redhat.

RED HAT®
OPENSTACK
PLATFORM

NFV use cases

Miki Kenneth, Yariv Rachmani

7 Feb. , 2017

IT Operations is Being Challenged

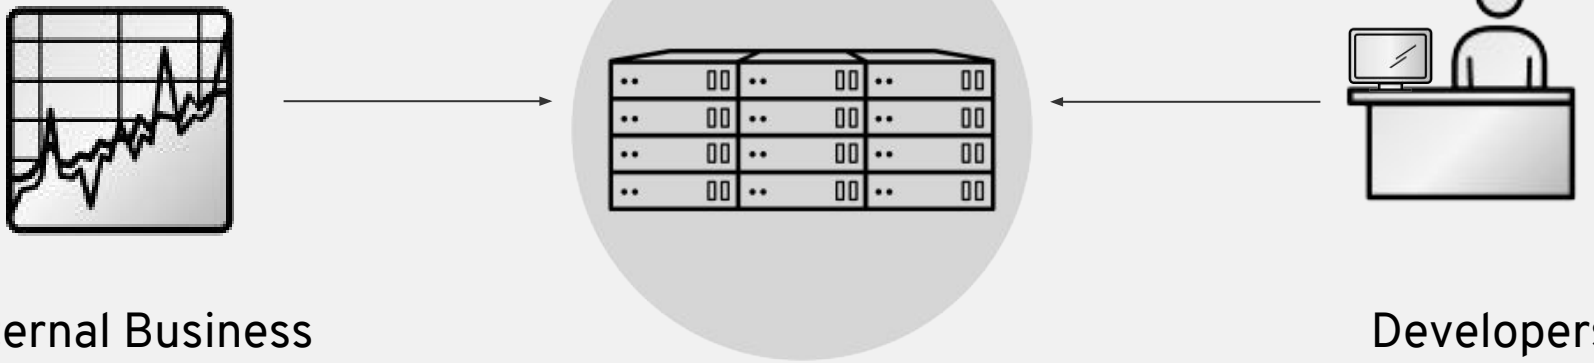

Internal Business

Developers

IT Operation
Challenges

Modern Apps & IaaS Enable the Digital Business

Streamlined and automated

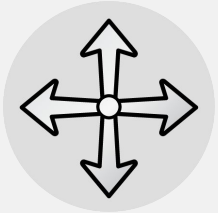

Elastic and scalable

Agile and responsive

Utility-like

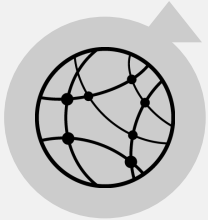

Orchestration, configuration management

Massive, fault-tolerant infrastructure

Rapid deployment

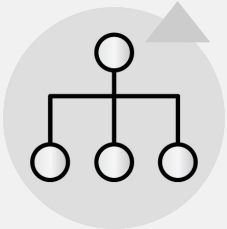

Managed, policy driven, & always on

Why This Evolution?

Existing infrastructure is not designed to cope with the demand

- **Data is too large**
  - We're producing vast amounts of unstructured data
  - Scaling UP no longer works. Scaling OUT is a necessity

- **Too many service requests**
  - More client devices coming online – Laptops, tablets, phones, watches, etc...
  - New Application demand generation is here

- **Applications and infrastructure were not designed for this level of demand**
  - Traditional capabilities are being exhausted

OpenStack Powers Digital Business

OpenStack meets the demands of "scale-out" digital business

- Brings public cloud-like capabilities into your datacenter
- Provides massive on-demand (scale-out) capacity
  - 1,000's → 10,000's → 100k's of VMs
- Removes vendor lock-in
  - Open source provides high-degree of flexibility to customize and interoperate
- Community development = higher "feature velocity"
  - Features and functions you need, faster to market over proprietary software
- Greater automation, resource provisioning, and scaling

OpenStack

- Provides **infrastructure and application** cloud services such as object storage, elastic scaling, load balancing, database, analytics, and search
- Is designed for **maximum scale and cost efficiency**
- Focuses on **availability by zones instead of per virtual machine** (VM), with more of the availability responsibility as part of the application architecture
- Leans toward **inexpensive commodity hardware**
- Focuses on **new apps written and optimized for cloud** environments
- Is designed to **create services, which apps are then written to utilize**

OpenStack

The other primary type of cloud system software is an evolution of traditional enterprise server virtualization seen more in private cloud deployments that:

- Is **compatible with existing applications** and operating systems
- Focuses on **ensuring mission-critical** VMs never go down, with availability
- primarily a function of the **infrastructure on a per-VM basis**
- **Accommodates enterprise hardware** and architectures (fault-tolerant servers, SANs, VLANs, etc.)
- Is designed and tuned to meet **individual application requirements**

What is OpenStack?

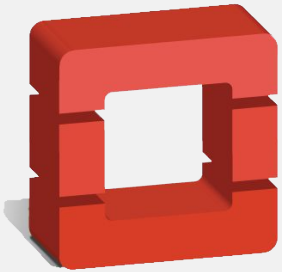

openstack™

CLOUD SOFTWARE

Service

- NOVA - Compute
- NEUTRON - Networking
- SWIFT - Object Storage
- CINDER - Block Storage
- KEYSTONE - Identity
- GLANCE - Image Service

- Horizon Dashboard
- Ceilometer Telemetry
- Heat Orchestration
- Trove Database
- Sahara Elastic Map Reduce
- Ironic Bare-Metal Provisioning
- Zaqar Messaging Service
- Manila Shared Filesystems
- Designate DNS Service
- Barbican Key Management
- Magnum Containers
- Murano Application Catalog
- Congress Governance

Nova

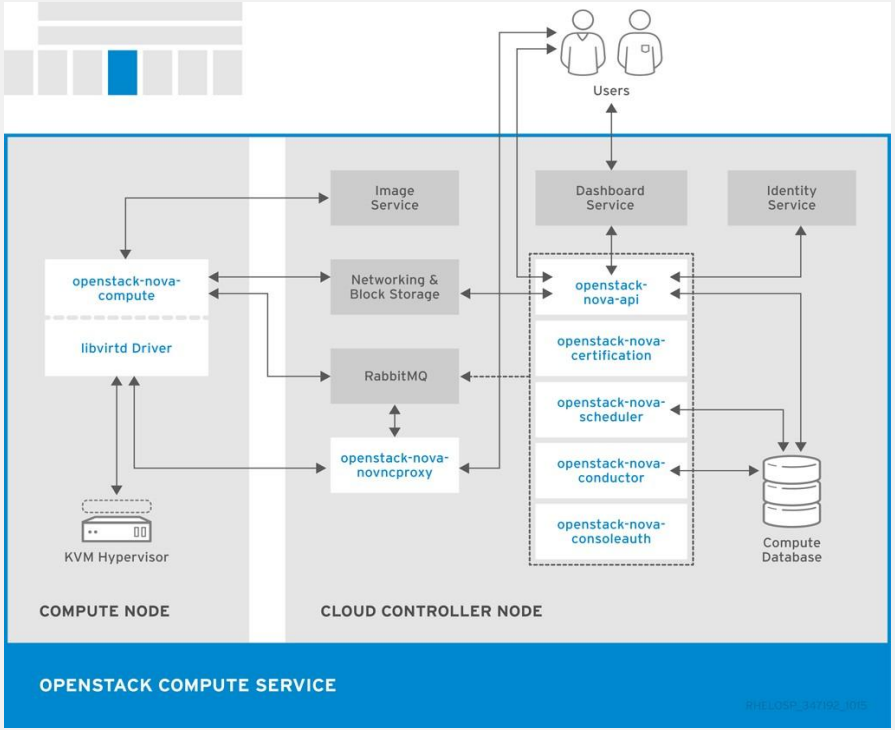

No need to manage hypervisors individually, due to distributed design of OpenStack, at any scale. Supports KVM and VMWare (vCenter)

Defines which choices are available to tenants: flavors offering specific capabilities and carefully planned capacity and overcommit ratios.

Networking (Neutron)

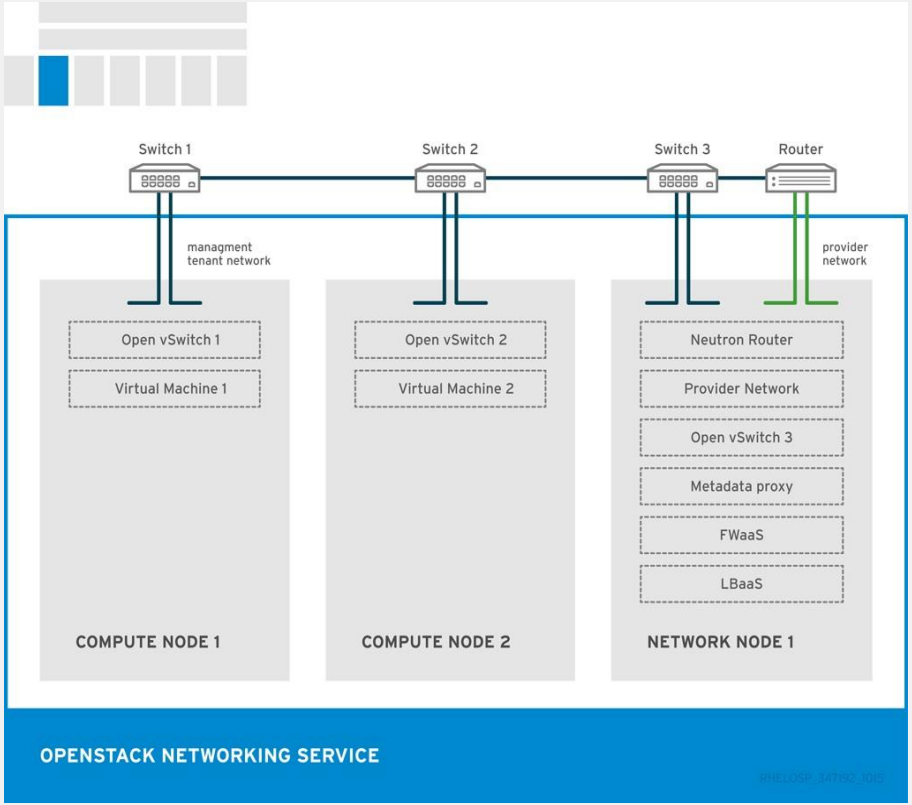

Create,Remove,Update,Delete (CRUD) networks,

subnets and ports, for basic L2 and L3 with IP Address Management (DHCP)
Define a tenant network (overlay)
Additionally:

- Provider networks
- Quotas
- Security Groups (per port)
- East/West L3 routing with tenant-defined routers
- External gateway, NAT, floating IPs
- Load balancing, VPN and Firewall

Red Hat Cloud Suite for Applications

Integrated DevOps Platform for the enterprise

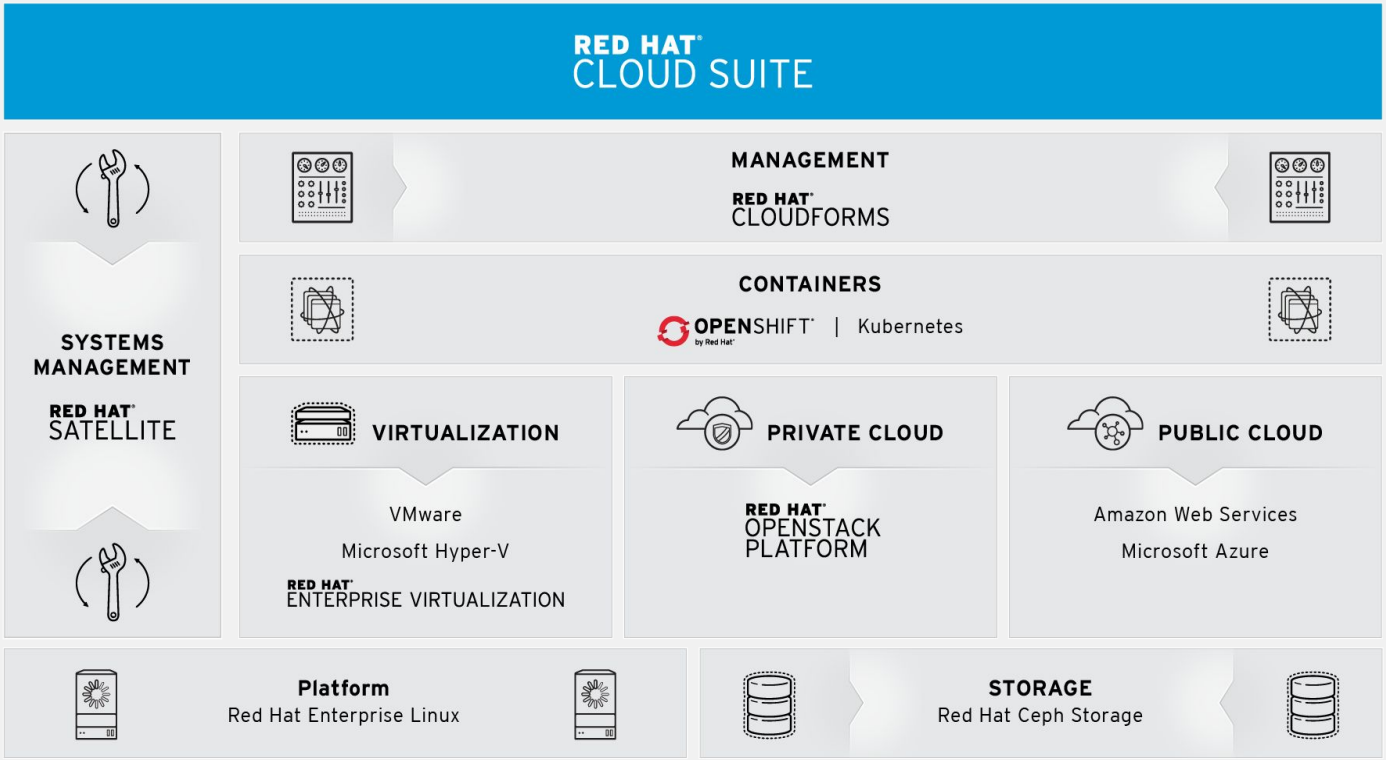

Red Hat OpenStack Platform - Open, Production-Ready Framework for the Cloud

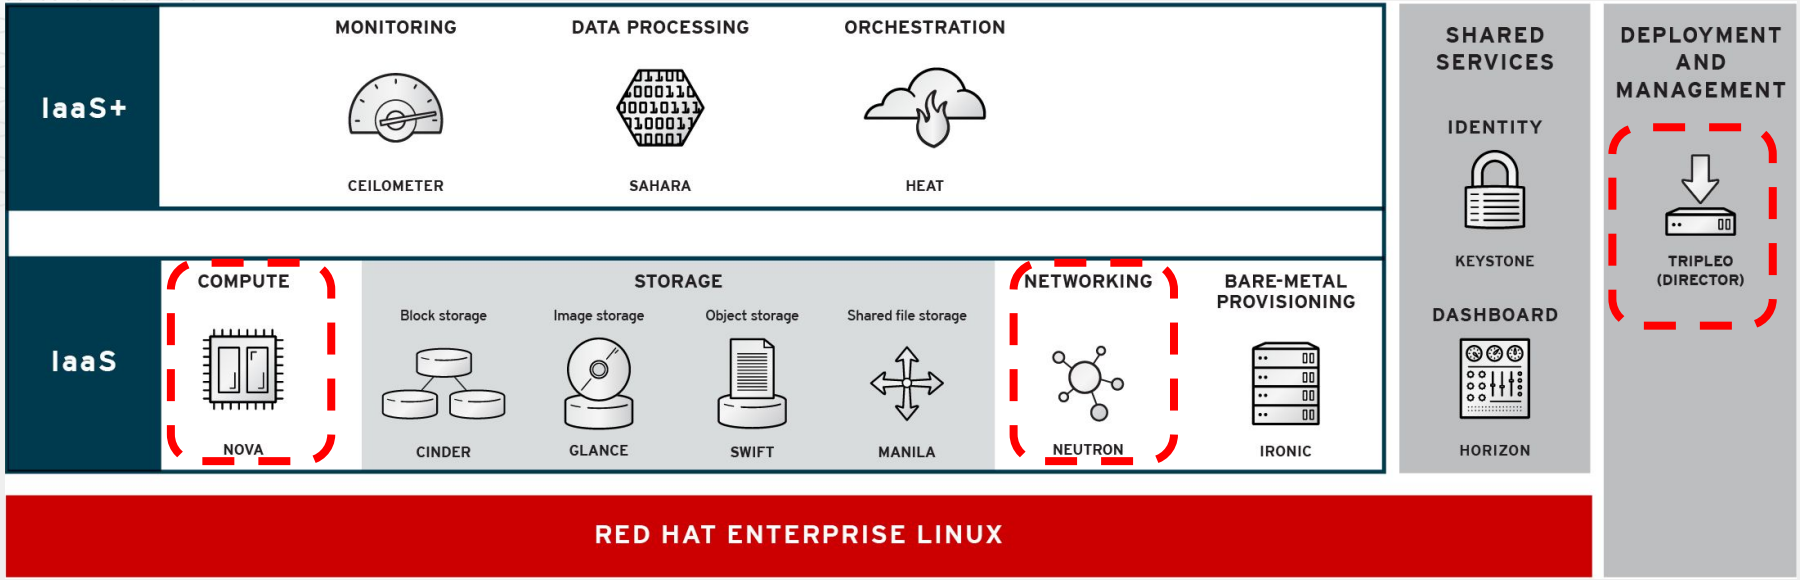

- Built on, integrated with Red Hat Enterprise Linux
- Co-engineered with Red Hat Enterprise Linux
- Tested and Certified by Red Hat for production workloads.

RED HAT® OPENSTACK® PLATFORM

redhat.

RH OSP Director/TripleO

OPENSTACK ORCHESTRATION

PLANNING

- Network topology
- Service parameters
- Resource capacity

DEPLOYMENT

- Deployment orchestration
- Service configuration
- Sanity checks

OPERATIONS

- Updates and upgrades
- Scaling up and down
- Change management

RED HAT® OPENSTACK® PLATFORM

redhat.

Red Hat OpenStack Platform - 10k Foot View
Roadmap at a Glance

	9 **RHOSP 9** *(Summer 2016)*	10 **RHOSP 10 (Newton)** *(December 2016)*	11+ *(Spring 2017 and beyond)*
Private IaaS	- Update & Upgrades - CloudForms ITB - Ceph Integration - DR capabilities - Nova in containers *(tech preview)*	- **Tighter CloudForms integration** - **Composability & Split stack** - **GUI & Validations** - **Simplified, Ops-friendly HA** - **Neutron DVR** - **Operator Tool Clients**	- Containerized services - Multi-site support - OpenShift OOB - Operator Tools
NFV - Telco	- RT-KVM *(tech preview)* - OVS-DPDK *(tech preview)* - OpenDaylight *(tech preview)*	- RT-KVM *(tech preview)* - **OVS-DPDK full support** - **SR-IOV director support** - OpenDaylight *(tech preview, now available in director)* - Hyper-Converged Infra *(limited availability)*	- RT full support - ODL full support - Carrier Grade - Hyper-Converged Infra - VNF Containers - VM & Containers service chaining

RED HAT® OPENSTACK® PLATFORM

redhat.

Reference architecture

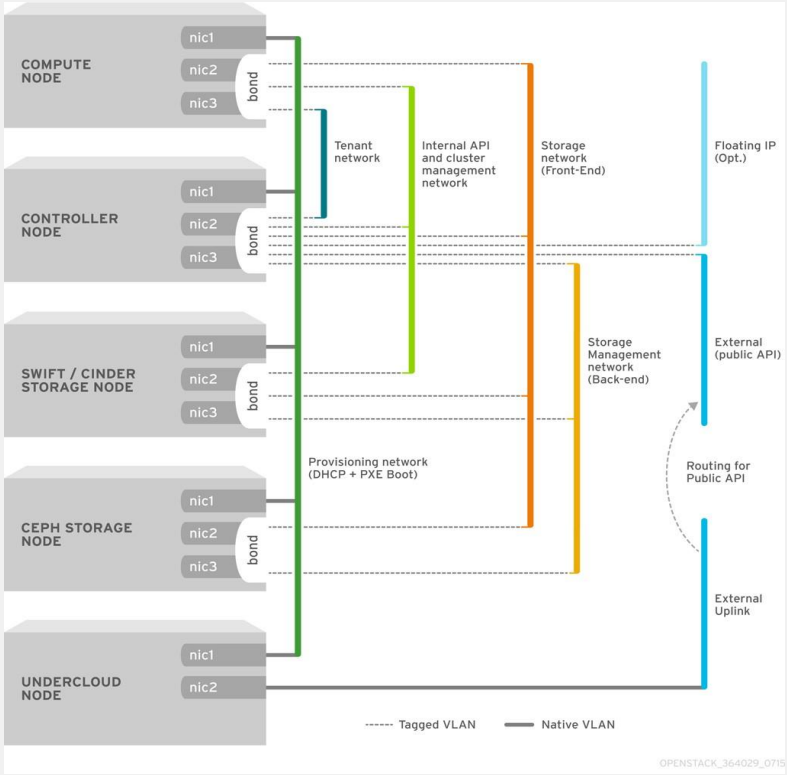

SDN

SDN is way to manage networks that separates the control plane from the forwarding plane. SDN is a complementary approach to network functions virtualization (NFV) for network management.

SDN controllers offered by BigSwitch, Nuage, Cisco, ODL
Part of the controllers are proprietary other manages by community.

SDN Controllers

The SDN Controller platform typically contains a collection of “pluggable” modules that can perform different network tasks

OpenStack Neutron interacts with many ML2 plugins for the different SDN controllers.

RH offers ODL, community, controller within its TripleO/Director, deployer, through Neutron ODL ML2 plugin.

NFV is on the road map of most Service Providers

"ECOMP is a stake in the ground. It's a declaration that networks of the future will be software-centric, that they'll be faster, more responsive to customer needs, and more efficient,"

— Chris Rice, Senior Vice President, Domain 2.0 Architecture and Design, AT&T

Deployment shifted from POC to commercialized service:

- *AT&T launched network on Demand deployments across 76 countries*
- *Orange Business Services SDN connectivity service for multinational companies available in 7- countries*
- *China, Japan Ties 1s Launced large scale vIMS, vEPC services*

NFV Achitecture

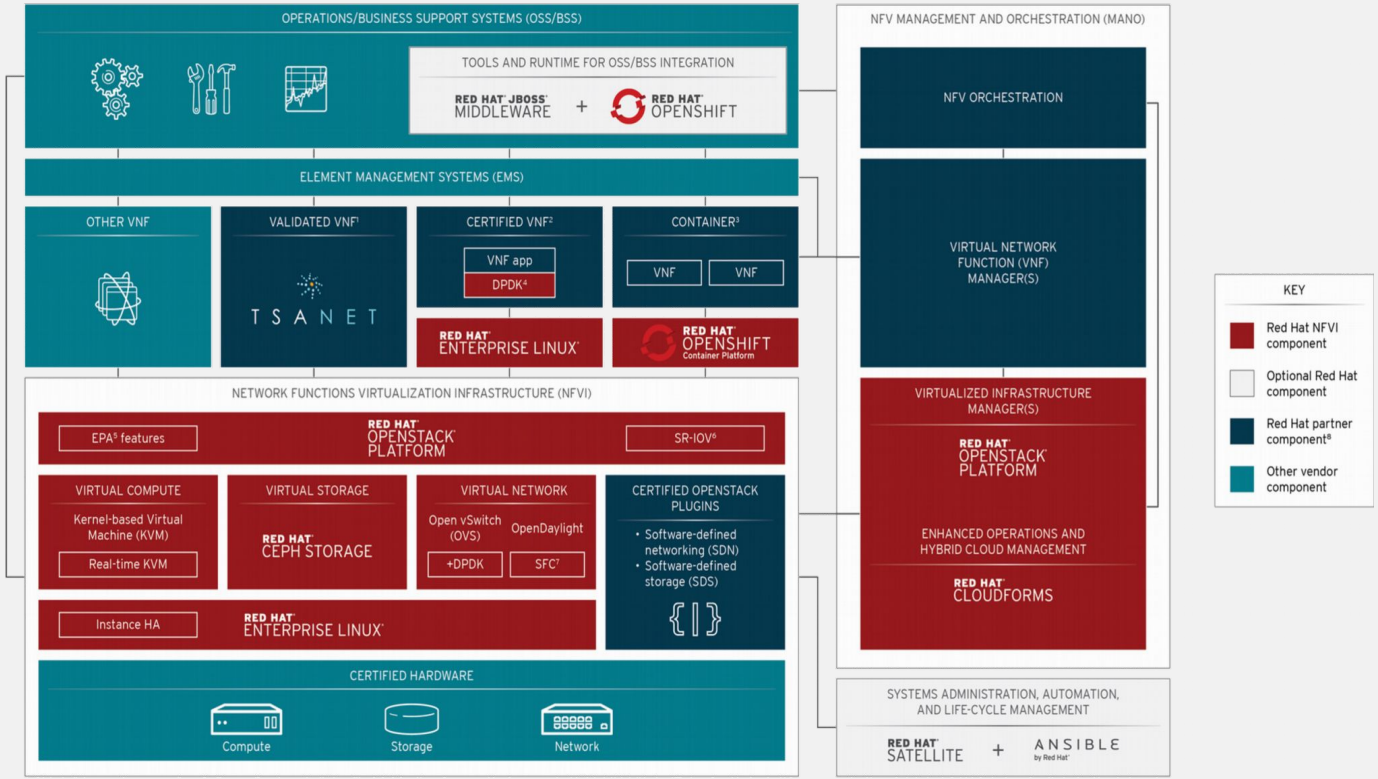

What is NFV?

Network functions virtualization (NFV) offers a new way to design, deploy and manage networking services.
NFV decouples the network functions, such as network address translation (NAT), firewalling, intrusion detection, domain name service (DNS), caching, vSwitches and vRouters to name a few, from proprietary hardware appliances.
The VNF(Virtual Network Function), can run as software service on vms.

NFV Application types

- No DPDK, RT Applications, uses RT Kernel at the hypervisor and at the vm
- DPDK, L2/L3 applications, such as vSwitches and vRouters running on the VM

NFV HW/SW enablers

NFV IAAS requires dedicated NICs in Computes run with unique SW approach:

- Each ports could handle 10G to 100G DPDK/SRIOV enabled with.

Vendors such as Intel, Mellanox, Qlogic, etc

- DPDK http://dpdk.org/ is a set of SW libraries and drivers for fast packet processing. It designed to run on any processors
- Current DPDK driver uses PMD, Pool Mode Driver, which utilise 100%CPU

at any moment

NFV SRIOV, DPDK applications

DPDK applications in the VM move from kernel to user space

No Kernel networking APIs, using DPDK library through PMD to read/write packets from NIC ports.

To achieve high performance every PMD core must be isolated from interrupts from other CPUs or IRQ events, it also must interact inside CPU NUMAs, to ensure 0 packet loss with minimum latency.

EPA, Enhanced Platform Awareness

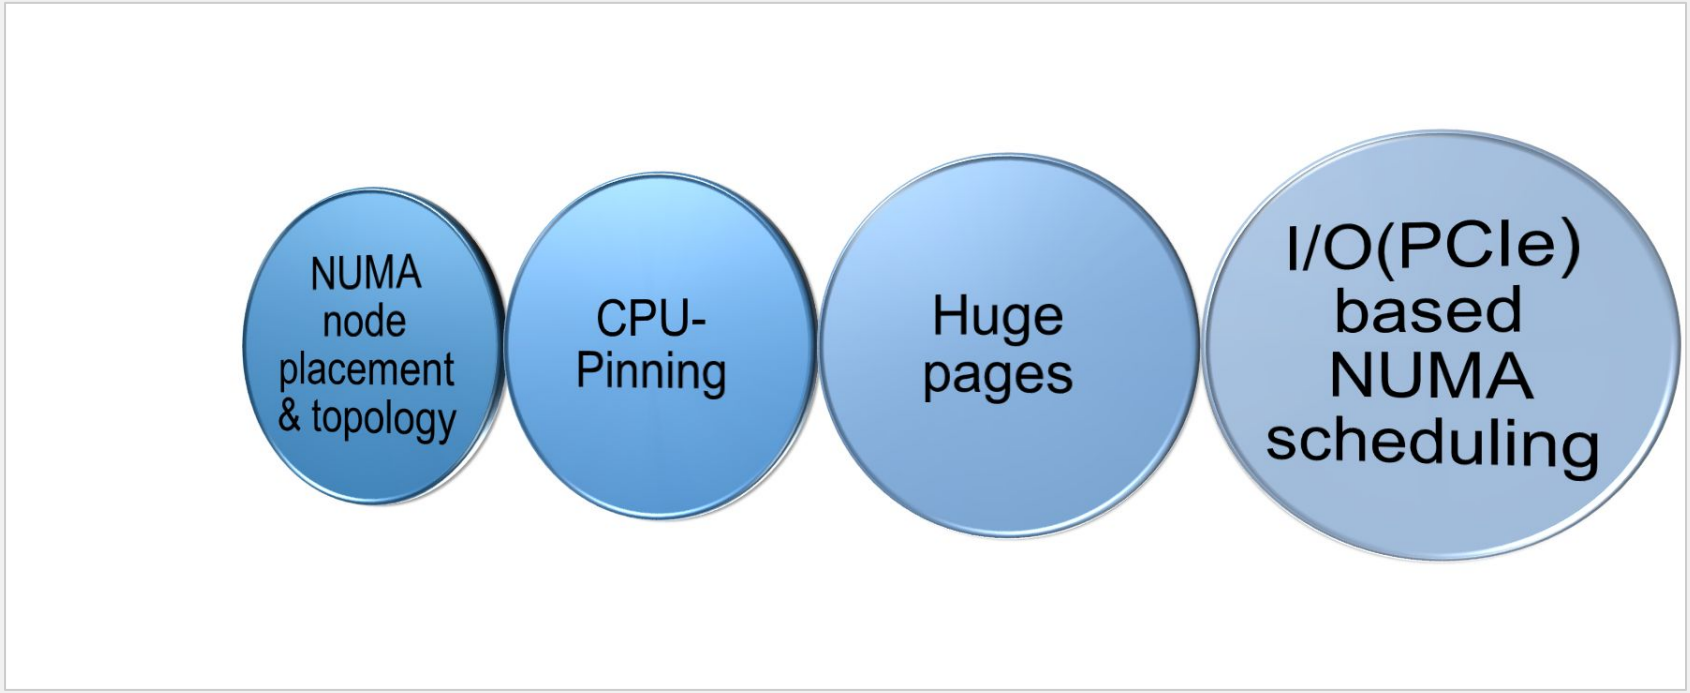

RED HAT® OPENSTACK® PLATFORM 10

redhat.

Numa Awareness

Scheduler Features

[DEFAULT]

scheduler_default_filters = RetryFilter, AvailabilityZoneFilter, RamFilter, ComputeFilter, ComputeCapabilitiesFilter,ImagePropertiesFilter, PCIPassthroughFilter, NUMATopologyFilter

NUMA flavor parameters

```
$ openstack flavor set FLAVOR-NAME \
    --property hw:numa_nodes=FLAVOR-NODES \
    --property hw:numa_cpus.N=FLAVOR-CORES \
    --property hw:numa_mem.N=FLAVOR-MEMORY
```

```
$ openstack flavor create --ram 2048 --disk 1 --vcpus 4 m1.numa
$ openstack flavor set m1.numa \     # configure guest node 0
        --property hw:numa_cpus.0=0,1 \
        --property hw:numa_mem.0=2048
$ openstack flavor set m1.numa \     # configure guest node 1
        --property hw:numa_cpus.1=2,3,4,5 \
        --property hw:numa_mem.1=4096
$ openstack flavor show m1.numa
```

RED HAT® OPENSTACK® PLATFORM 10

redhat.

CPU Pinning

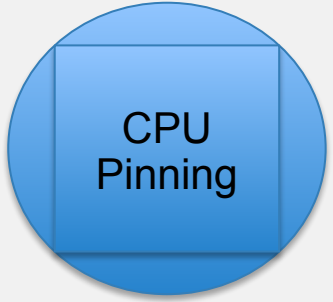

- vcpu_pin_set = 0,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19
- reserved_host_memory_mb = 1024
- NovaSchedulerDefaultFilters= “RamFilter, ComputeFilter, AvailabilityZoneFilter, ComputeCapabilitiesFilter, ImagePropertiesFilter, PciPassthroughFilter, NUMATopologyFilter, AggregateInstanceExtraSpecsFilter”
- hw:cpu_policy=shared | dedicated
- Host_aggregates is used to separate hosts for CPU pinning

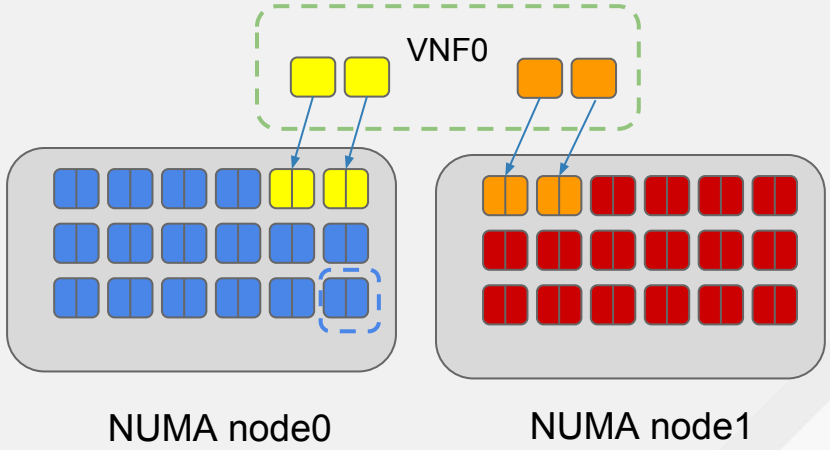

RED HAT OPENSTACK PLATFORM 10

redhat.

Huge pages

- hw:mem_page_size=large|small|any|<custom page size in KiB>
- X86_64: custom page size (2048KB, 1048576KB).
- No memory overcommit.
- Different Hugepage size for VNFs.

I/O based PCIe scheduling

- Intelligent NUMA node placement for guests.
- Affinitizes guest to NUMA node as allocated PCI device.
- Works with instances with NUMA topology defined.

RED HAT® OPENSTACK® PLATFORM 10

NFV SR-IOV Compute deployment

Depending on the VNF, it will be connected to one or more fabric network, typically:

- One fabric for control plane
- One fabric for user plane (end user datas: mobile traffic, web, ...)

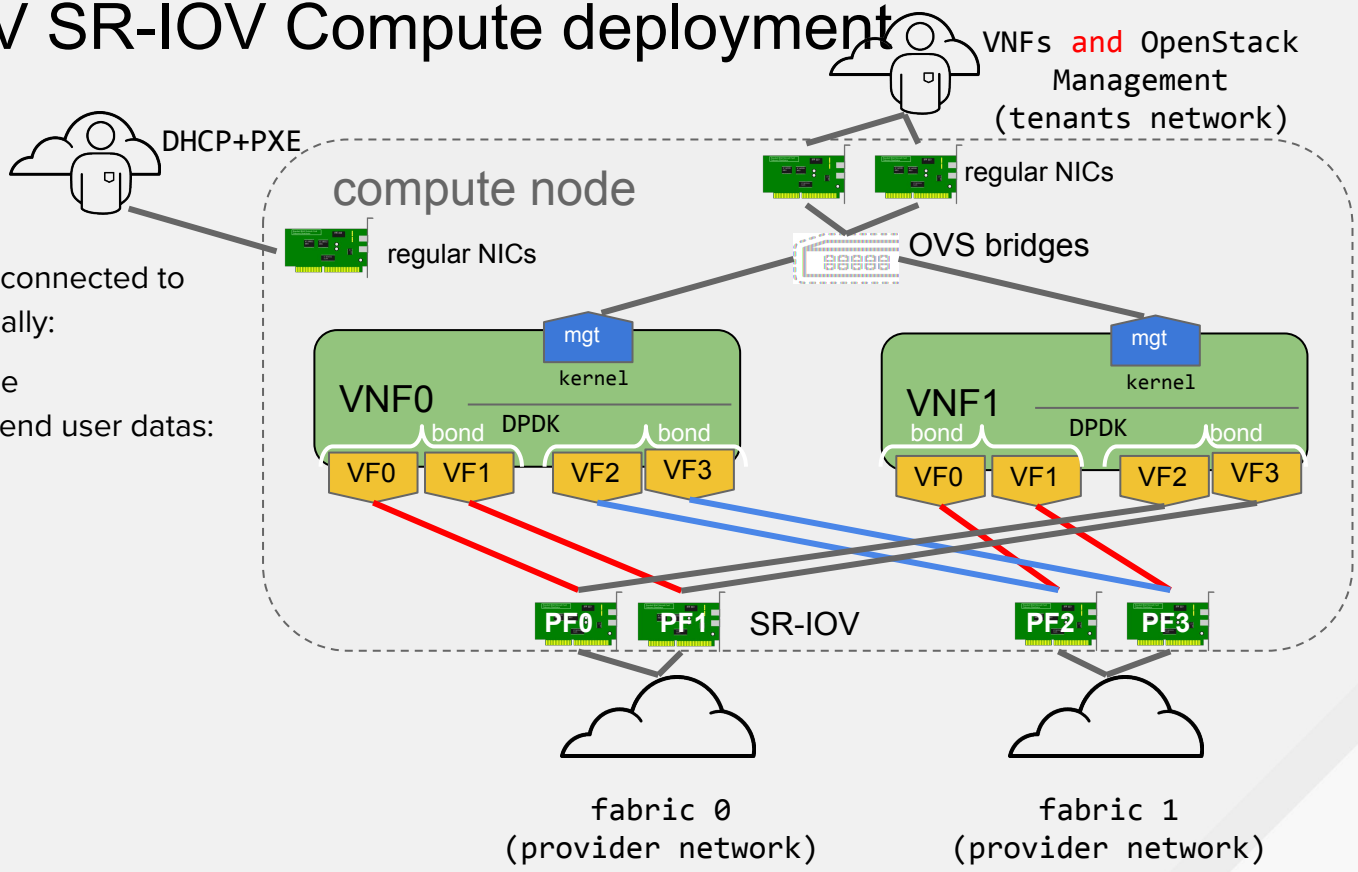

NFV OVS-DPDK Compute deployment

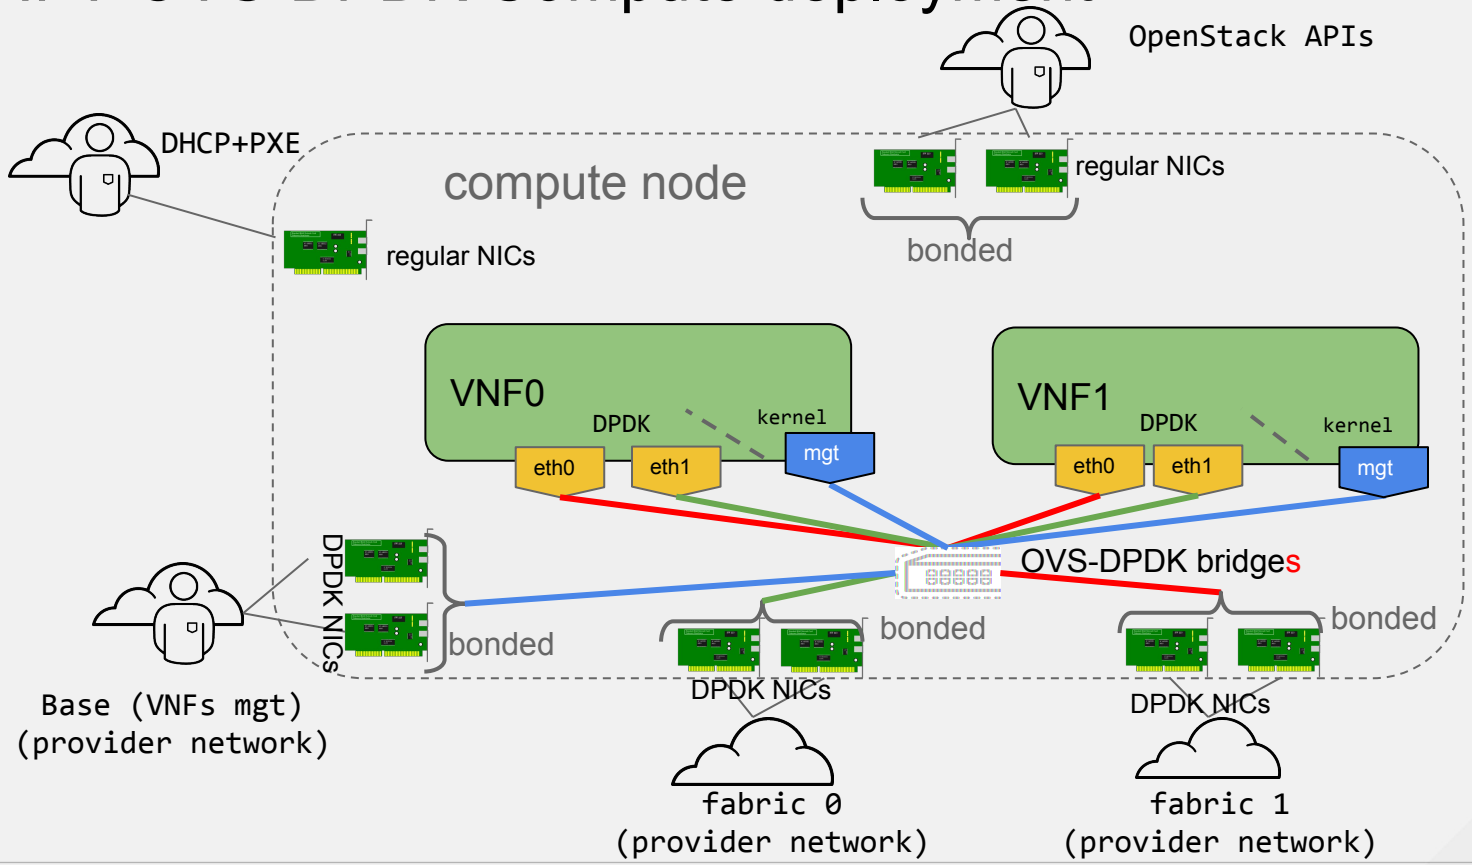

Red Hat OpenStack Platform - Open, Production-Ready Framework for the Cloud

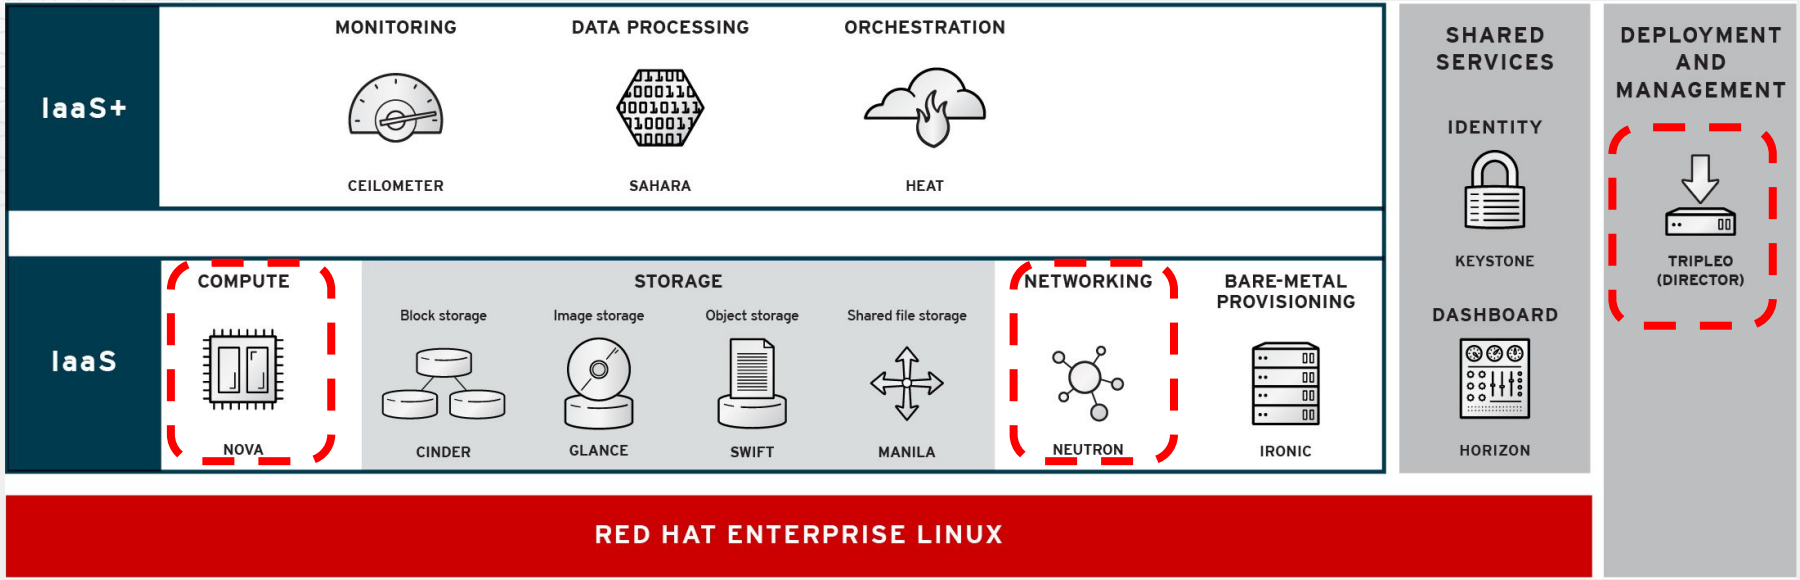

- Built on, integrated with Red Hat Enterprise Linux
- Co-engineered with Red Hat Enterprise Linux
- Tested and Certified by Red Hat for production workloads.

RED HAT® OPENSTACK® PLATFORM

redhat.

redhat.

THANK YOU

plus.google.com/+RedHat
linkedin.com/company/red-hat
youtube.com/user/RedHatVideos

facebook.com/redhatinc
twitter.com/RedHatNews