
redhat.
OpenStack, OpenDaylight, and OPNFV
Chris Wright
Chief Technologist
Red Hat
Feb 3, 2015

Agenda

- Introduction
- Open Source NFV
- OPNFV
- OpenDaylight
- OVS+DPDK
- Red Hat and NFV

Who am I ?

- Chris Wright
- Chief Technologist at Red Hat
- Board Member of OpenDaylight and OPNFV Project
- During nearly 20 years as a software engineer he has worked in the telecom industry on high availability and distributed systems and in the Linux industry on security and virtualization.

Community and Red Hat

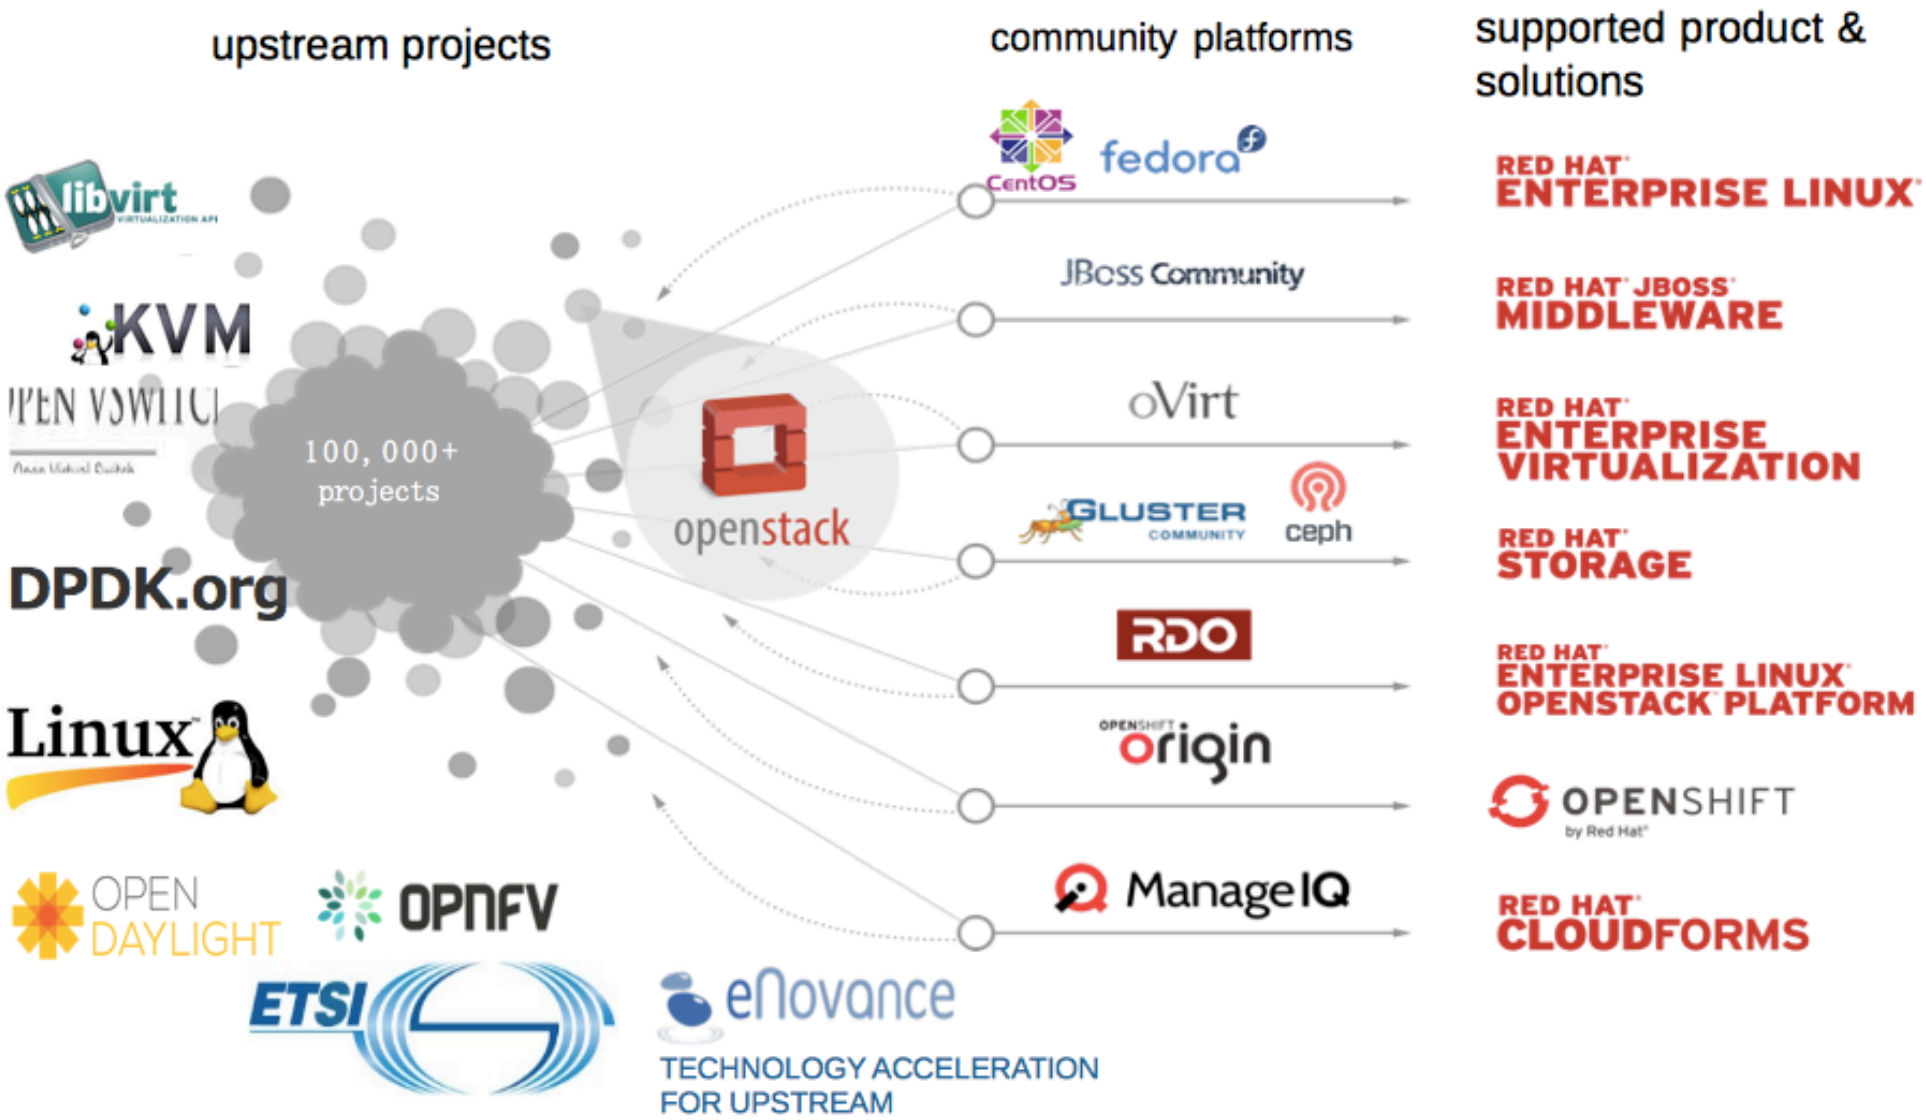

Open Source NFV

Network Functions Virtualization (NFV)

- Network Functions are trapped in function specific HW
  - Expensive to procure, integrate, deploy and operate
- Virtualize Network Functions (VNFs)
- Distribute VNFs on COTS-based IaaS – a Cloud
  - Bring modern data center architecture to Telco network
- Steer traffic with SDN
  - Traffic must traverse chain of functions in well-defined order

Why NFV?

- Reduce time-to-market for new services
  - Improve business agility
  - Compete with over-the-top web-based services
- Reduce CAPEX and OPEX costs

Virtual Network Function as a Service

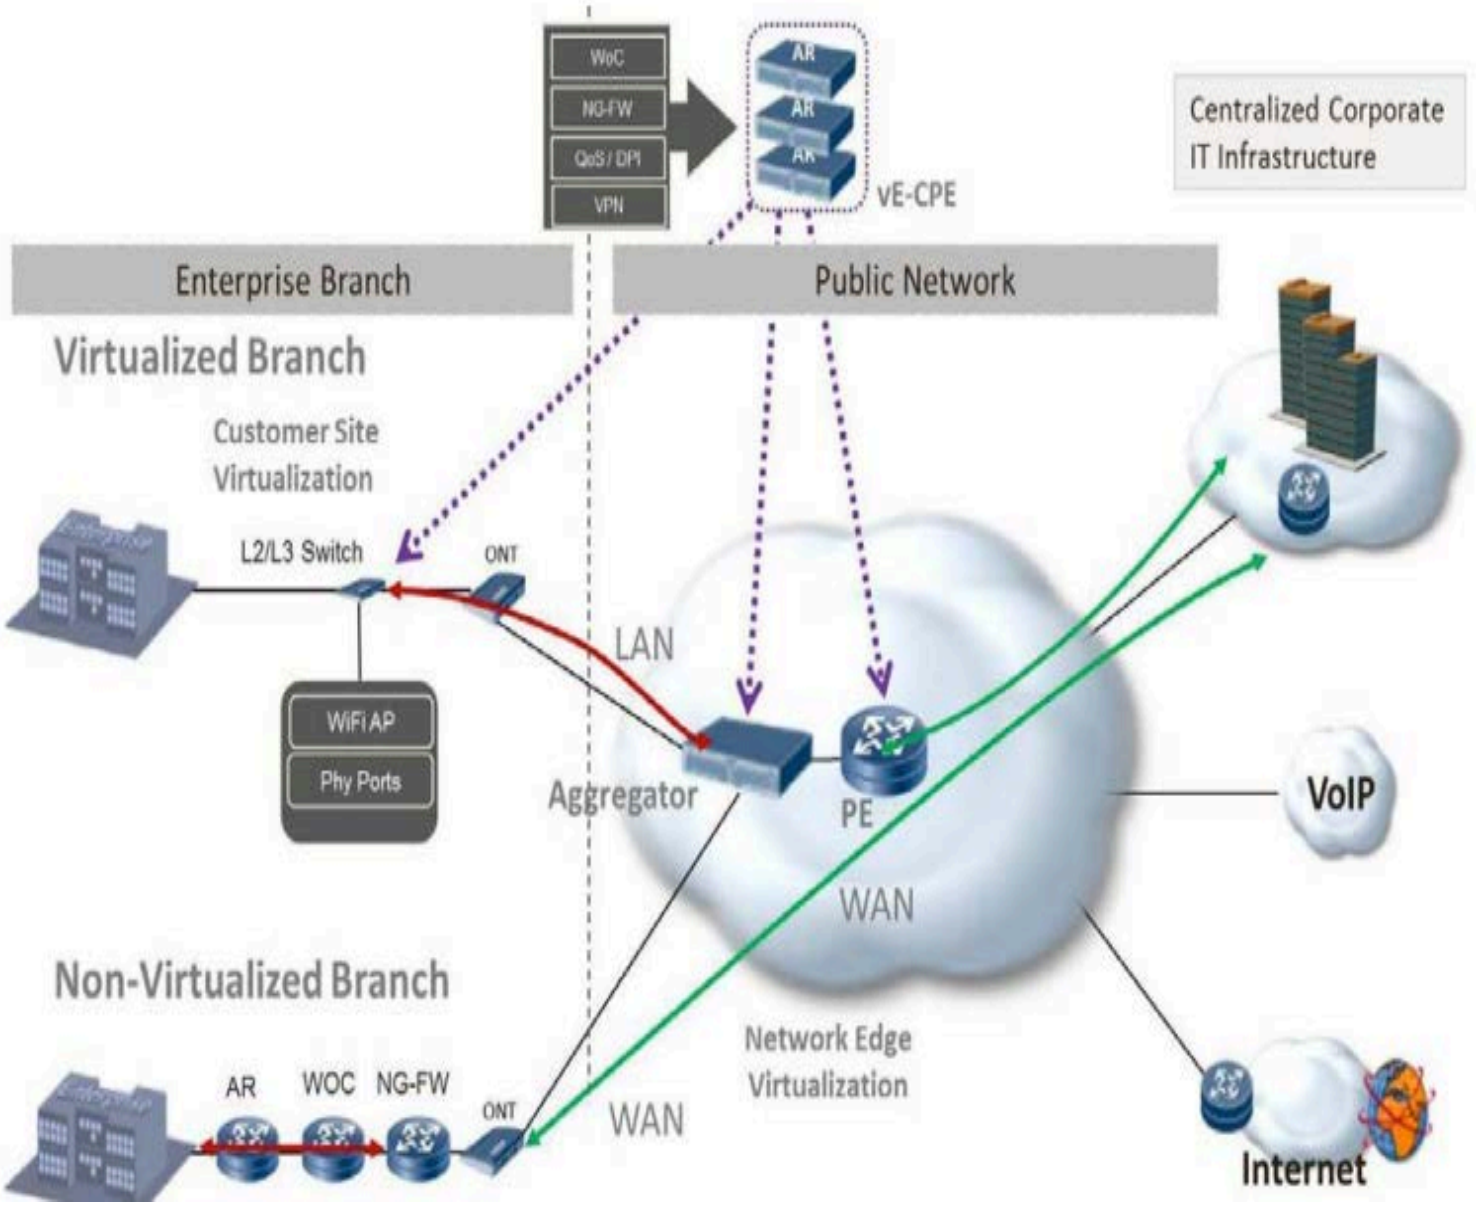

Network Functions

- Access Router
- WAN Optimization
- Firewall
- CG-NAT
- ADC¥IPS
- NG-Firewall
- DPI
- Monitoring QoS

(ETSI #2)

Example: VNFaaS

- Requires service chaining support in OpenDaylight, OpenStack Neutron
  - Enable Nova instances as nodes in a Neutron service chain
  - SFC support in OpenDaylight
- Performance
  - DPDK-accelerated Open vSwitch
- Reliability
  - HA instances in Nova

Virtualized Mobile Base Station

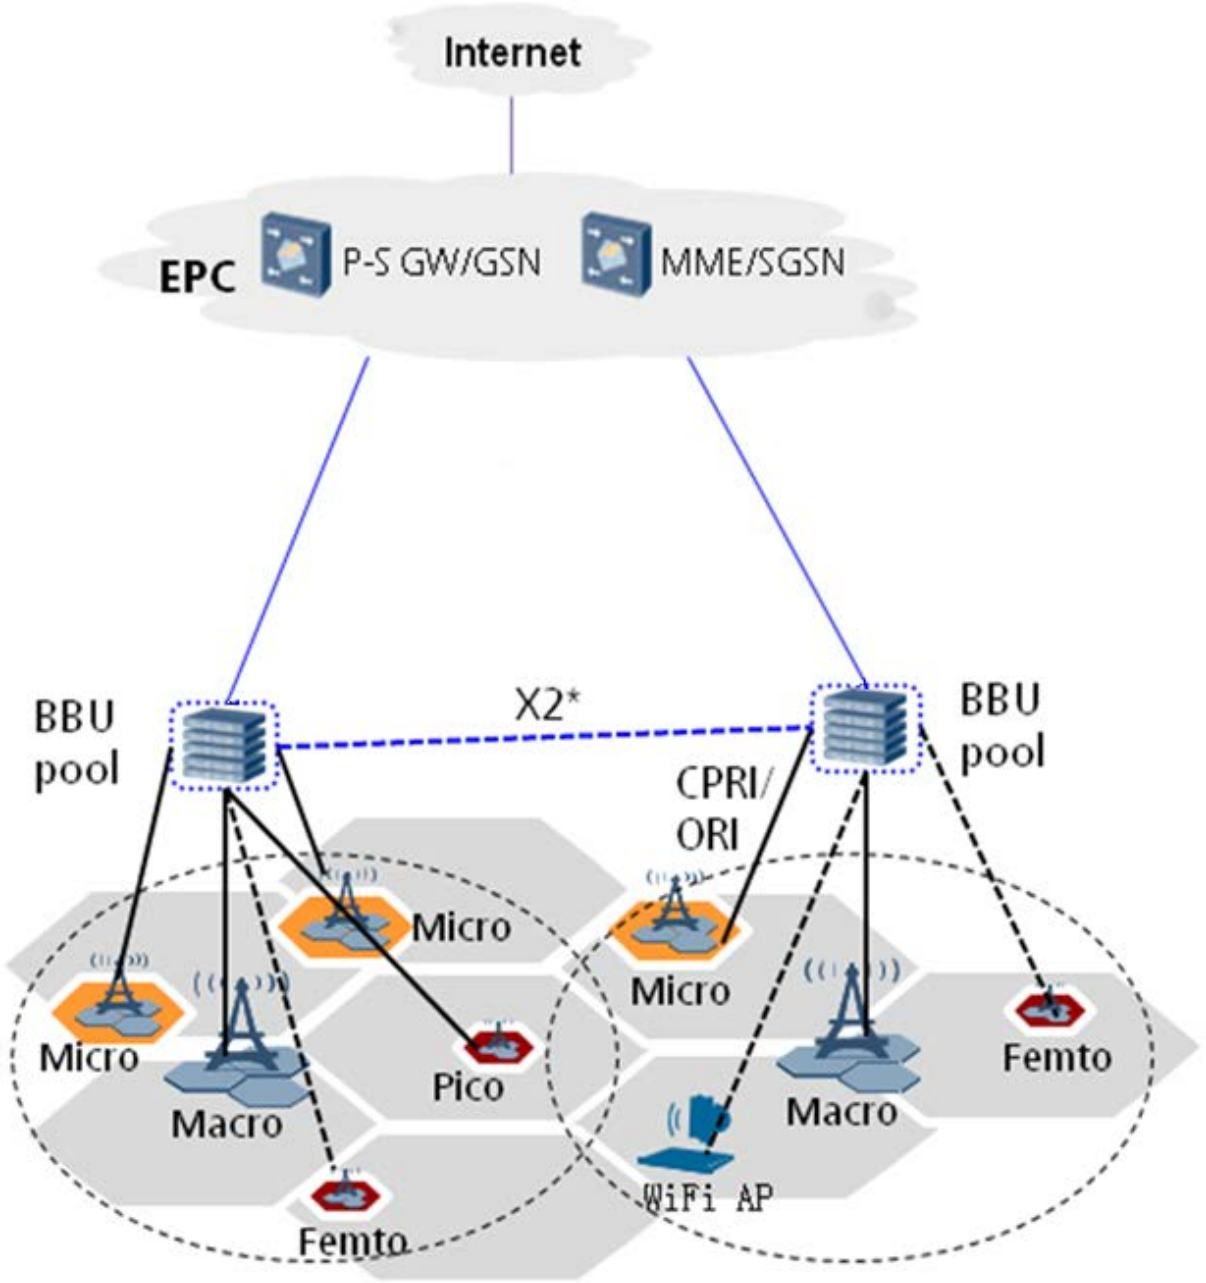

- Multiple RAN technologies from multiple vendors to be consolidated on a single BS to improve utilization
- BS virtualization can share resources among multiple logical RAN nodes from different systems.
- C-RAN for efficient resource utilization among different physical BSs.
- Baseband Unit (BBU) pool with high performance servers and real-time processing for signaling capacity

(ETSI #4)

Example: C-RAN/vRAN

- Determinism
  - Real-time patches to KVM
- Performance
  - DPDK accelerated Open vSwitch
- Reliability at scale
  - HA instances in OpenStack Nova
  - HA OpenStack
  - Fault management and reporting
  - IPv6 support – kernel through Nova, Neutron

NFV is not just OpenStack

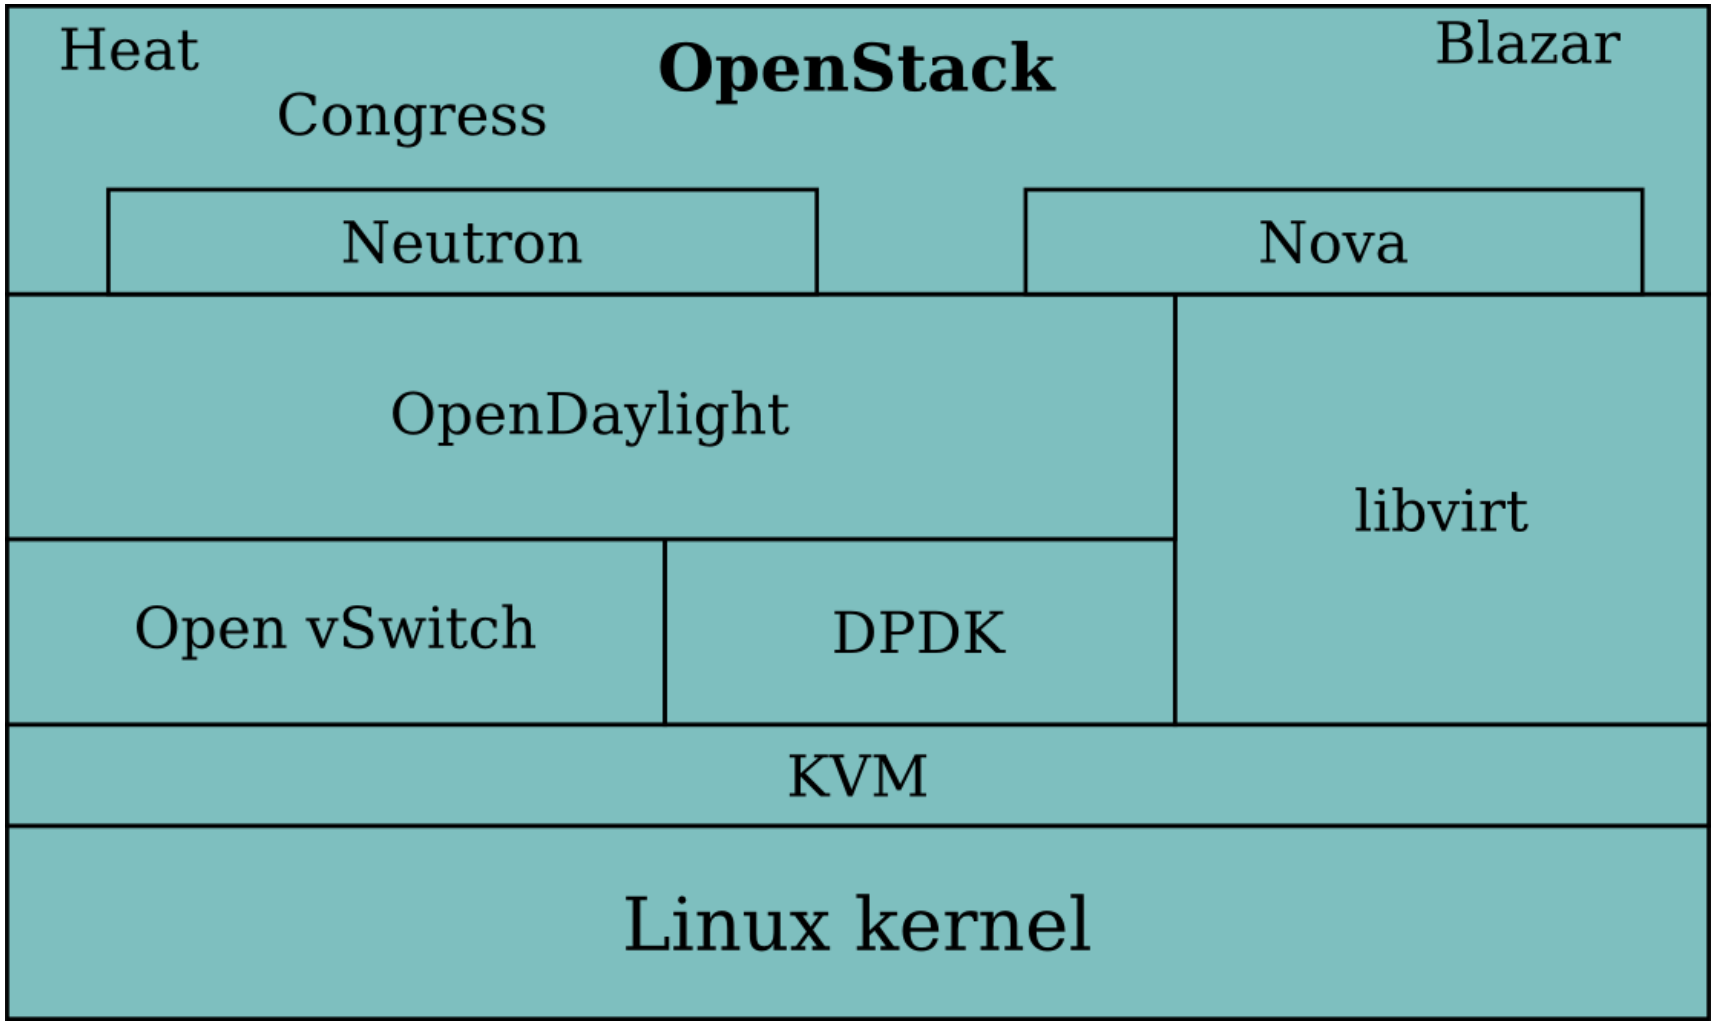

NFV Communities

ETSI NFV	driving industry trends, documenting functional requirements
OPNFV	integrating open source components, communicating needs upstream
OpenStack	abstraction/integration of virtual compute, network, storage
libvirt, qemu/KVM, Ceph, Open Daylight, Open vSwitch	Underlying infrastructure management

Upstream first

Upstream first

- Red Hat believes changes should be designed, written and proposed with upstream
- Maintaining significant out-of-tree patches across multiple projects is unsustainable

Not always easy!

Upstream First

- Requires discipline (small changes, discussed first)
- Relationships (your priority might not be their priority)
- Many communities – not just OpenStack!

Open Source NFV Summary

- NFV requires a collection of multiple projects
- Use-cases require features which touch many components
- Maintaining forks of multiple projects is unsustainable
- Multiple service providers and NEPs have similar requirements
- Engaging with upstream projects early is required for success

What is OPNFV?

- Open Source NFV Reference Implementation
  - Open, Transparent, Merit-based
  - Upstream first
  - Integrate and validate
- Consortium
  - Promote NFV use-cases upstream
  - Develop and test features to fill gaps
  - Red Hat is Platinum Founding member

OPNFV

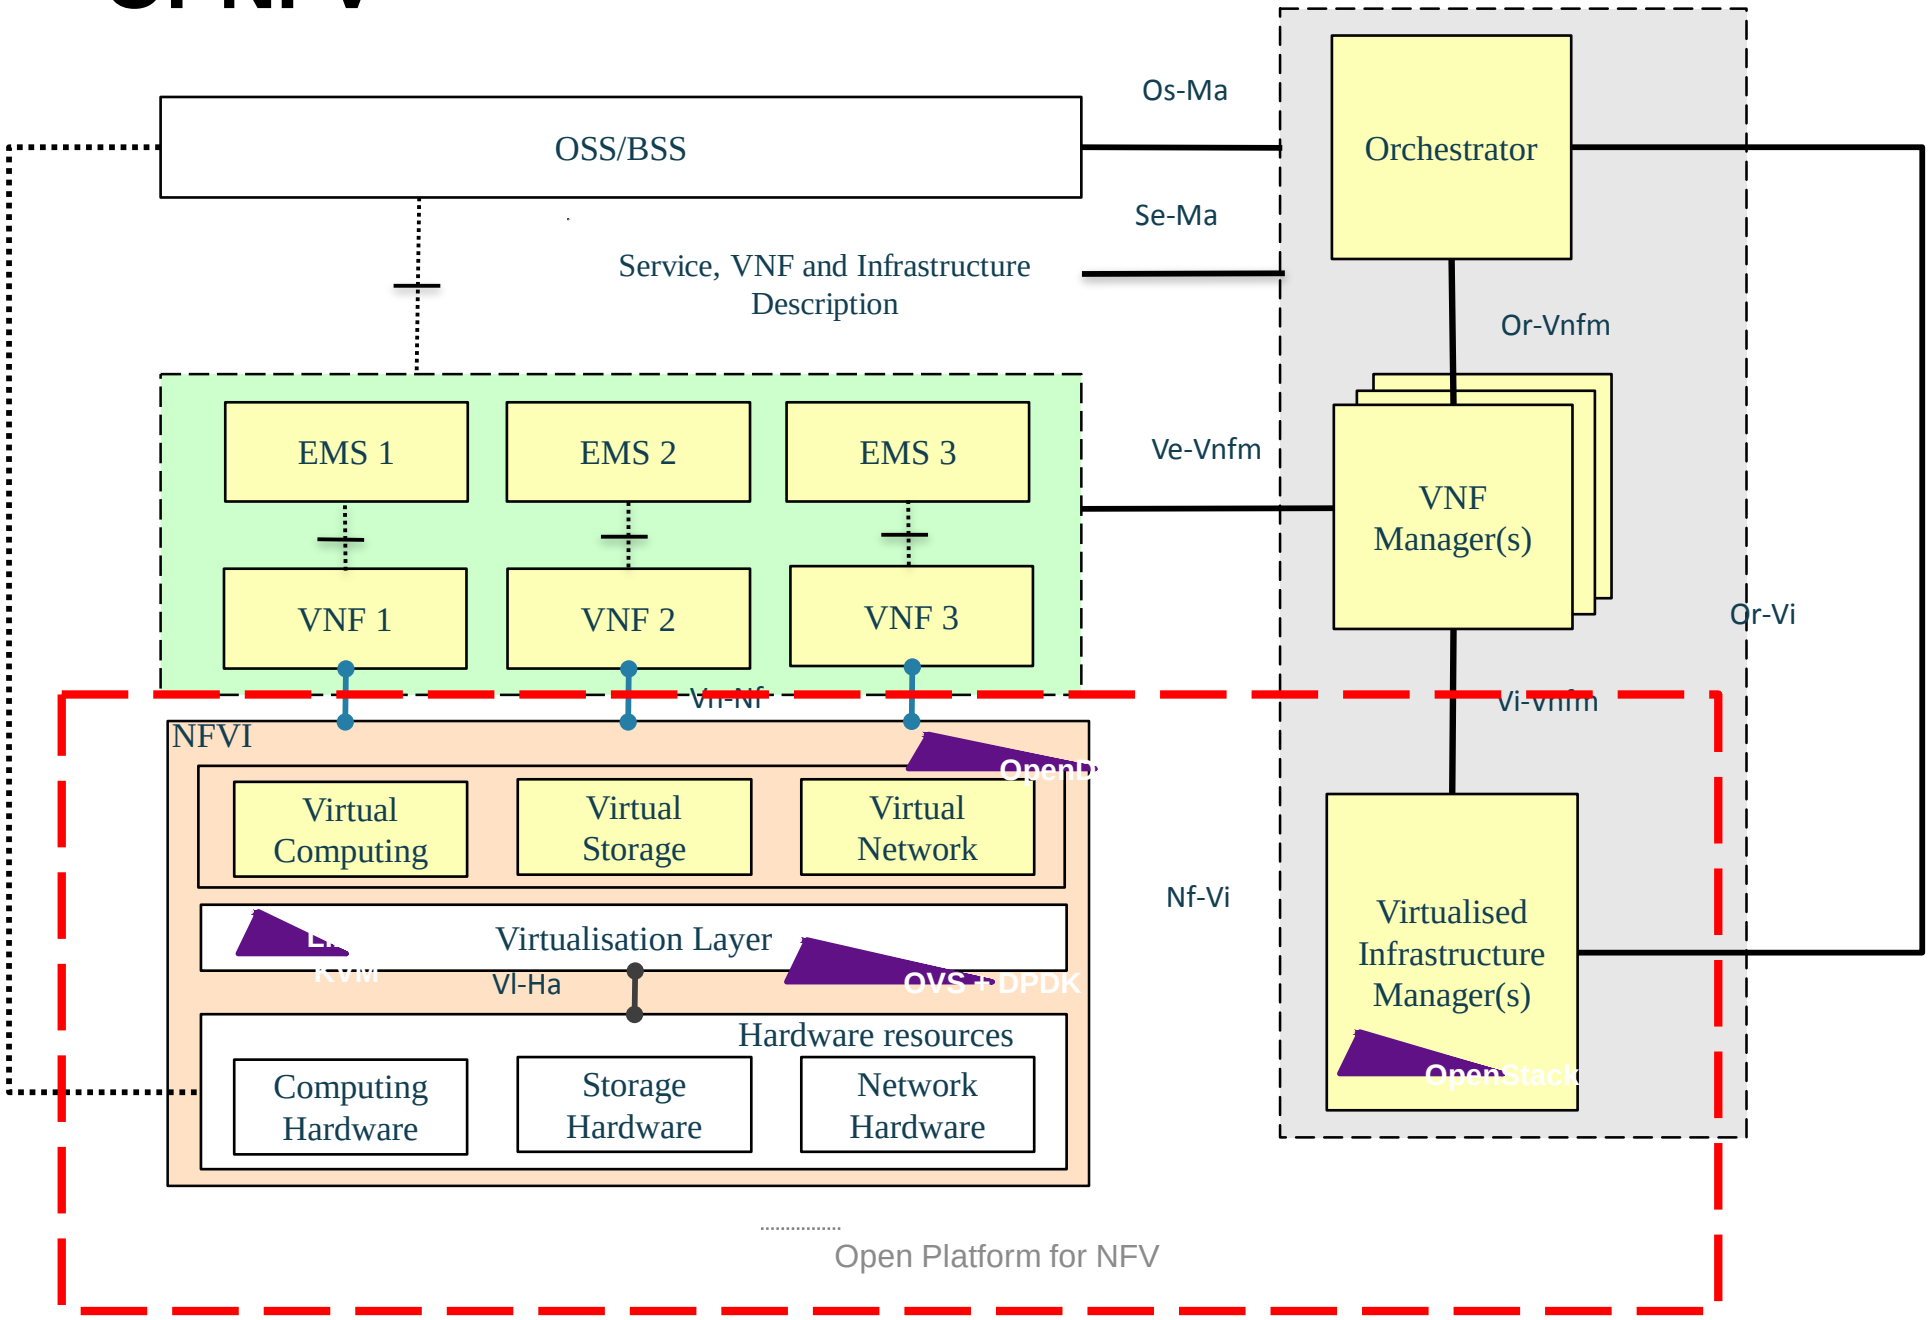
OSS/BSS
Os-Ma
Orchestrator
Se-Ma
Service, VNF and Infrastructure Description
Or-Vnfm
EMS 1
EMS 2
EMS 3
Ve-Vnfm
VNF Manager(s)
VNF 1
VNF 2
VNF 3
Or-Vi
Vn-Nf
Vi-Vnfm
NFVI
Virtual Computing
Virtual Storage
Virtual Network
Nf-Vi
Virtualisation Layer
KVM
Vl-Ha
OVS+DPDK
Virtualised Infrastructure Manager(s)
OpenStack
Hardware resources
Computing Hardware
Storage Hardware
Network Hardware
Open Platform for NFV

Goals of OPNFV

- Help OPNFV members engage relevant upstream communities
- Understand NFV requirements (from ETSI NFV), translate them for upstream developers
- Key communities:
  - OpenStack Telco Team
  - Open Daylight, Open vSwitch, DPDK
  - CentOS NFV SIG
  - ETSI NFV ISG (PoC)

OPNFV Projects

- Fault Management (Doctor)
- Continuous Integration (Octopus)
- Bootstrap/Get-started
- Virtualized Infrastructure Deployment Policies (Copper)
- Resource Management (Promise)
- High availability for VNFs
- IPv6-enabled OPNFV
- Characterize vSwitch Performance for Telco NFV Use Cases
- Software Fastpath Service Quality Metrics

Red Hat OPNFV plans

- Focus on identifying and fixing gaps upstream
- Red Hat's OPNFV distribution will be Red Hat Enterprise Linux OpenStack Platform
- We expect upstream projects to satisfy NFV use-cases, and will drive change upstream to ensure this
- Desire to avoid “Carrier Grade” forks of OpenStack

What is OpenDaylight?

- Open Source SDN
  - Open
  - Transparent
  - Merit-based
- Consortium
  - Facilitate
  - Advocate
  - Support
  - Red Hat is Platinum Founding member

OpenDaylight SDN Platform

- Modular, extensible, pluggable
- Java/OSGi/Karaf based platform
- Evolving towards model driven using YANG
- Multi-protocol
- Eclipse Public License

LEGEND

AAA: Authentication, Authorization & Accounting
AuthN: Authentication
BGP: Border Gateway Protocol
COPS: Common Open Policy Service
DLUX: OpenDaylight User Experience
DDoS: Distributed Denial Of Service
DOCSIS: Data Over Cable Service Interface Specification
FRM: Forwarding Rules Manager
GBP: Group Based Policy
LISP: Locator/Identifier Separation Protocol
OVSDB: Open vSwitch DataBase Protocol
PCEP: Path Computation Element Communication Protocol
PCMM: Packet Cable MultiMedia
Plugin2OC: Plugin To OpenContrail
SDNI: SDN Interface (Cross-Controller Federation)
SFC: Service Function Chaining
SNBI: Secure Network Bootstrapping Infrastructure
SNMP: Simple Network Management Protocol
TTP: Table Type Patterns
VTN: Virtual Tenant Network

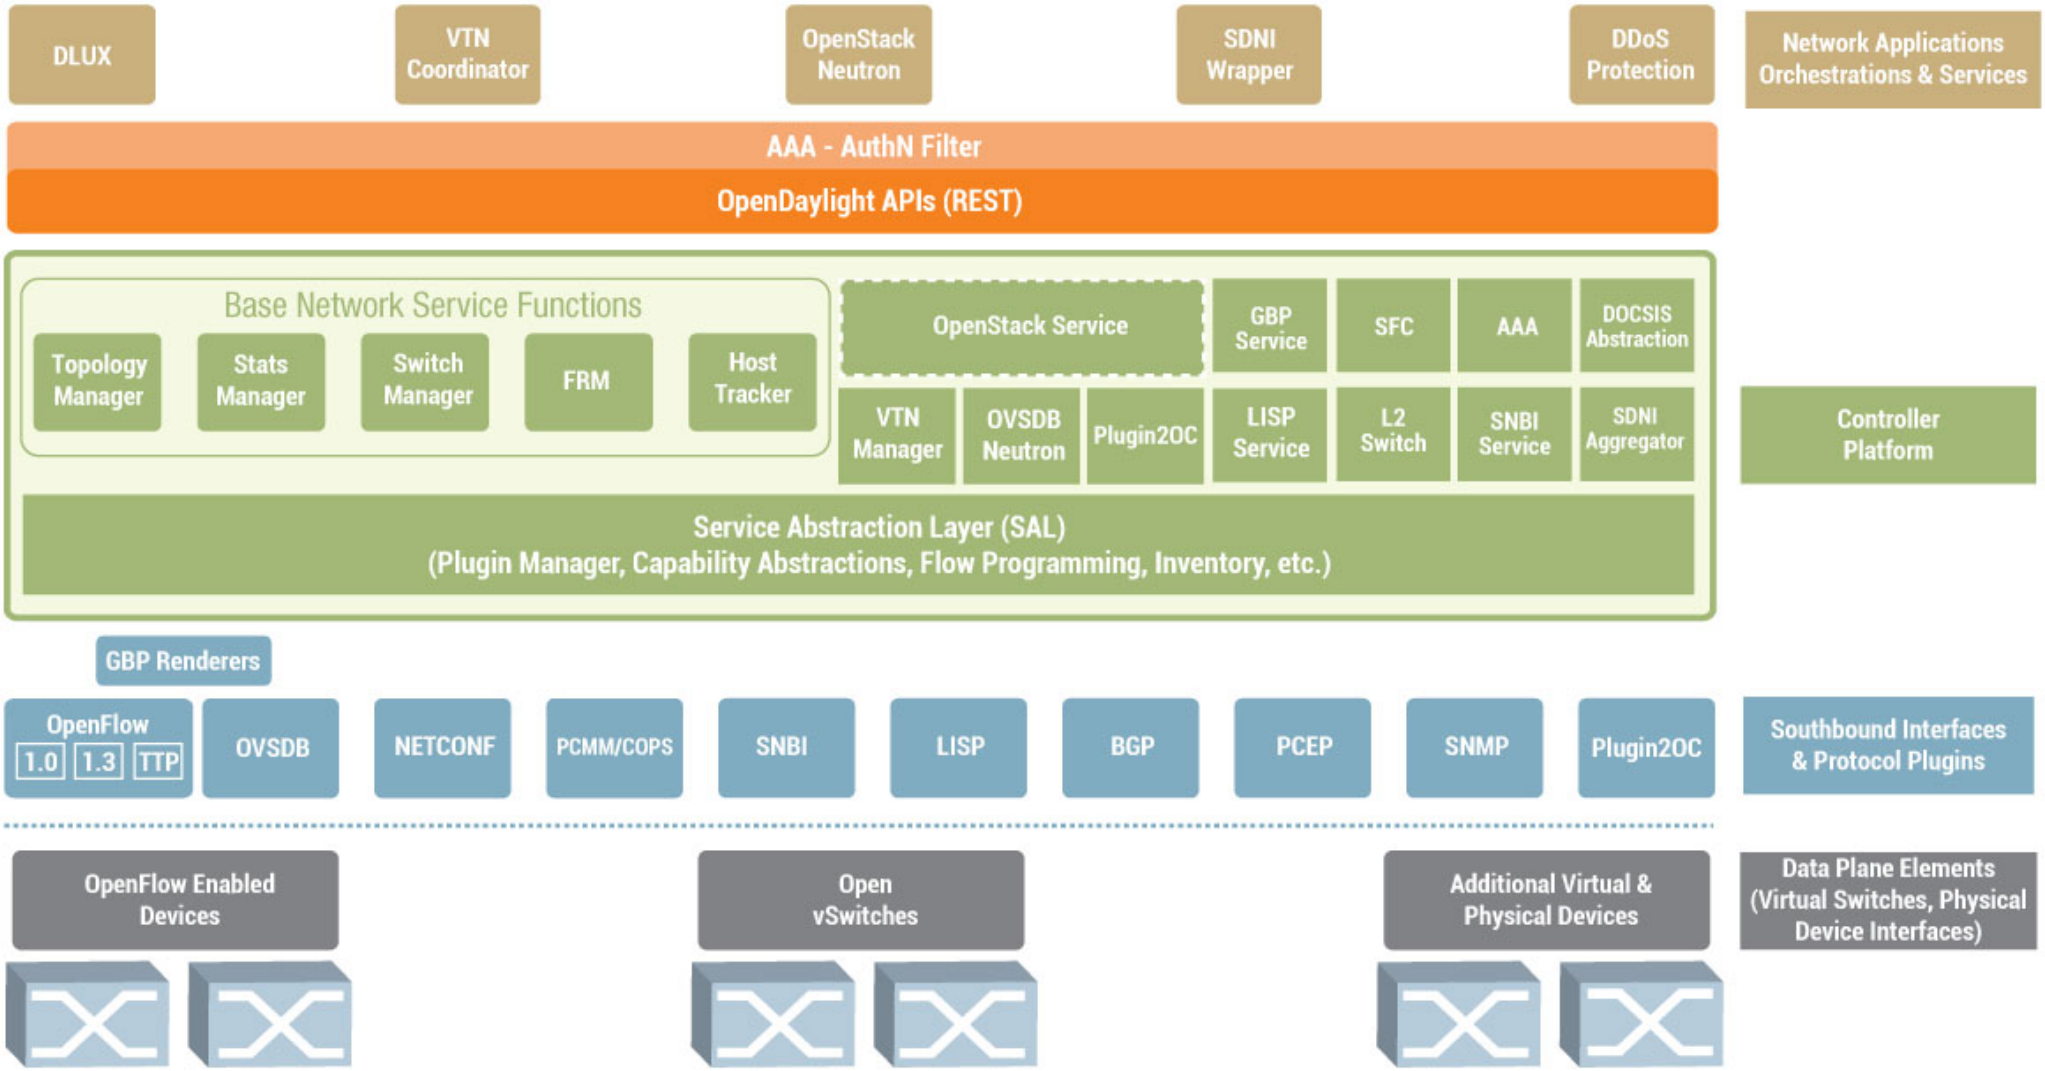

Red Hat ODL Focus

- Network Virtualization for OpenStack
  - ML2 ODL driver + extensions (L3, *aaS)
- Overlay networks (including HW)
  - Add OVSDB HW_VTEP schema support
  - Underlay informing (e.g. QoS)
- MD-SAL, AAA, OpFlex, SFC
- Infrastructure (testing and performance)

OVSDB Lithium Roadmap

- Cleanups – Karaf, Infra (Jenkins and Sonar), Test coverage, Eclipse integration, devstack and packstack integration
- Move to MD-SAL and clustering improvements
- Neutron service completion
- L3 service completion
- LBaaS feature work, VPNaaS and FWaaS
- DPDK integration
- SFC integration
- L4-7 appliance integration
- HW VTEP integration

Open vSwitch (OVS)

- Multi-layer virtual switch
  - Configuration managed w/ OVSDB
  - Flow tables controlled by OpenFlow
- Provides connection between VMs on same host
- Provides uplink to physical network via host NIC
- Data fast path in-kernel
- Challenges
  - kernel networking stack can be bottleneck
  - 64 byte packet processing rates suffer
  - Microflows vs. megaflows

DPDK

- Library for userspace packet processing
- Directly manages NIC with userspace poll mode driver (PMD)
  - Drivers for both physical and virtual NICs
- Polls driver NIC for packets, NIC DMAs directly to application buffers
- Platform specific optimizations
  - Hugepages, NUMA and cacheline aware
  - Batched packet processing
  - CPU instructions (SSE4, AVX, etc)

DPDK Challenges

- API/ABI compatibility, difficult to package in distribution
- Duplicate driver stacks, limited driver support
- Compile time rather than runtime optimizations
- Currently x86-centric
- OVS integration disables kernel features

OVS + DPDK

- Intel reports improved packet processing rates

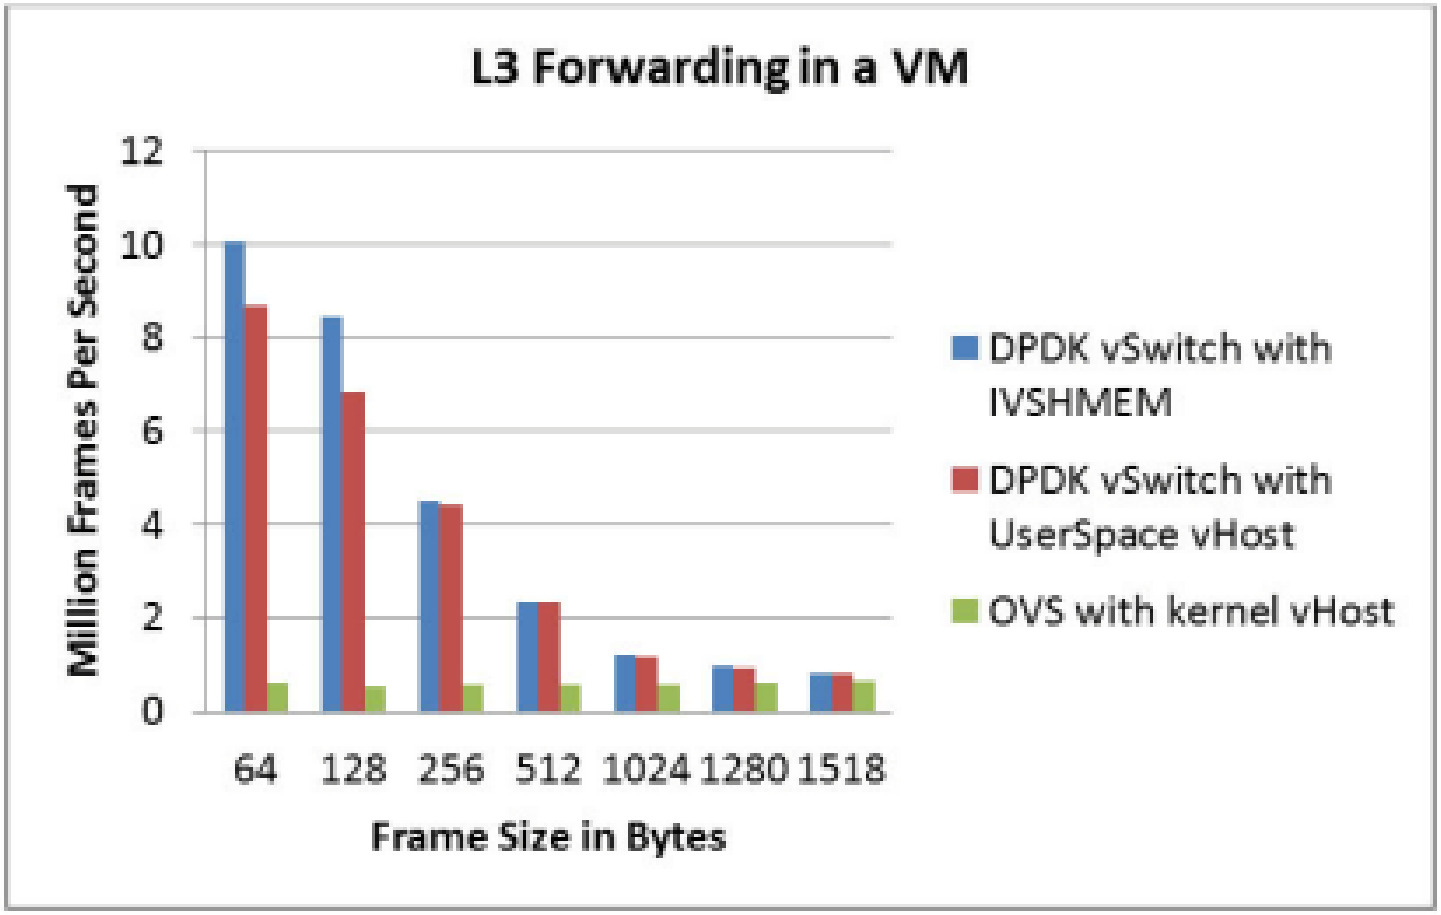

Putting it all together

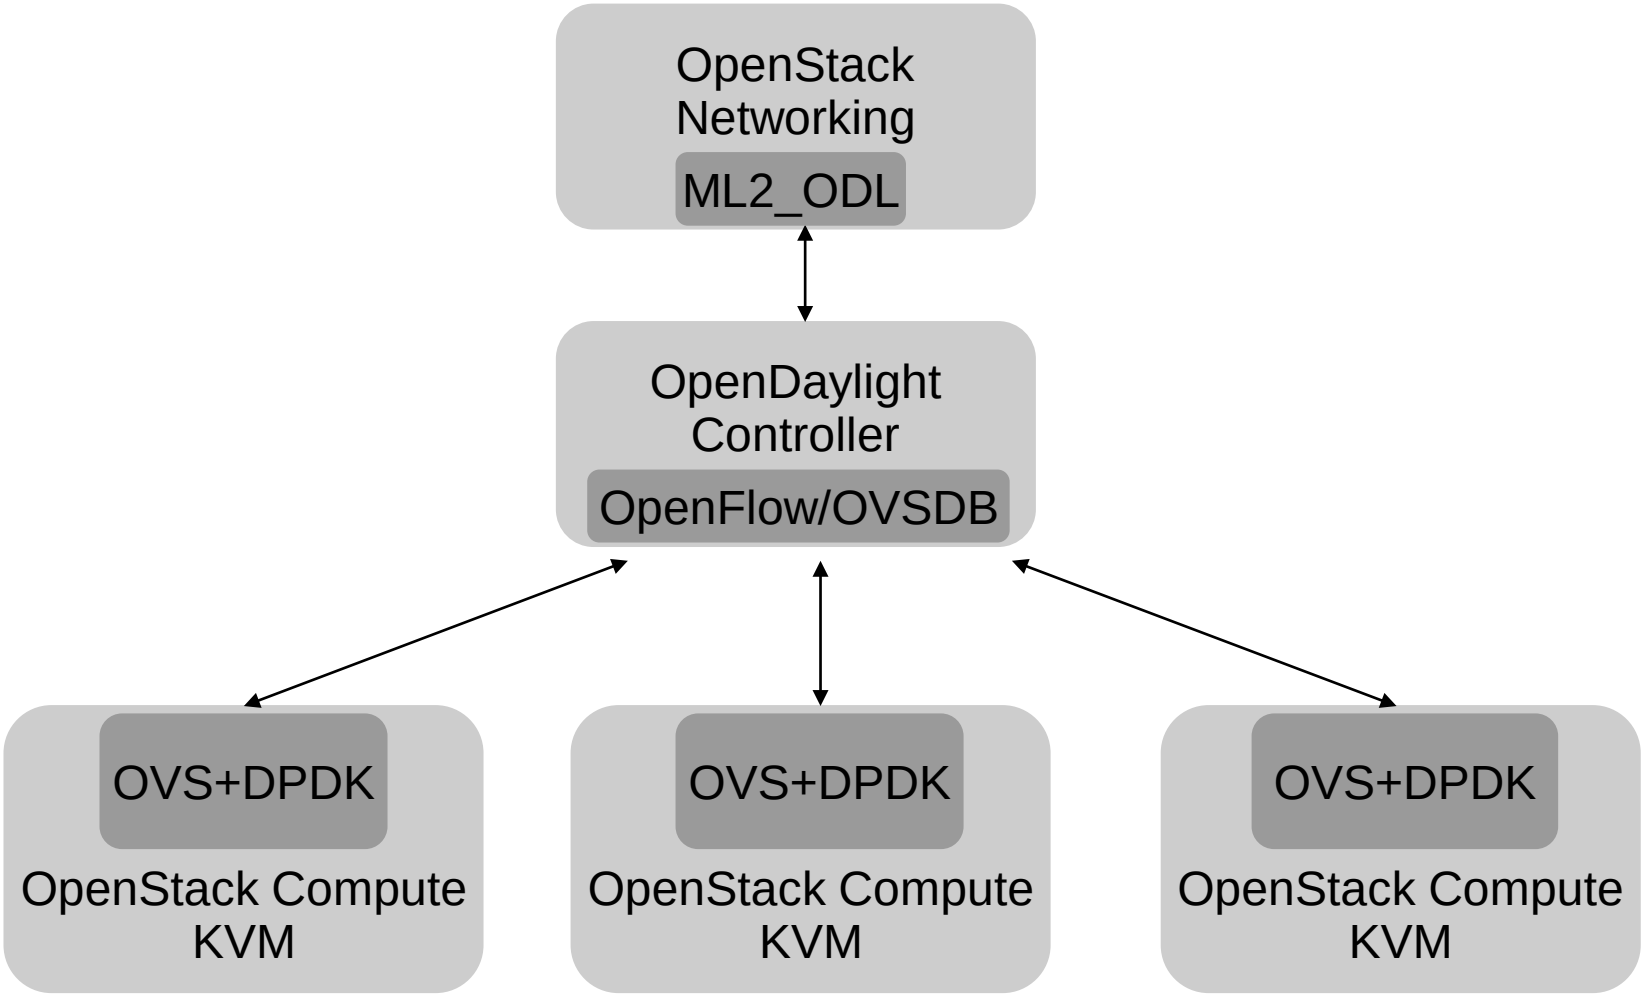

Making NFV and OpenStack real

Communities

- Created Sub-Team upstream to focus on NFV
  - https://wiki.openstack.org/wiki/TelcoWorkingGroup
  - Upstreamed NFV related patches to Juno cycle
- OPNFV brings additional resources to focus on NFV requirements in OpenStack
  - http://opnfv.org/

Red Hat and NFV

- Ability to effect change through entire Linux stack
- Active in Open vSwitch, DPDK
- Platinum, founding member of OpenStack Foundation, OpenDaylight Foundation and OPNFV
- Leading developer in Linux kernel, KVM, qemu, libvirt, OpenStack

Partner Ecosystem

- Bring production quality NFV platform to market rich functionality together with commercial partners

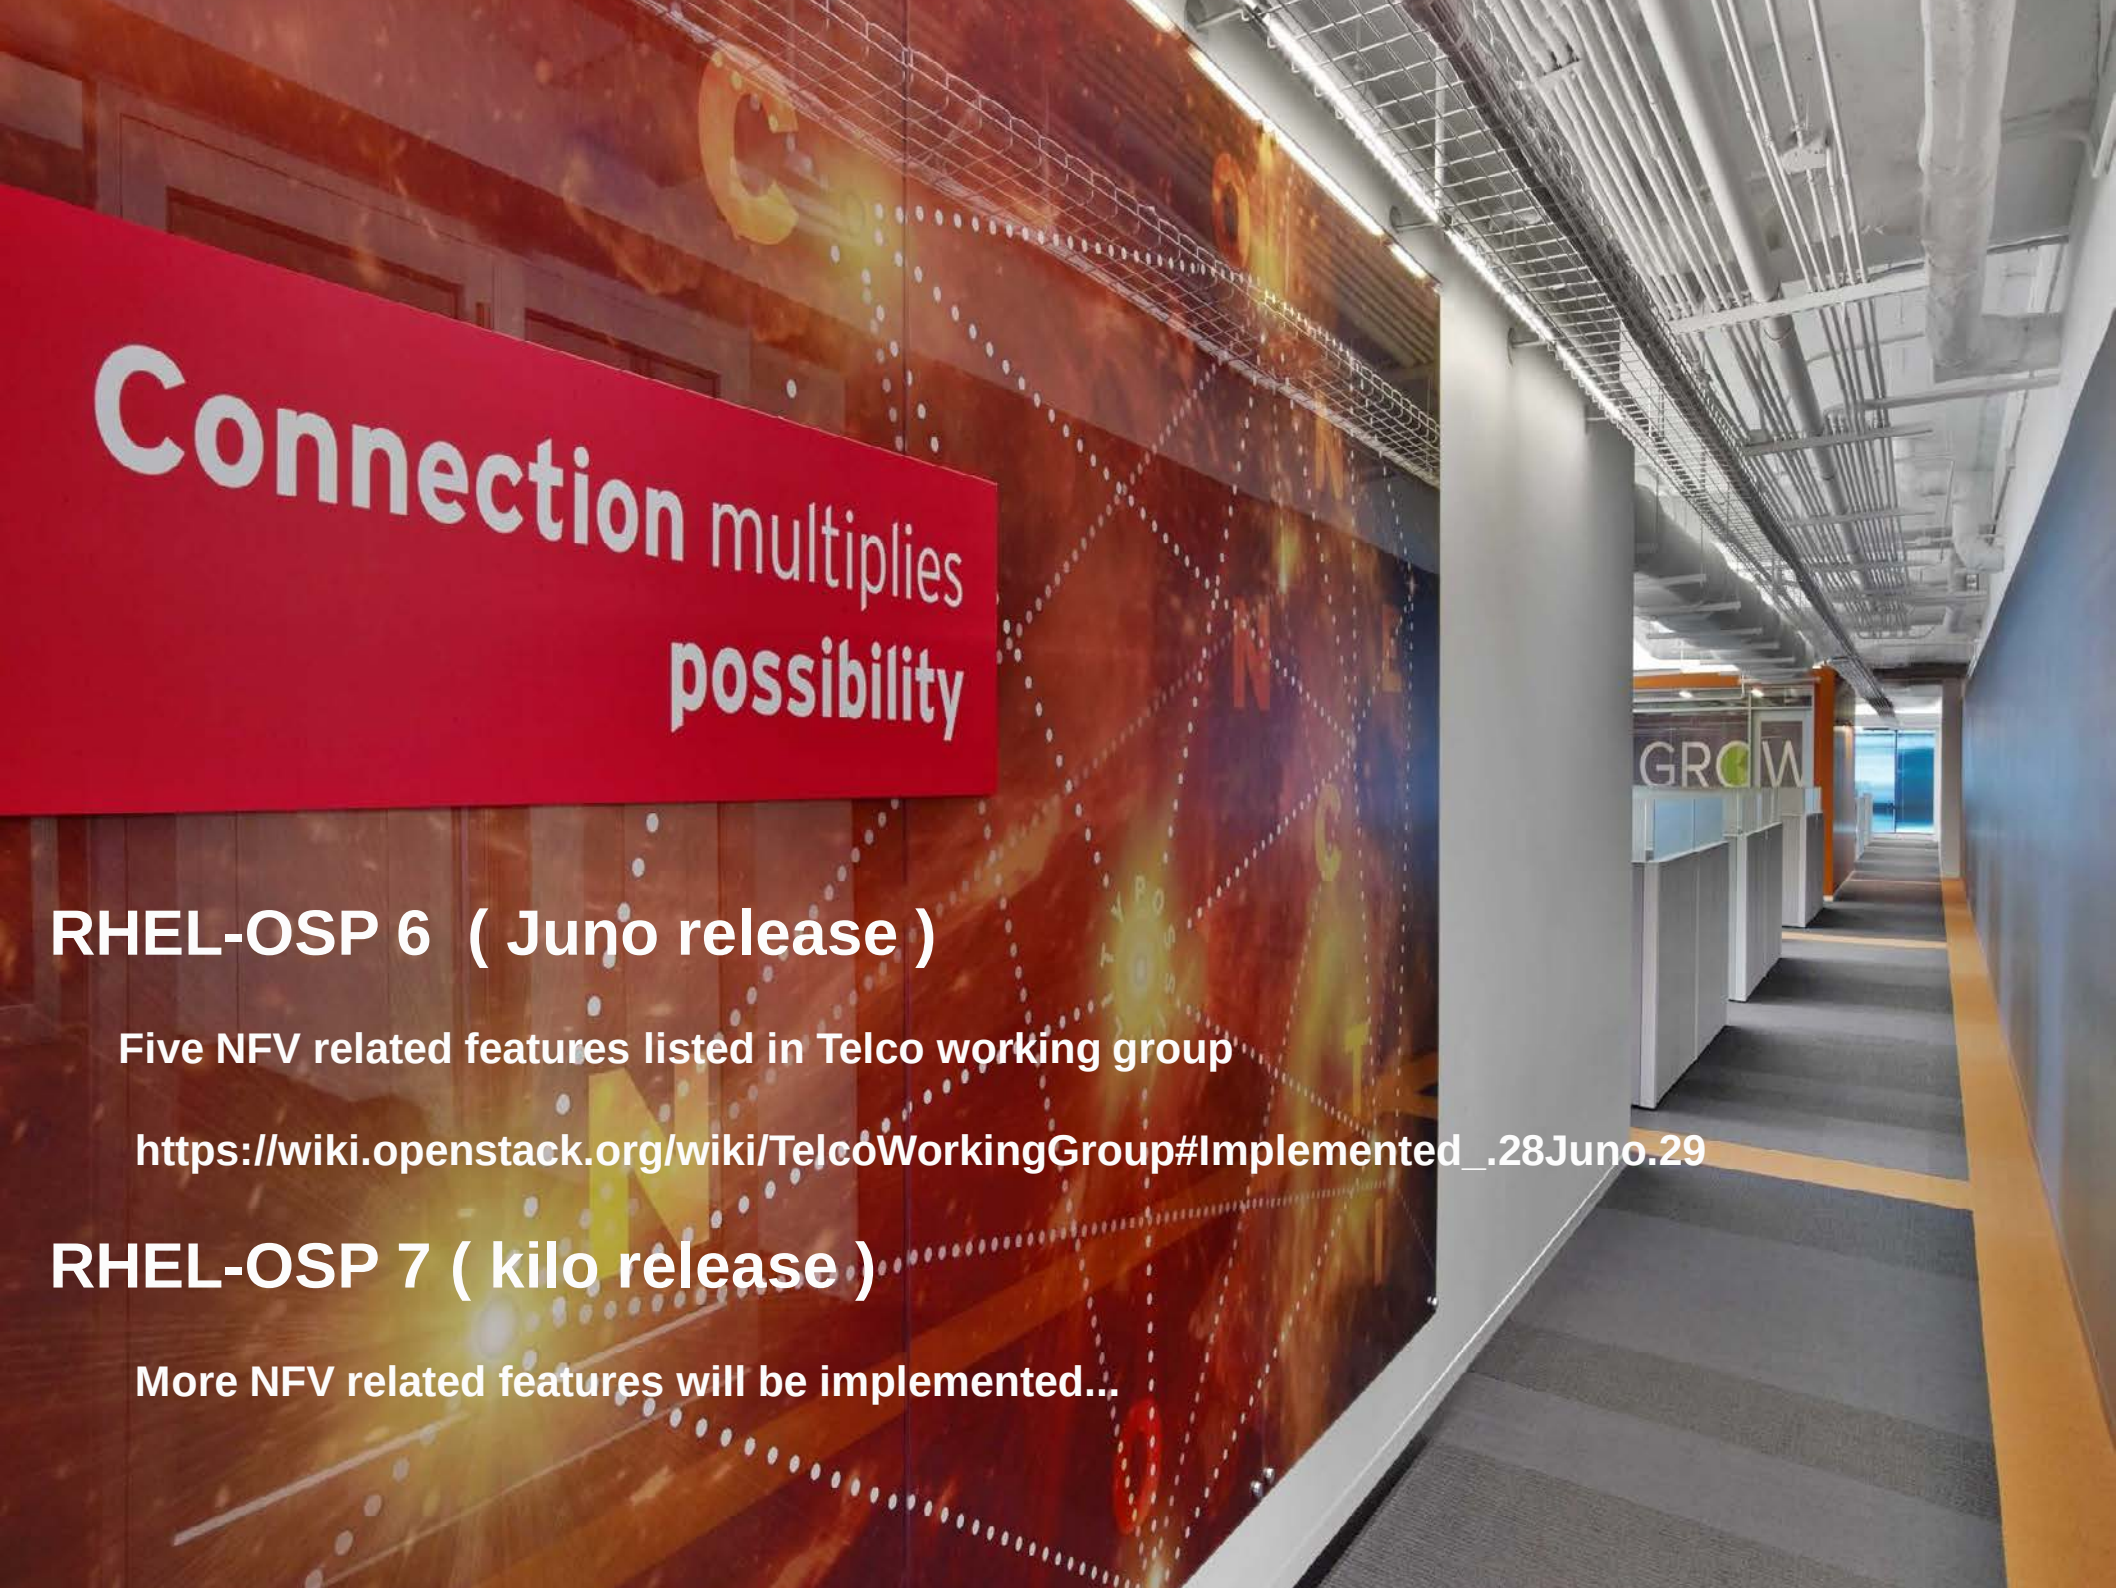

RHEL-OSP 6 (Juno release)

Five NFV related features listed in Telco working group

https://wiki.openstack.org/wiki/TelcoWorkingGroup#Implemented_.28Juno.29

RHEL-OSP 7 (kilo release)

More NFV related features will be implemented...

Thank you

NFV OpenStack Performance and Determinism

- CPU pinning
- NUMA aware CPU and memory scheduling
- NUMA aware I/O scheduling
- VM memory backed by hugepages
- OVS+DPDK accelerated packet processing
- SR-IOV accelerated packet processing
- Resource aware scheduling

NFV OpenStack Reliability

- All infrastructure deployed with HA
- VM HA (non-cloud aware application)
- Rich monitoring requirements
  - Fault detection, resource consumption
  - Ability to monitor Key Performance Indicators (KPIs)

NFV OpenStack Misc.

- Service VM, service insertion, and service chain APIs
- IPv6 support
- VLAN trunk to VM
- vNIC w/ no address
- 2 vNIC in one VM on same subnet
- OVF support
- Network QoS support
- Evacuate instance to scheduled host
- Heat multi-region support