



South-bound Interface for Programmable Data-plane Panel

fd.io
OpenSwitch
P4
POF

Panelists:

[Haoyu Song](#), Huawei

[Jan Medved](#), Cisco

[Michael Zayats](#), HPE

[Prem Jonnalagadda](#), Barefoot Networks

Moderator:

[Anu Mercian](#), HPE

Key points

- Introduction to the technologies
- Programmable data plane in OpenDayLight
- Integration efforts in SDN controllers for programmable data planes
- Use cases in general

Data-Programmability

- Fd.io – fast-data input/output

webpage: <https://fd.io/>

- OpenSwitch – Network Operating System

webpage: <http://www.openswitch.net/>

- P4 – Programming Protocol-independent Packet Processors

webpage: <http://p4.org/>

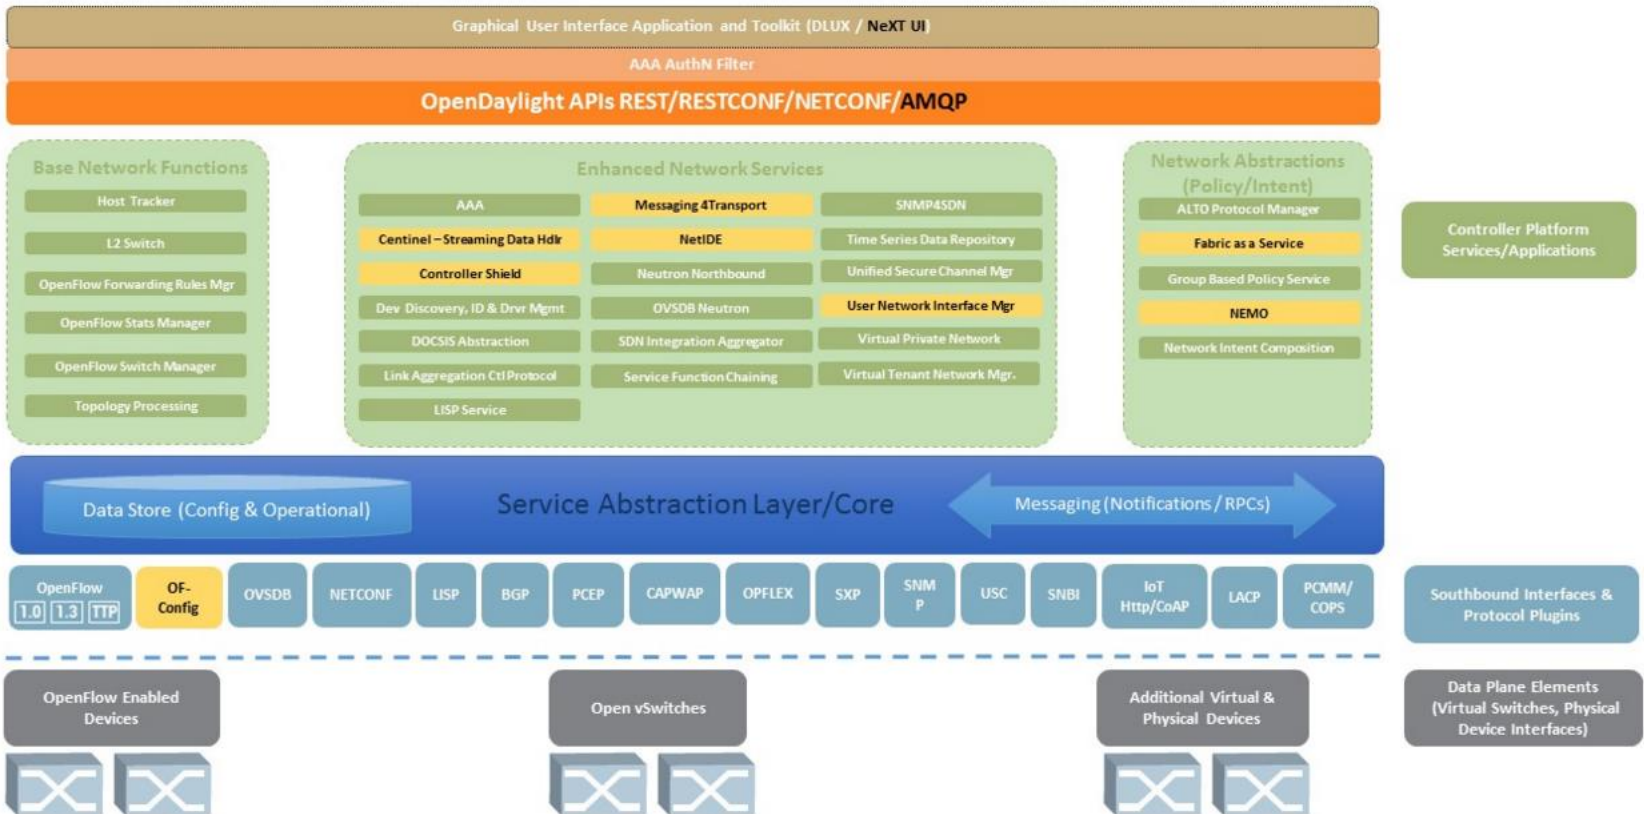
- POF – Protocol-Oblivious Forwarding

webpage: <http://www.poforwarding.org/>

Beryllium Release



4th Release "Beryllium"
Production-Ready Open SDN Platform



Thank you



OPEN NETWORKING
SUMMIT 2016
MARCH 14-17, 2016 | SANTA CLARA, CA



OpenSwitch

Michael Zayats

[Home](#)



OPEN NETWORKING
SUMMIT 2016
MARCH 14-17, 2016 | SANTA CLARA, CA



Motivation for a New NOS

- Open and community driven
- Support common network use-cases - HA, L2, L3, OpenFlow,...
- Consistent configuration, monitoring, troubleshooting
- Declarative, programmatic interfaces (REST/OVSDB) and CLI / GUI



OPEN NETWORKING
SUMMIT 2016
MARCH 14-17, 2016 | SANTA CLARA, CA



Open Community Enablement

- Active mailing list with an average of 25 mails per day
- Weekly + frequent on-demand IRC chats
- All code/design reviews are public using Gerrit
- New code is mostly Apache 2.0, hosted on git.openswitch.net and mirrored to GitHub
- Defined releases and schedules
- Yocto for image building and development environment
- Advanced virtual/physical test framework



OPEN NETWORKING
SUMMIT 2016
MARCH 14-17, 2016 | SANTA CLARA, CA

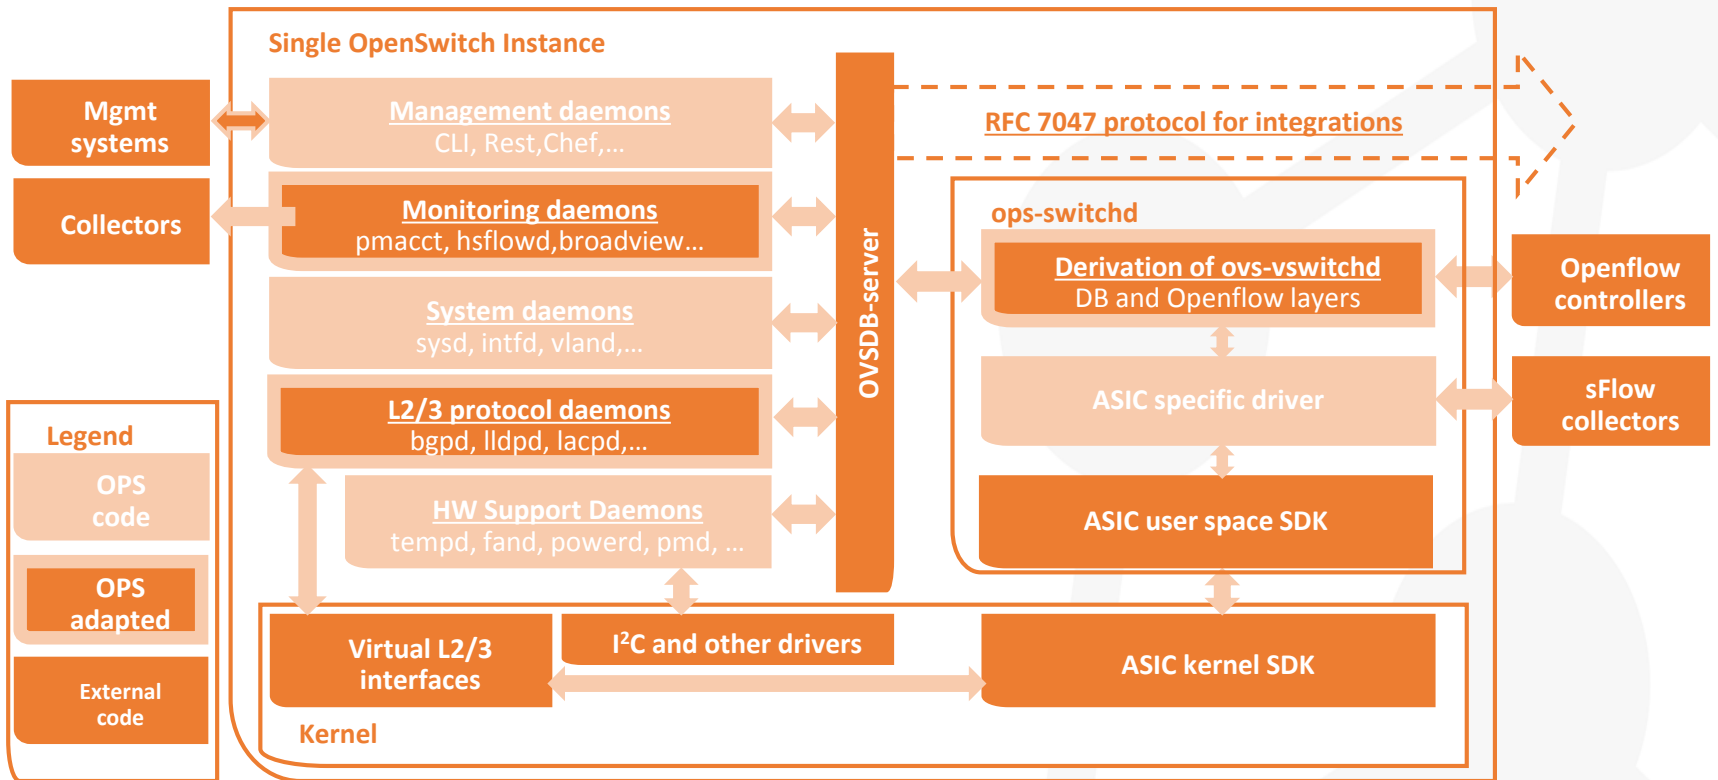


Corporate Participants





Architectural View





OPEN NETWORKING
SUMMIT 2016
MARCH 14-17, 2016 | SANTA CLARA, CA



Current Status

- Almost all the basics are in place – vlans, lags, routing, ECMP, BGP, OSPF, sFlow, NTP, DHCP, SVI, CLI, REST etc.
- ACLs, QoS, mirroring, MSTP and many others are coming in the next couple of months.
- Broadcom Trident II is supported (Accton 5712/6712), Tomahawk (Accton 7712) is on the way.
- Broadcom contributes Broadview monitoring support and develops OpenFlow support in its plugin.
- Barefoot contributes P4 plugin and emulator which runs inside VM and Docker containers.
- Cavium enables OPS on XPliant platforms.
- Additional ASIC vendors are working to enable their platforms.
- Code is in Alpha quality - features are being rapidly added
- Most developers will switch from feature development to hardening by end of March, with the goal of having operational system in June/July timeframe.

[Home](#)

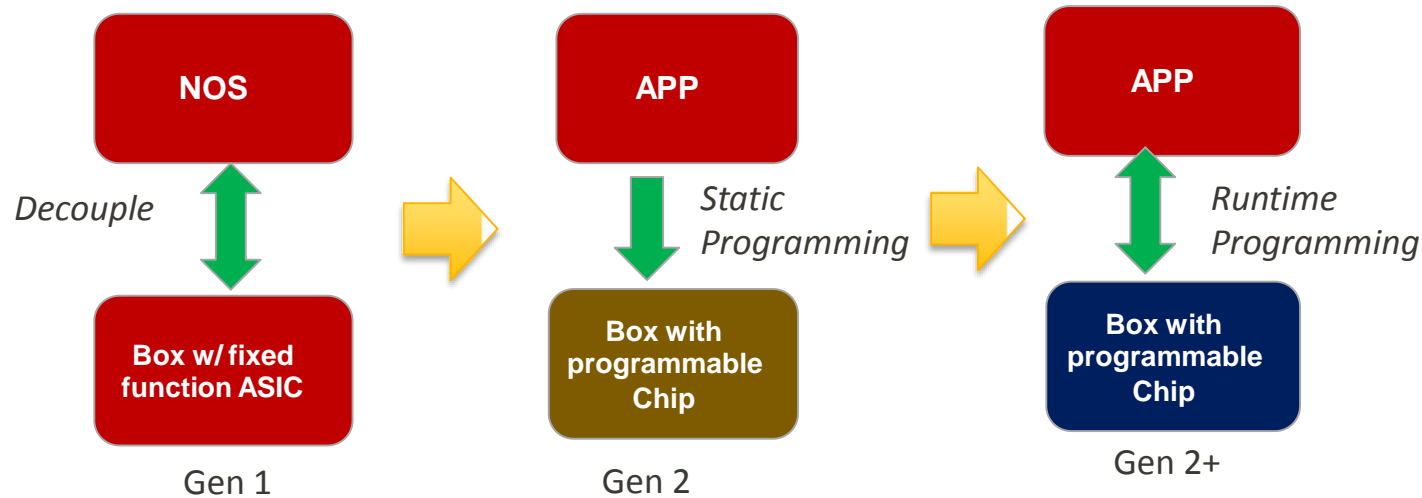


A Brief Introduction of Protocol Oblivious Forwarding (POF)

Dr. Haoyu Song
Huawei Technologies

[Home](#)

Towards Open Programmable Data Plane

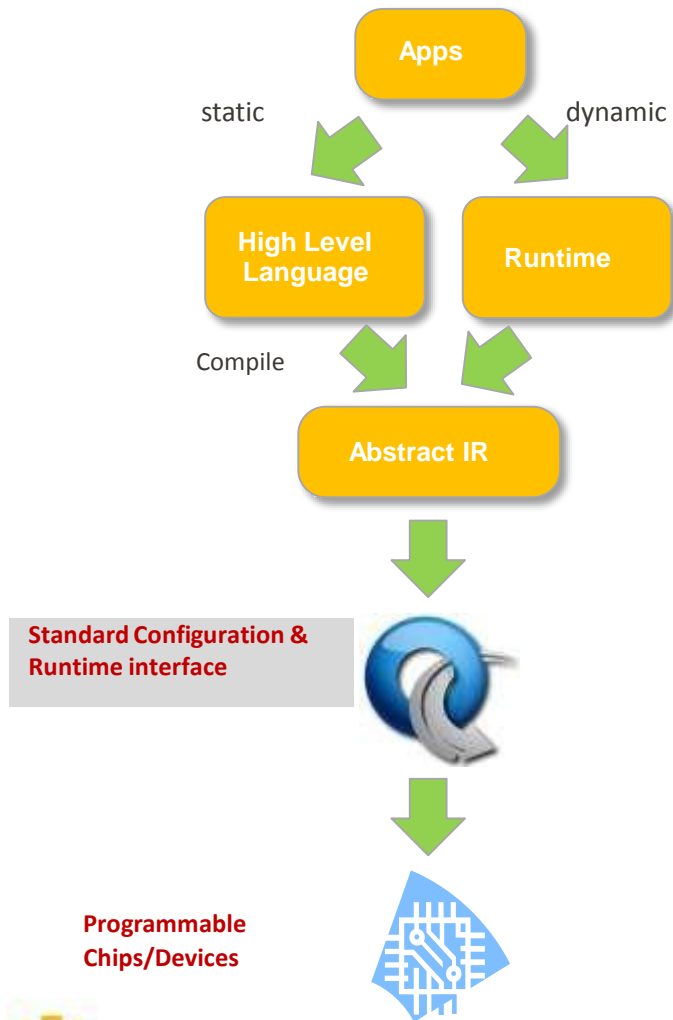


- Hide heterogeneity
- Application reuse

- Custom application
- One size fits all

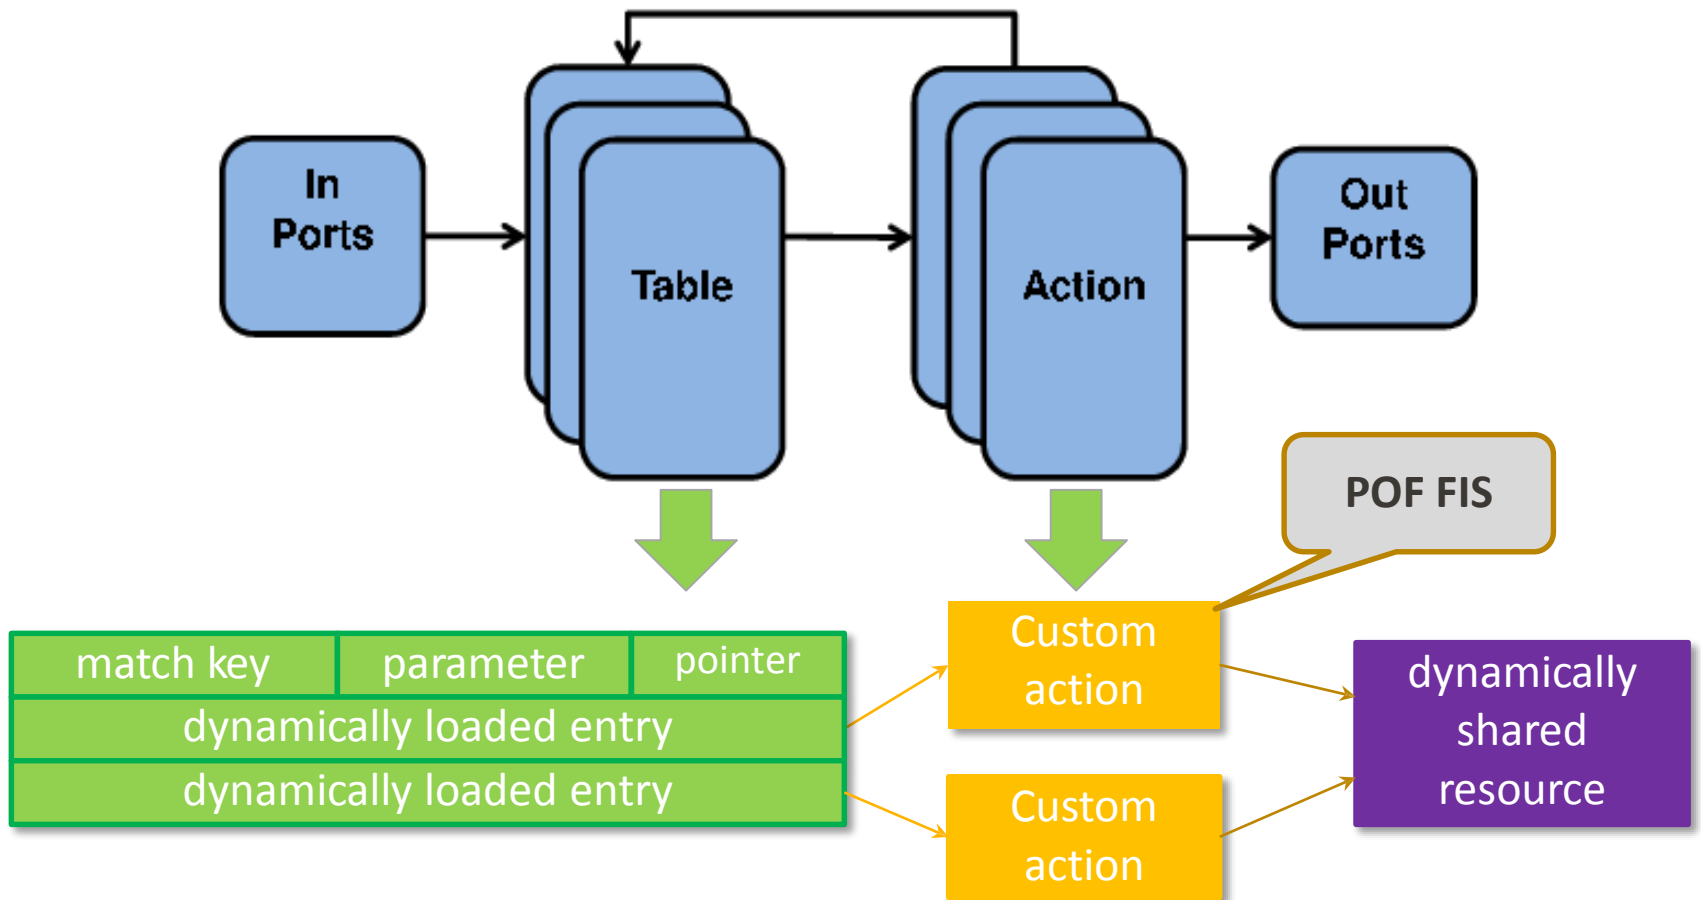
- Real time & on demand
- Interactive

POF Architecture & Characteristics

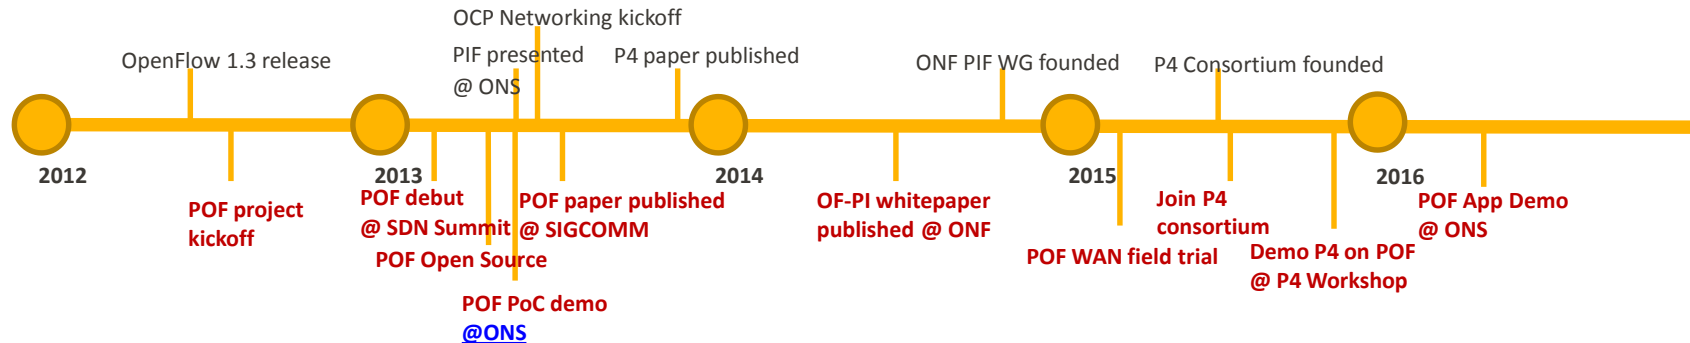


- Language Independent
- Protocol Independent
- Target Independent
- Interactive programming for real-time and on-demand applications

POF Data Plane Abstraction



POF Status Update



- Prototype: NE40E-based full-stack PoC
- ONF: Next gen OpenFlow & PIF
- Application: Dynamic network probes
- Next Step: Main-stream SDN controller integration

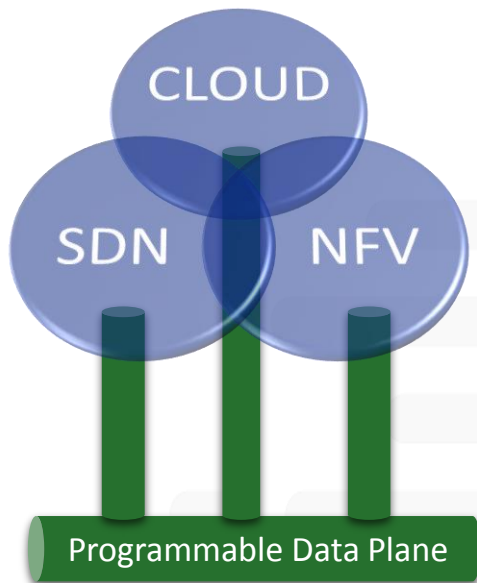


fd.io Intro

Jan Medved, Ed Warnicke

[Home](#)

Evolution of Programmable Networking



- Many industries are transitioning to a more dynamic model to deliver network services
- The great unsolved problem is how to deliver network services in this more dynamic environment
- Inordinate attention has been focused on the non-local network control plane (controllers)
 - Necessary, but insufficient
- There is a giant gap in the capabilities that foster delivery of dynamic Data Plane Services

Issues/Limitations with Existing Data Plane Solutions

- Known issues with Performance, Scalability & Stability
- Overly Complex Architectures
 - Hard to evolve
 - Slow rate of innovation
 - Steep learning curve
- Hard to deploy/upgrade/operate
 - slow cycles, too many kernel dependencies
- Lack of :
 - automated end-to-end system testing frameworks
 - leads to unpredictable system behavior
 - support for diverse/custom hardware
 - portability across compute platforms
 - optimal use of compute microarchitectures
 - network level instrumentation
 - Few debugability features
 - Few if any Statistics/Counters exposed



Introducing Fast Data: fd.io



fd.io Charter

- New project in Linux Foundation

- Multi-party
- Multi-project

- What does multi-party mean?

- Multiple members - Open to all

- What does multi-project mean?

- Multiple subprojects
- Subproject autonomy
- Cross project synergy
- Open to new subprojects
- Anyone can propose a subproject
- Allows for innovation

Create a Platform that enables Data Plane Services that are:

- Highly performant
- Modular and extensible
- Open source
- Interoperable
- Multi-Vendor

Platform fosters innovation and synergistic interoperability between Data Plane Services

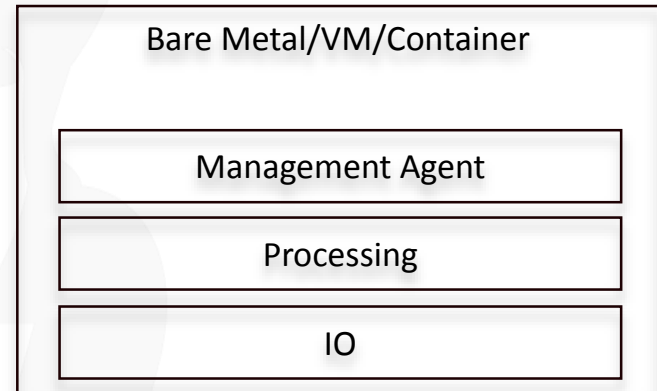
Source of Continuous Integration resources for Data Plane services based on the Consortium's project/subprojects

Meet the functionality needs of developers, deployers, datacenter operators

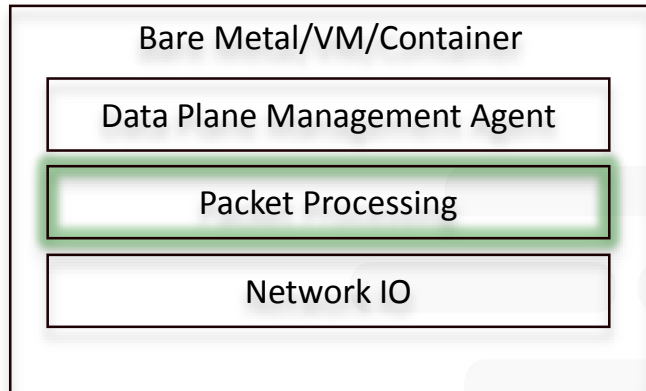


Fast Data Scope

- Fast Data Scope:
 - **IO**
 - Hardware/vHardware <-> cores/threads
 - **Processing**
 - Classify
 - Transform
 - Prioritize
 - Forward
 - Terminate
 - **Management Agents**
 - Control/manage IO/Processing



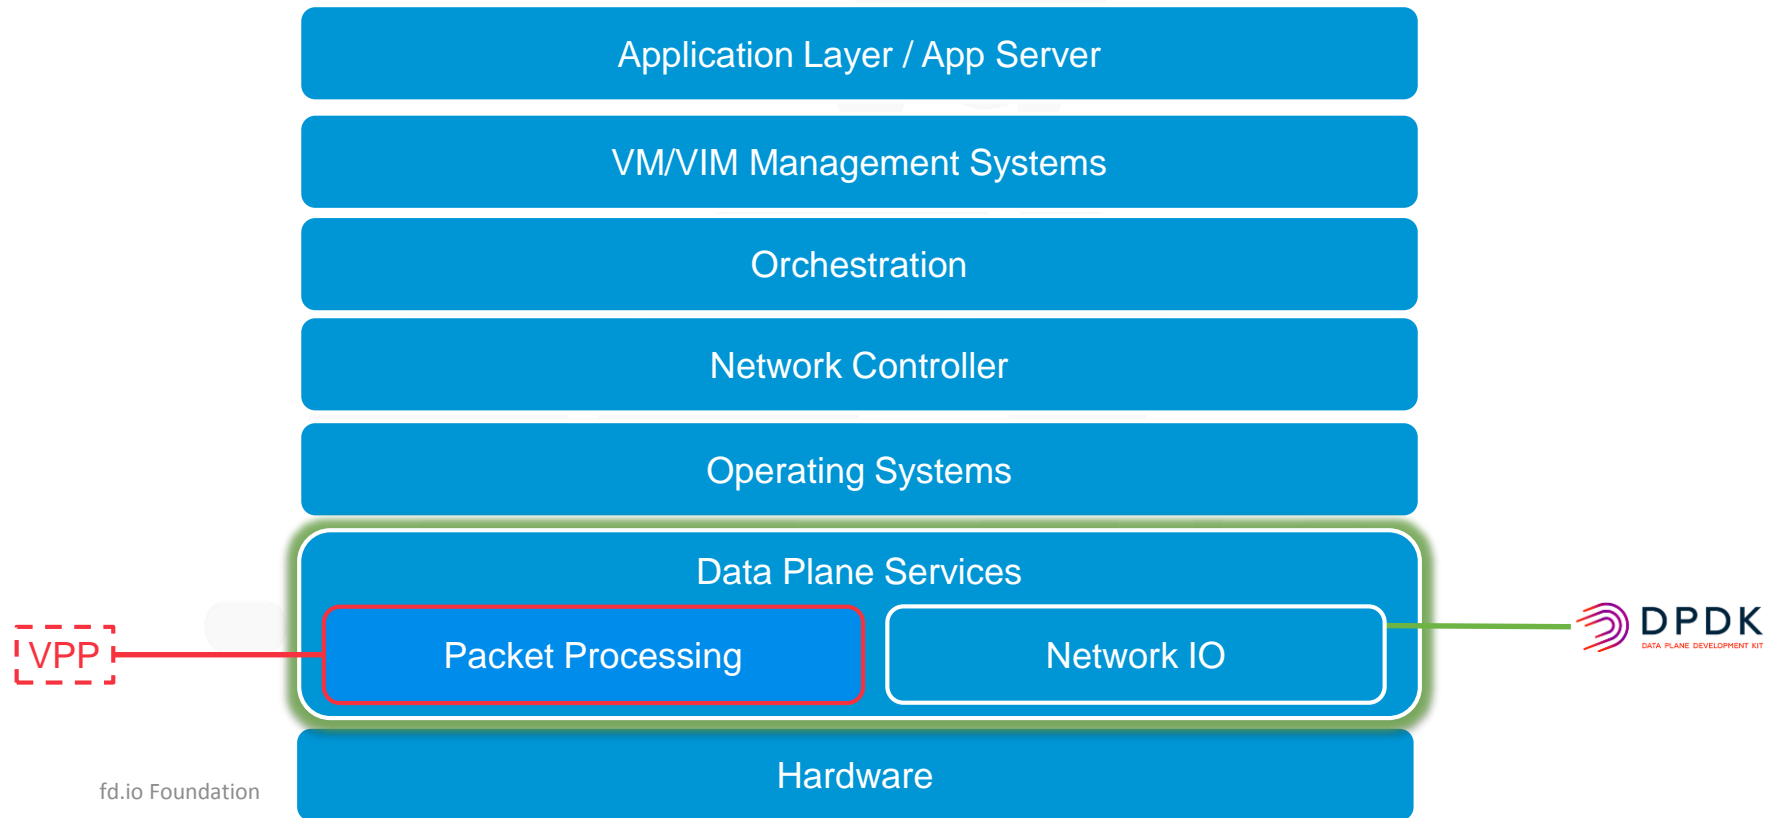
Introducing Vector Packet Processor - VPP



- VPP is a rapid packet processing development platform for highly performing network applications.
- It runs on commodity CPUs and leverages DPDK
- It creates a vector of packet indices and processes them using a directed graph of nodes – resulting in a highly performant solution.
- Runs as a Linux user-space application
- Ships as part of both embedded & server products, in volume
- Active development since 2002




VPP in the Overall Stack





Programming
Protocol-independent
Packet
Processors



How can you ...




Quickly deploy new protocols in your network?

See what your forwarding plane is doing?

Own your forwarding plane & intellectual property?

Tailor your network to meet their needs?

Write code that is portable across many devices in your networks?



Need a couple of things ...




1. Programmable Switch Chips

Giving end-users a programmable target for their forwarding plane

2. Industry-wide Programming Language

Allowing end-users to define and modify their forwarding plane



P4



High-level programming language

Networking domain specific

Protocol Independent


Abstract Forwarding Model

Growing adoption and ecosystem

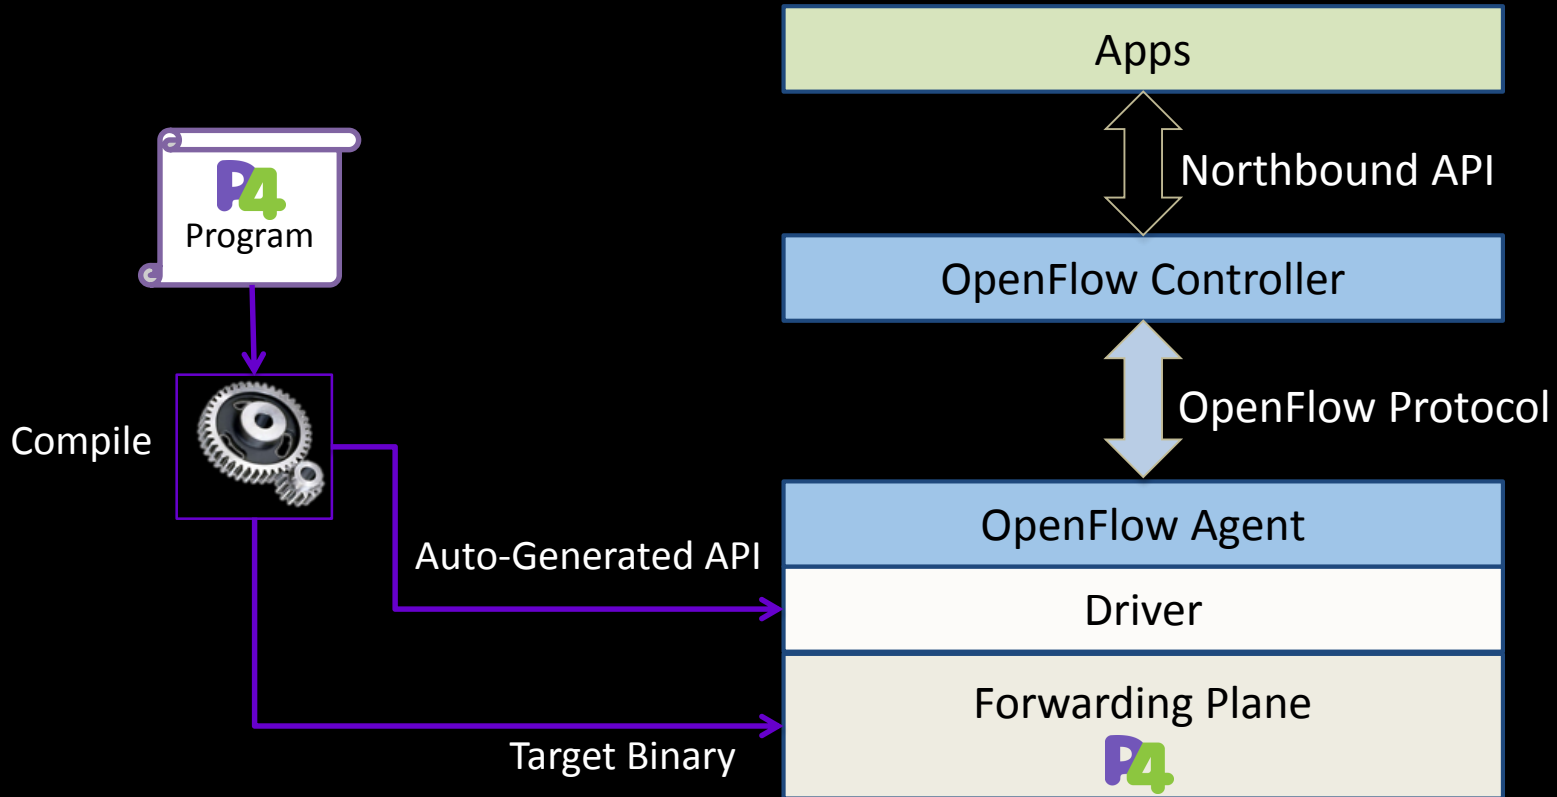
```
/* Router MAC lookup */
action rmac_hit() {
    modify_field(13_metadata.rmac_hit, TRUE);
}


action rmac_miss() {
    modify_field(13_metadata.rmac_hit, FALSE);
}

table rmac {
    reads {
        13_metadata.rmac_group : exact;
        12_metadata.lkp_mac_da : exact;
    }
    actions {
        rmac_hit;
        rmac_miss;
    }
    size : ROUTER_MAC_TABLE_SIZE;
}
```



P4 & OpenFlow





P4 Language Consortium (P4.org)



Independent CA non-profit


Membership is FREE

Open source

Apache 2.0 CLA

Website: <http://p4.org>

GitHub: <https://github.com/p4lang>



P4.org Membership




40 organizations from Industry and Academia and growing ...

Cloud Service Provider, Web 2.0, Telco Operator, OEM, ODM, NPU, FPGA, ASIC, Software, ...

<http://p4.org/join-us>





Developer Resources



Spec

<http://p4.org/spec/>

Compiler

<https://github.com/p4lang/p4c-bm>


L2/L3 Switch

<https://github.com/p4lang/switch>

Packet Test Framework (PTF)

<https://github.com/p4lang/ptf>

much more at <https://github.com/p4lang>



Now you can ...



Quickly add new protocols.

Remove unused protocols.


Flexibly assign memories to features.

Create new diagnostics, telemetry, OAM etc.

Compose forwarding behavior from libraries.

Specify forwarding behavior once; compile to many devices.

Keep your IP from going to your competitors!



Get involved!



Join P4.org - <http://p4.org/join-us>

Try P4 development tools and programs (switch.p4)

Including P4-programmable S/W switches and test framework

Exciting apps for network monitoring, analysis, diagnostics, and control

Join us at the following events

OpenSwitch Booth at ONS March 15-17, 2016

P4 Introduction & Tutorial on May 23, 2016

P4 Workshop 2016 on May 24, 2016

Join the mailing lists - <http://lists.p4.org>