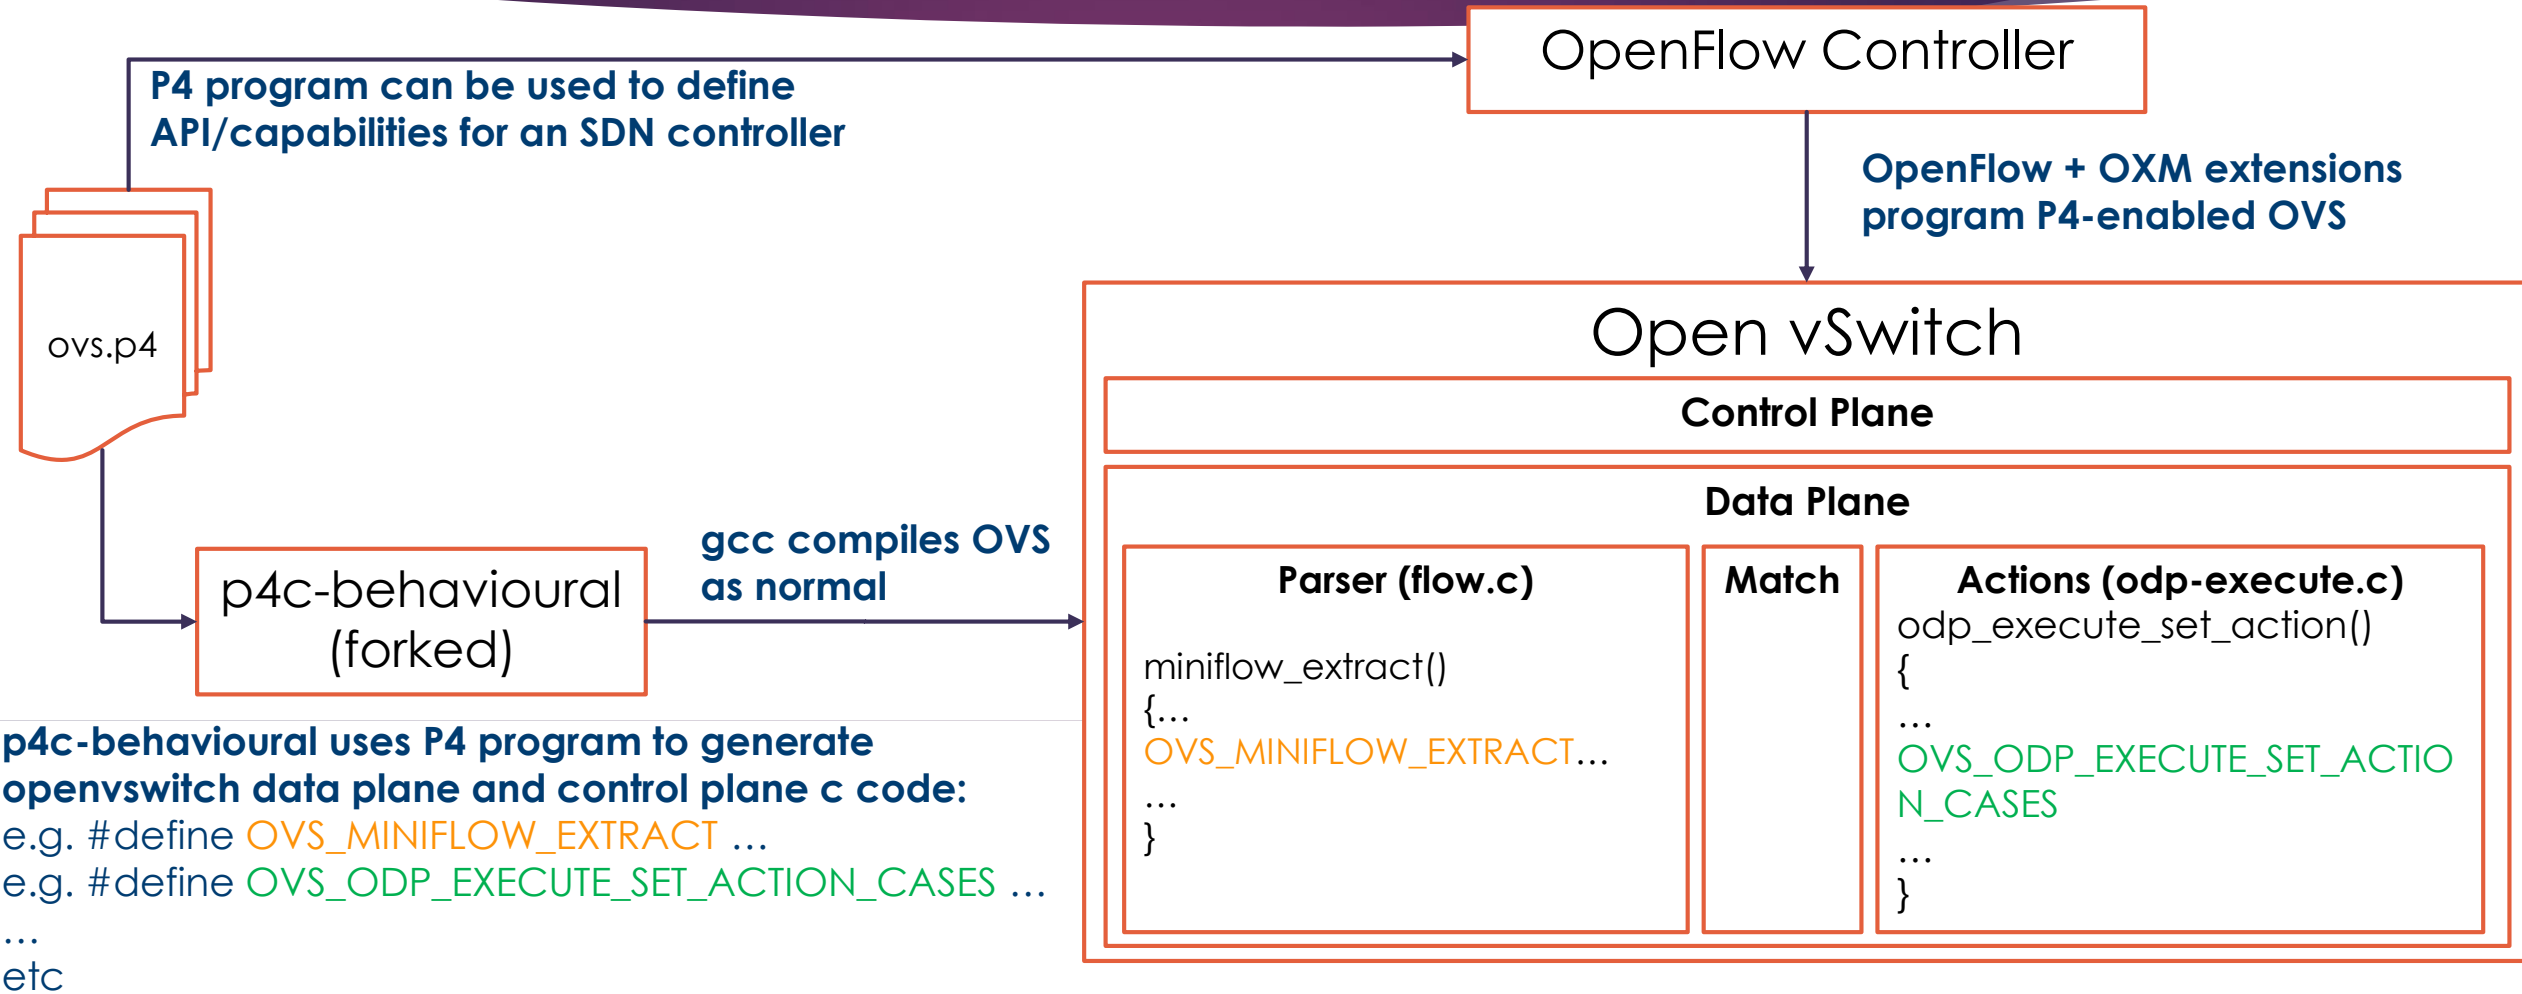
Questions?

Cian Ferriter

cian.ferriter@intel.com

Support Slides

OVS P4 Detailed Architecture



- Main web portal
 - <http://pisces.cs.princeton.edu/>
- Pisces Papers
 - <http://pisces.cs.princeton.edu/papers/sigcomm16-pisces.pdf>
- OVS Trees
 - <https://github.com/mshahbaz/ovs/tree/p4>
 - <https://github.com/mshahbaz/p4c-behavioral/tree/ovs>
- OVS Patches
 - <http://openvswitch.org/pipermail/dev/2016-September/079647.html>
 - <http://openvswitch.org/pipermail/dev/2016-September/079647.html>