VSPERF and Open vSwitch for OPNFV

Mark Gray, Thomas Herbert and Maryam Tahhan.

LINUX FOUNDATION
COLLABORATIVE PROJECTS

What is OPNFV?

OPNFV is a carrier-grade, integrated, open source platform to accelerate the introduction of new NFV products and services.

OPNFV

OPNFV Membership List

Platinum Members

at&t BROCADE 中国移动 China Mobile CISCO DELL NTT docomo EMC² ERICSSON hp HUAWEI IBM

intel JUNIPER NETWORKS NEC NOKIA redhat TELECOM ITALIA vodafone ZTE中兴

Silver Members

6WIND ADVA Optical Networking Alcatel·Lucent ARM Array NETWORKS BROADCOM CableLabs ubuntu Supported by Canonical CAVIUM CenturyLink ciena CITRIX

ClearPath networks CONTEXTREAM Coriant cyan Dialogic dorado SOFTWARE ENEA H3C ixia kt Metaswitch Networks midokura MIRANTIS

ooredoo orange OVERTURE QOSMOS sandvine SK telecom SPIRENT Sonus Sprint Stratus Technologies vmware WIND RIVER XILINX

OPNFV

What does OPNFV do?

OPNFV

OPNFV Platform Overview

OPNFV is a carrier-grade, integrated, open source reference platform

Work with upstream projects to coordinate continuous integration and testing

Fill development gaps

Establish an ecosystem for NFV solutions based on open standards and software

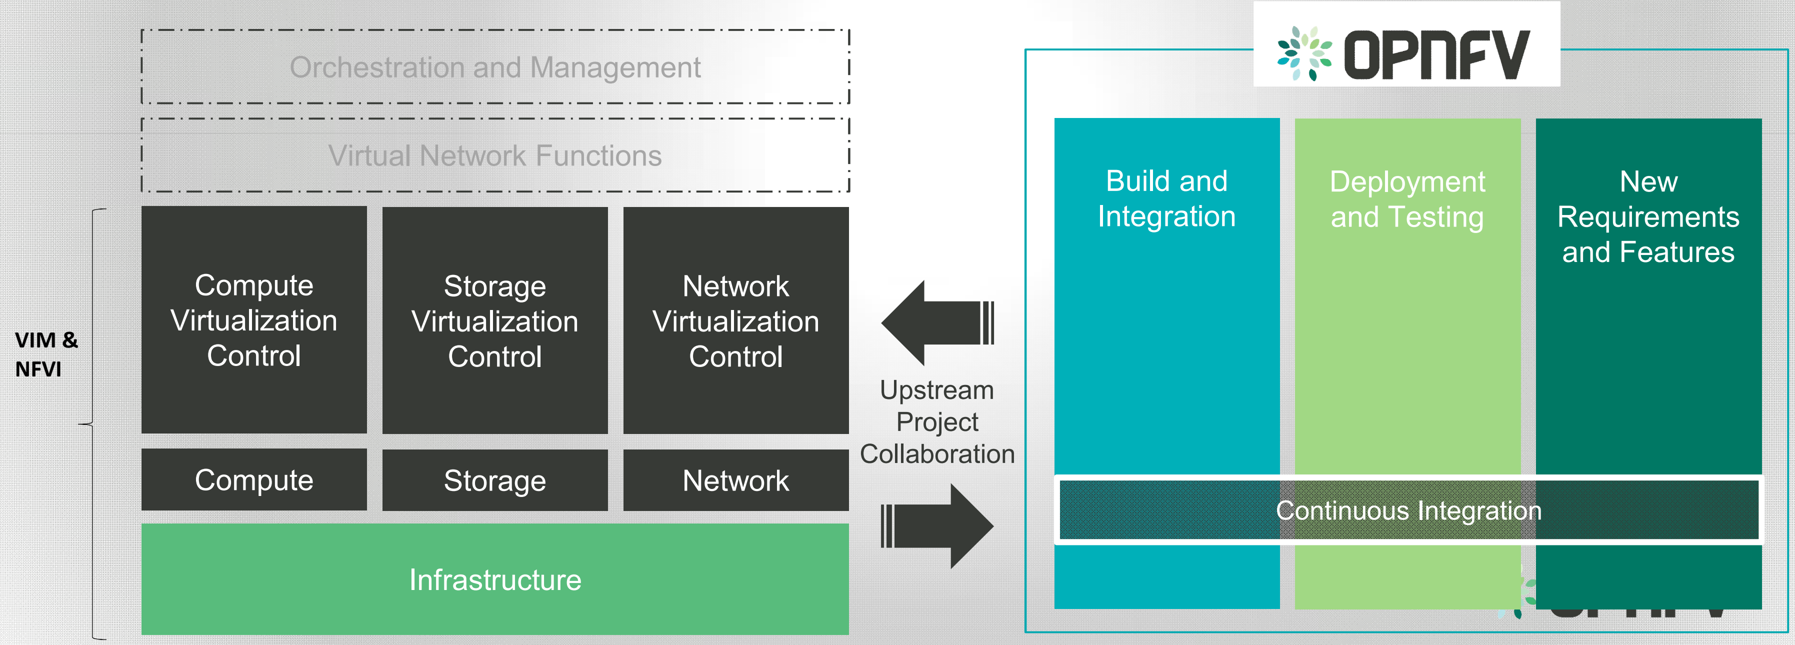

How does it do it?

OPNFV

Open Source Community Working with Upstream Communities

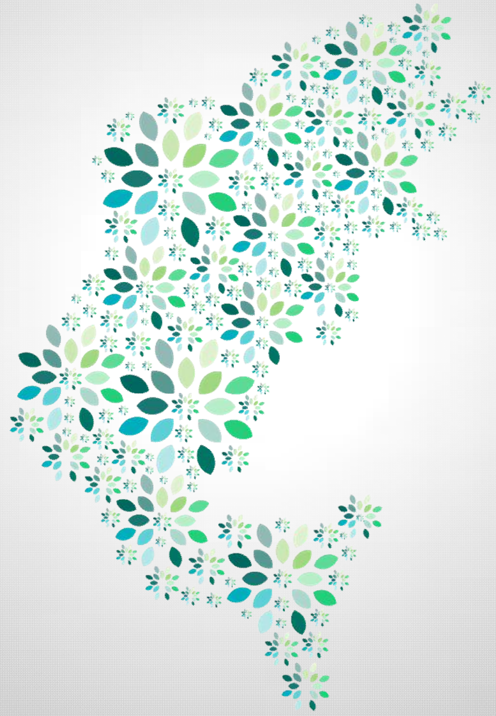

OPNFV

OPNFV Releases

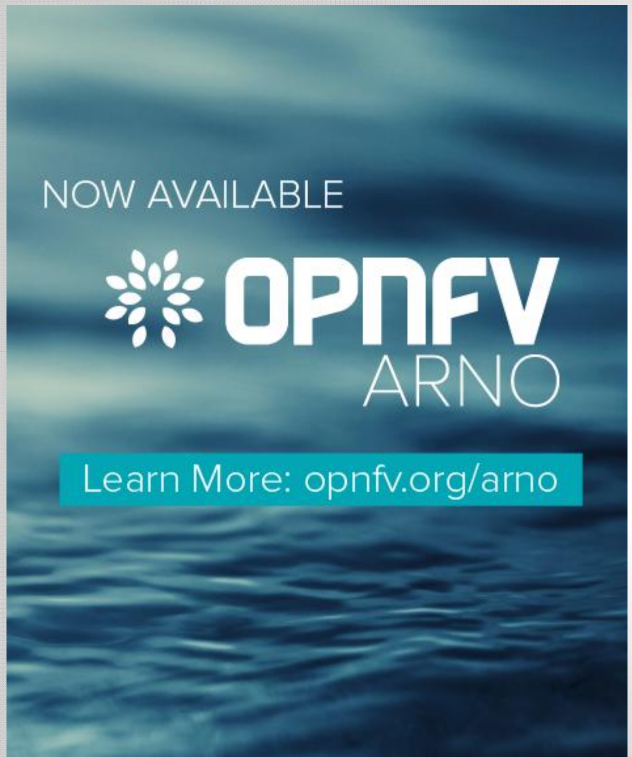

Brahmaputra Release

Coming Soon

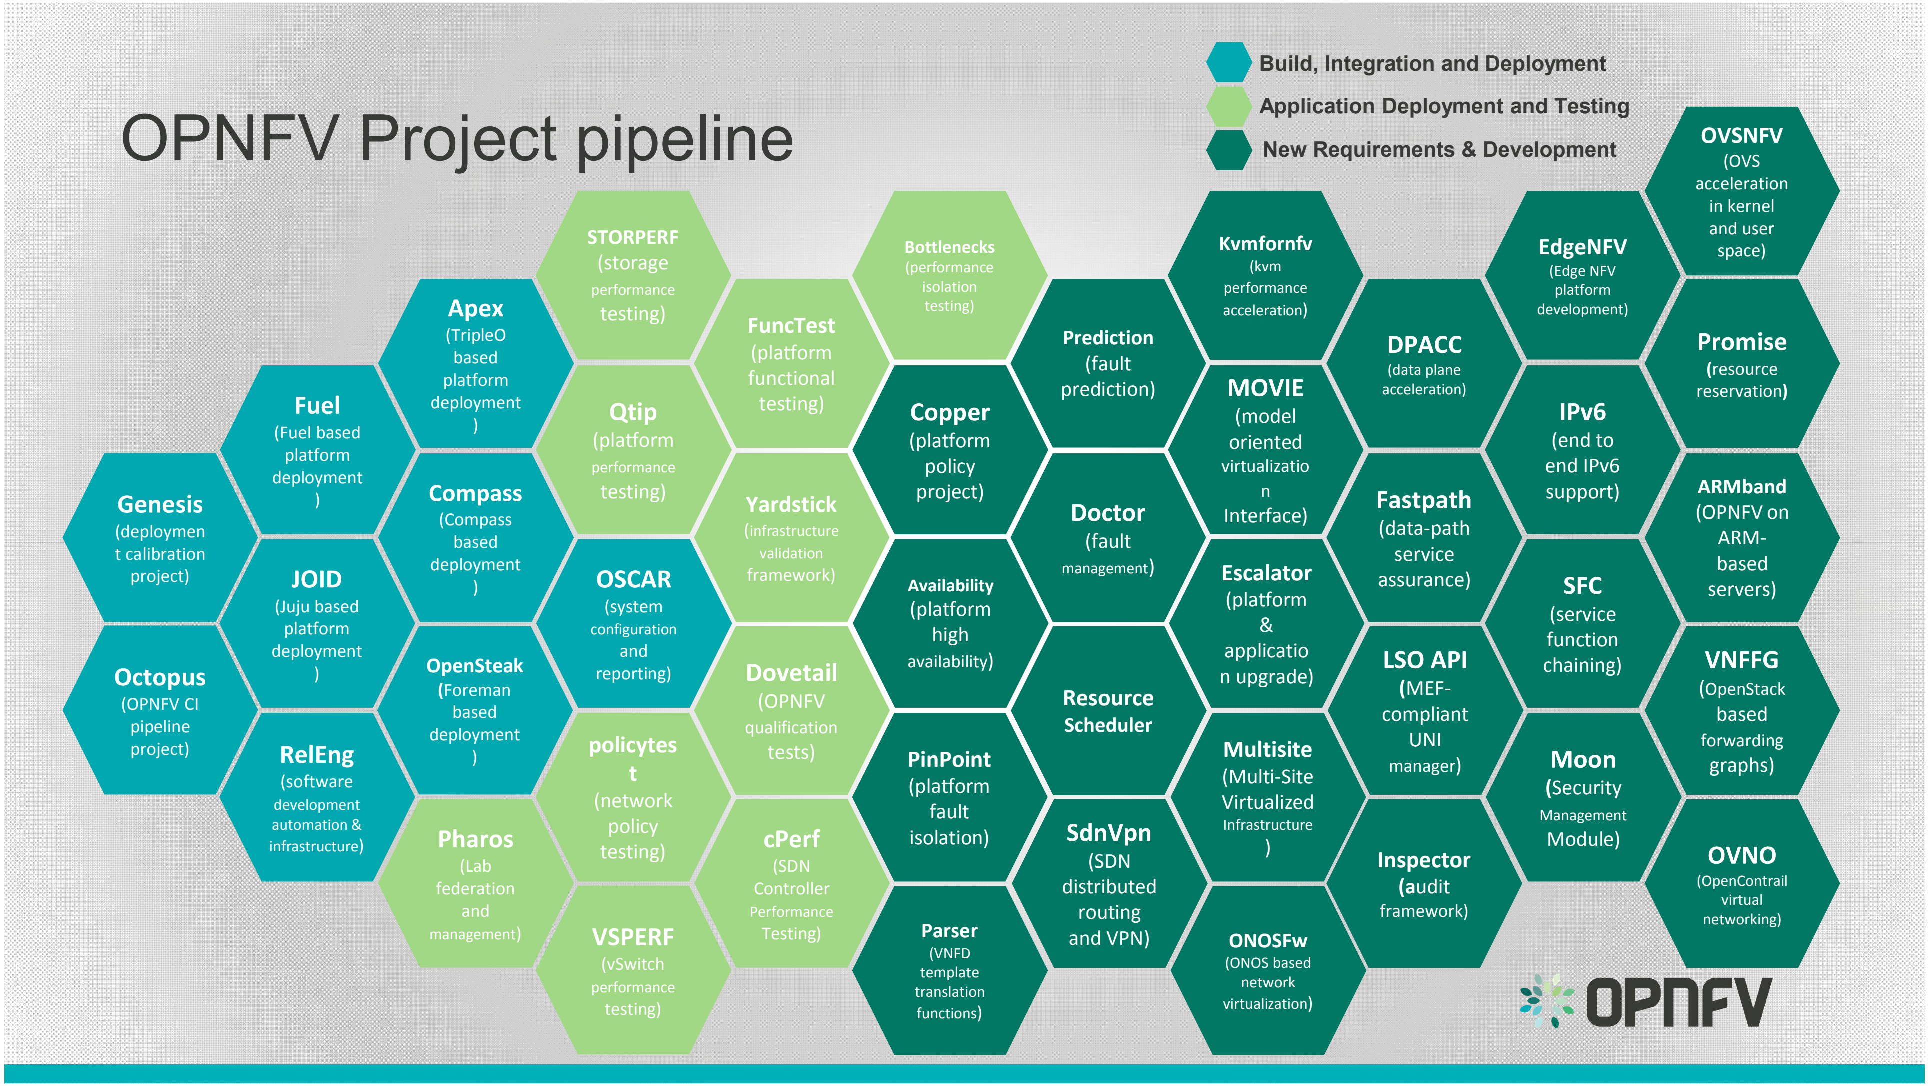
OPNFV Project pipeline
Build, Integration and Deployment
Application Deployment and Testing
New Requirements & Development
Genesis (deploymen t calibration project)
Octopus (OPNFV CI pipeline project)
Fuel (Fuel based platform deployment)
JOID (Juju based platform deployment)
RelEng (software development automation & infrastructure)
Apex (TripleO based platform deployment)
Compass (Compass based deployment)
OpenSteak (Foreman based deployment)
Pharos (Lab federation and management)
STORPERF (storage performance testing)
Qtip (platform performance testing)
OSCAR (system configuration and reporting)
policytes t (network policy testing)
VSPERF (vSwitch performance testing)
FuncTest (platform functional testing)
Yardstick (infrastructure validation framework)
Dovetail (OPNFV qualification tests)
cPerf (SDN Controller Performance Testing)
Bottlenecks (performance isolation testing)
Copper (platform policy project)
Availability (platform high availability)
PinPoint (platform fault isolation)
Parser (VNFD template translation functions)
Prediction (fault prediction)
Doctor (fault management)
Resource Scheduler
SdnVpn (SDN distributed routing and VPN)
Kvmfornfv (kvm performance acceleration)
MOVIE (model oriented virtualizatio n Interface)
Escalator (platform & applicatio n upgrade)
Multisite (Multi-Site Virtualized Infrastructure)
ONOSFw (ONOS based network virtualization)
DPACC (data plane acceleration)
Fastpath (data-path service assurance)
LSO API (MEF-compliant UNI manager)
Inspector (audit framework)
EdgeNFV (Edge NFV platform development)
IPv6 (end to end IPv6 support)
SFC (service function chaining)
Moon (Security Management Module)
OVSNFV (OVS acceleration in kernel and user space)
Promise (resource reservation)
ARMband (OPNFV on ARM-based servers)
VNFFG (OpenStack based forwarding graphs)
OVNO (OpenContrail virtual networking)
OPNFV

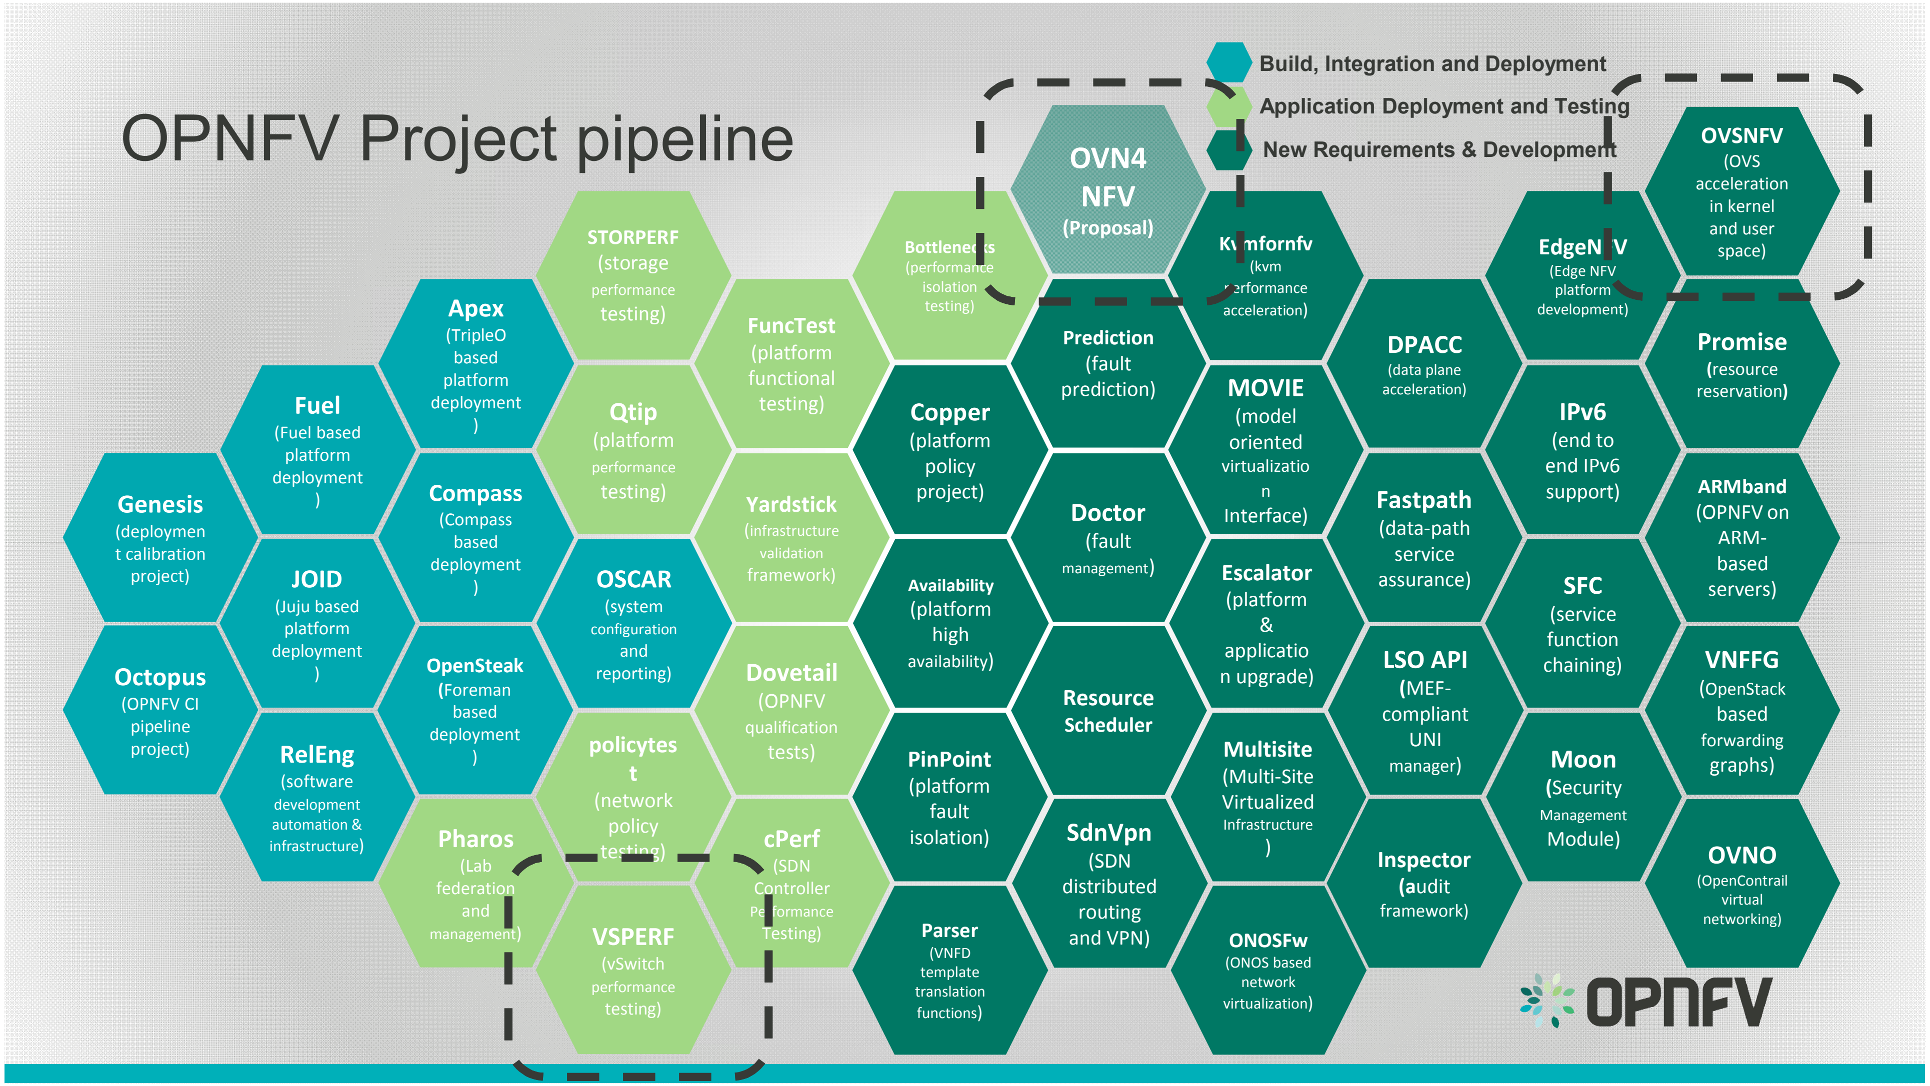

OPNFV Project pipeline
Build, Integration and Deployment
Application Deployment and Testing
New Requirements & Development
Genesis (deployment calibration project)
Octopus (OPNFV CI pipeline project)
Fuel (Fuel based platform deployment)
JOID (Juju based platform deployment)
RelEng (software development automation & infrastructure)
Apex (TripleO based platform deployment)
Compass (Compass based deployment)
OpenSteak (Foreman based deployment)
Pharos (Lab federation and management)
STORPERF (storage performance testing)
Qtip (platform performance testing)
OSCAR (system configuration and reporting)
policytest (network policy testing)
VSPERF (vSwitch performance testing)
FuncTest (platform functional testing)
Yardstick (infrastructure validation framework)
Dovetail (OPNFV qualification tests)
cPerf (SDN Controller Performance Testing)
Bottlenecks (performance isolation testing)
Copper (platform policy project)
Availability (platform high availability)
PinPoint (platform fault isolation)
Parser (VNFD template translation functions)
OVN4 NFV (Proposal)
Prediction (fault prediction)
Doctor (fault management)
Resource Scheduler
SdnVpn (SDN distributed routing and VPN)
Kvmfornfv (kvm performance acceleration)
MOVIE (model oriented virtualization Interface)
Escalator (platform & application upgrade)
Multisite (Multi-Site Virtualized Infrastructure)
ONOSFw (ONOS based network virtualization)
DPACC (data plane acceleration)
Fastpath (data-path service assurance)
LSO API (MEF-compliant UNI manager)
Inspector (audit framework)
EdgeNFV (Edge NFV platform development)
IPv6 (end to end IPv6 support)
SFC (service function chaining)
Moon (Security Management Module)
OVSNFV (OVS acceleration in kernel and user space)
Promise (resource reservation)
ARMband (OPNFV on ARM-based servers)
VNFFG (OpenStack based forwarding graphs)
OVNO (OpenContrail virtual networking)
OPNFV

Open vSwitch for NFV

OVS and Independent Data Plane

- OVS Architecture Supports Independent Data Planes
  - DPDK
  - Linux Kernel Data Plane
- OVS with Accelerated Data Plane
  - OVS with DPDK
  - Currently the Most Widely Adopted
  - The Most Promise for the Near Future

Open vSwitch Architecture and DPDK

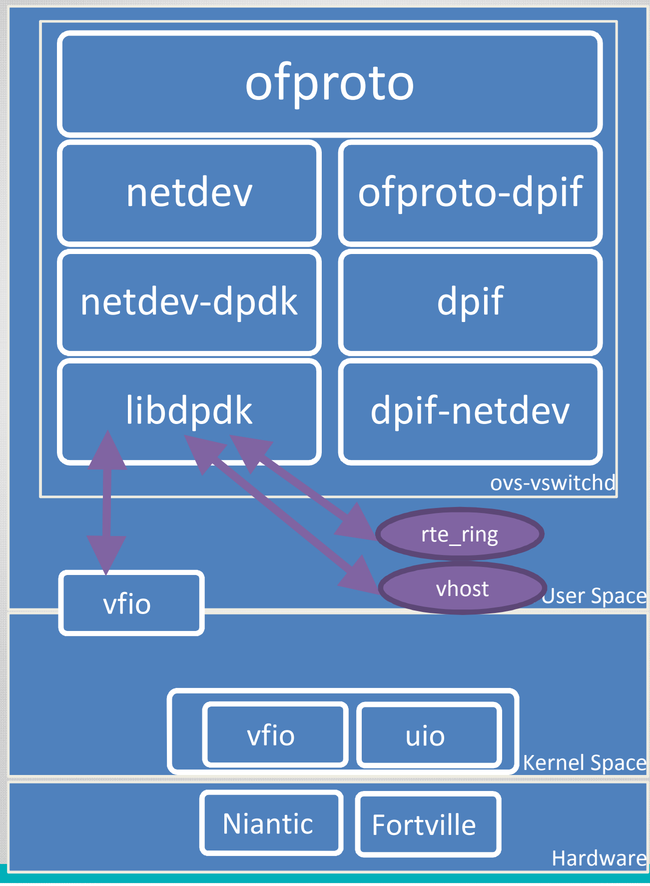

DPDK – Open vSwitch

- DPDK – Data Plane Development Kit
  - About 4 Years Old
  - First Integrated with OVS from 2013
  - Fast Packet Forwarding
  - Poll Mode Drivers
  - Uses Commodity Hardware
  - Multiple Threads and Cores
- Up to 12X Speed Improvement for small packets
- Over 15mpps Forwarding
  - Small Packets

- Disadvantages WRT Linux Kernel
- Linux Data Plane Has
  - Complete TCP/IP Stack
  - 20 years of development
  - Rich Debugging Options
    - Promiscuous IFs
  - Access to Wide Variety of Network IF's and VF's
  - Tunnels and Endpoints

OPNFV

OVSNFV – Phase 1 (Build, Integrate, Deploy, Test)

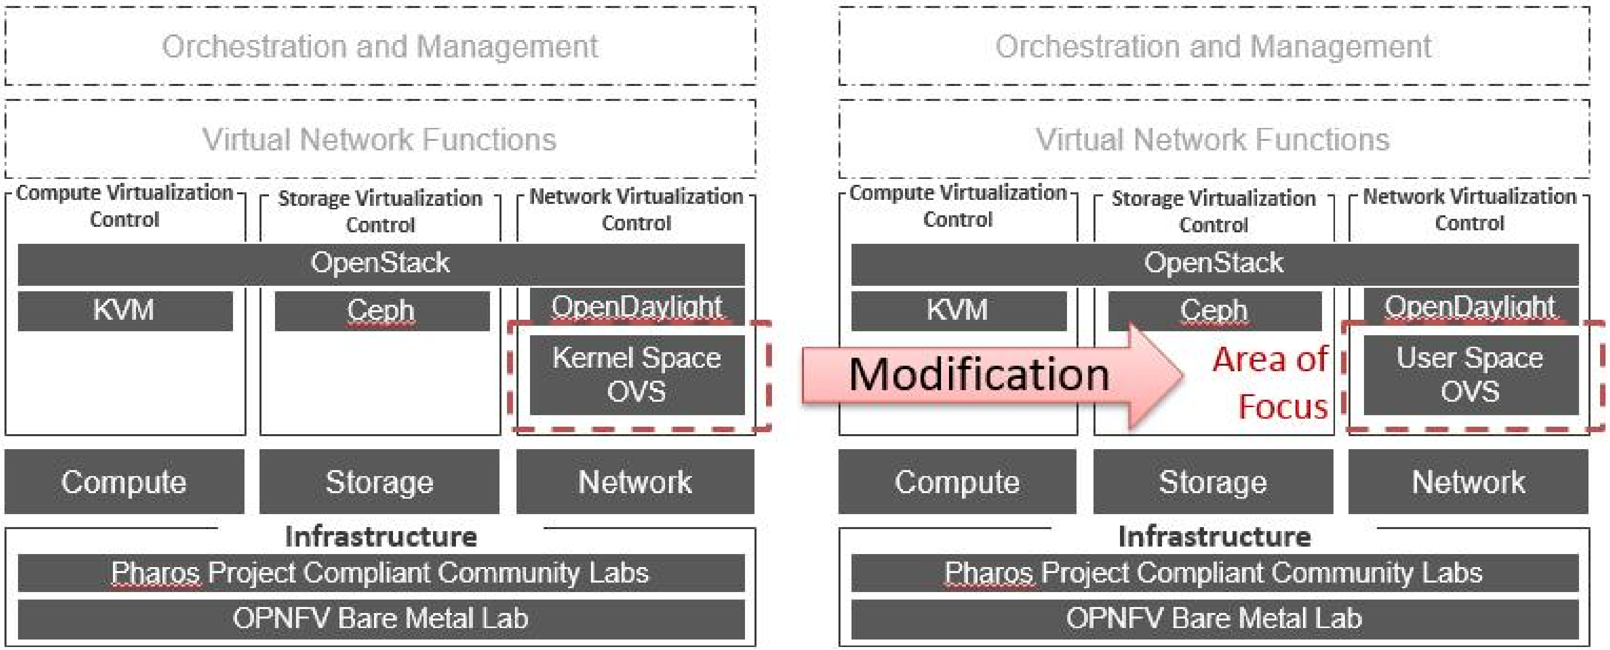

OPNFV

OVSNFV Project

- Collaborative Development
  - Incubation Stage
- Overall goal:
  - provide Open vSwitch with user space accelerated data plane for deployment within the OPNFV ecosystem.
- Take OVS and DPDK from the upstream projects
- Deploy OVS/DPDK as Package for use by
  - VSPERF
  - SFC
  - General Use as Deployed OVS in OPNFV
- Test and Verify Assumption of DPDK Use Case in OPNFV
- Provide Alternative OVS-Linux Kernel for Comparison

OVSNFV Project

- Project Wiki Page
  - https://wiki.opnfv.org/ovsnfv
- Project Lead
  - Mark Gray (Intel)
- Committers
  - Mark Gray (Intel)
  - Joseph Gasparakis (Intel)
  - Billy O Mahony (Intel)
  - Hongbo Tianhongbo (Huawei)
  - Thomas F Herbert (Red Hat)

OPNFV

OVSNFV Project

- Fed by Two Upstream Projects
  - Open vSwitch
  - DPDK
- We are NOT Forking Either DPDK or OVS
- Strive For Upstream Enablement for Easier OPNFV Integration
  - Upstream: Maintain “Similar” Semantics for Both
- Although We May Use Patches before They Are Merged Upstream
  - To Support Specific Required Use Cases

OVSNFV Project – Upstream Issues

- DPDK Device Management
  - Driverctl Utility Preferably with Systemd patch
  - http://dpdk.org/ml/archives/dev/2015-November/028121.html
- NSH patch from Intel (Danny Zhou)
- “Alternate” RPM Install
  - Separate Glance Images for Ironic Compute Node Install
- ML2 Mechanism Driver Update for DPDK/OVS
- OpenStack ODL change to add Vhost-User Port Names
  - https://review.openstack.org/#/c/215612/

OPNFV

OVSNFV – Looking Forward

- Discover Requirements and Needed Features
- Deployment of OVS/DPDK in OPNFV

- Get Feedback from OPNFV Ecosystem
  - Gather Missing Required Features

- Push Issues Upstream to Improve
  - DPDK
  - OVS

- Merged into DPDK and OVS

OPNFV

OVSNFV – Looking Forward

- Discover Requirements and Needed Features
- Deployment of OVS/DPDK in OPNFV
- Get Feedback from OPNFV Ecosystem
  - Gather Missing Required Features
- Push Issues Upstream to Improve
  - DPDK
  - OVS
- Merged into DPDK and OVS

VSPERF Overview

Define, implement and execute an appropriate set of tests in order to objectively measure the current Telco characteristics of a virtual switch in the NFVI

Drive standardization

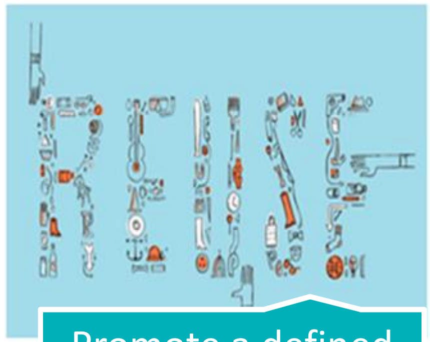

Promote a defined platform and reuse

Establish best practice

30+ committers and contributors

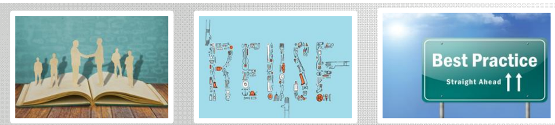

VSPERF Standardization and Open Source Projects

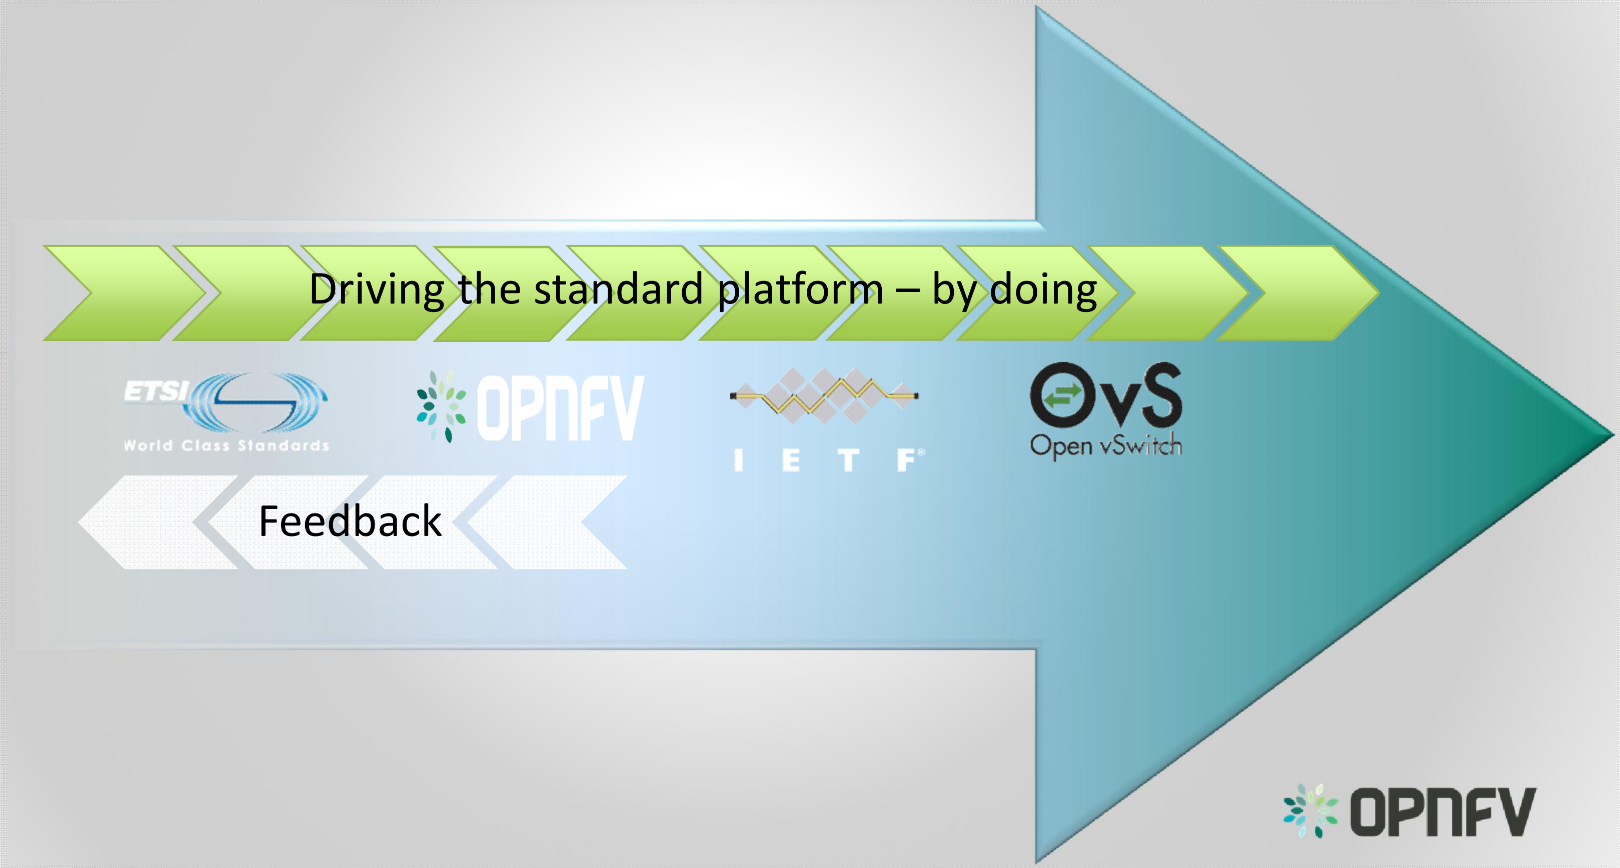

OPNFV

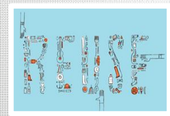

VSPERF Deliverables

IETF Draft

[Docs] [txt|pdf|xml] [Tracker] [Email] [Diff1] [Diff2] [Nits]

Versions: 00 01

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 16, 2016

M. Tahhan
B. O'Mahony
Intel
A. Morton
AT&T Labs

October 14, 2015

Benchmarking Virtual Switches in OPNFV
draft-vsperf-bmwg-vswitch-opnfv-01

Abstract

This memo describes the progress of the Open Platform for NFV (OPNFV) project on virtual switch performance "VSWITCHPERF". This project intends to build on the current and completed work of the Benchmarking Methodology Working Group in IETF, by referencing existing literature. The Benchmarking Methodology Working Group has traditionally conducted laboratory characterization of dedicated physical implementations of internetworking functions. Therefore, this memo begins to describe the additional considerations when virtual switches are implemented in general-purpose hardware. The expanded tests and benchmarks are also influenced by the OPNFV mission to support virtualization of the "telco" infrastructure.

Test Specification

Table Of Contents

1. CHARACTERIZE VSWITCH PERFORMANCE FOR TELCO NFV USE CASES LEVEL TEST DESIGN

Table of Contents

Modular Test Framework

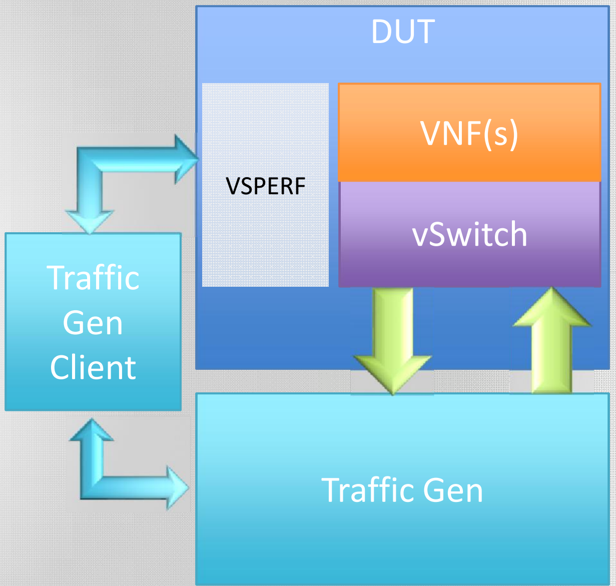

Consumable by: IETF

Consumable by: OPNFV, Open vSwitch

VSPERF 3x4 Matrix LTD Coverage

	SPEED	ACCURACY	RELIABILITY	SCALE
Activation	• RFC2889. AddressLearningRate • RFC2889. AddressCachingCapacity • InitialPacketProcessingLatency • LatencyAndLatencyVariation	• CPDP.Coupling.Flow.Addition	• RFC2544.SystemRecoveryTime • RFC2544.ResetTime	• RFC2889.AddressCachingCapacity
Operation	• RFC2544.PacketLossRatio • RFC2544.PacketLossRateFrmMod • RFC2544.BackToBackFrames • RFC2889.MaxForwardingRate • RFC2889.ForwardPressure • RFC2889.BroadcastFrameForwarding • RFC2889 Broadcast Frame Latency test • CPU.RFC2544.0PacketLoss • RFC2544.WorstN-BestN • InterPacketDelayVariation.RFC5481	• RFC2889.ErrorFramesFiltering • RFC2544.Profile	• RFC2889.Soak • RFC2889.SoakFrameModification • PacketDelayVariation.RFC3393.Soak	• Scalability.RFC2544.0PacketLoss • MemoryBandwidth.RFC2544.0PacketLoss.Scalability
De-Activation				

RFC2544 Benchmarking Methodology for Network Interconnect Devices
RFC2889 Benchmarking Methodology for LAN switching Devices

OPNFV

Future Work

- Integrating multiple traffic gens: IXIA, Spirent, Moongen and Xena.
- Methodology extensions: iterations for the short trial tests.
- Prove out and refine methodology and tests through the framework.
- Add more tests to the LTD and the framework.
- Continuous Integration support.

OVS call to action

- So join us in OPNFV to help establish an Open Source, carrier grade, integrated platform that includes a carrier grade OVS.

OPNFV

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: www.intel.com/design/literature.htm

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families: Go to: Learn About Intel® Processor Numbers

Intel, the Intel logo, Itanium, Intel Atom, Intel Xeon Phi, Intel AppUp, and Xeon 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.

OPNFV