

Red Hat Enterprise Linux OpenStack Platform:

Life-cycle, roadmap, & partner ecosystem

Stephen Gordon (@xsgordon),
Sr. Technical Product Manager, OpenStack
Red Hat

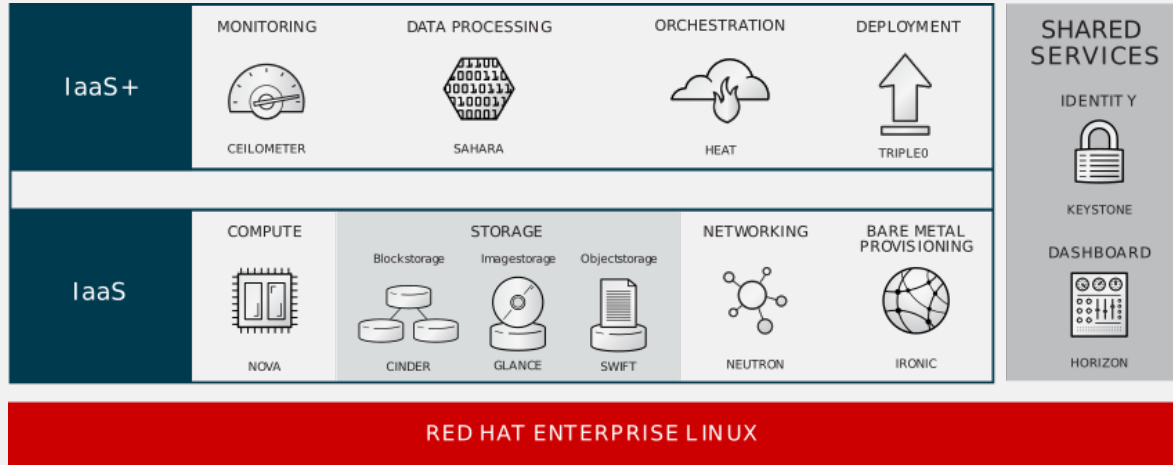
Bill Mason
Director, Partner Strategy and Programs
Red Hat

Agenda

- OpenStack Refresher
- Life-cycle
- Roadmap
- Partner Ecosystem

OpenStack Refresher

OpenStack Architecture



- OpenStack is made up of individual autonomous components
- All of which are designed to scale-out to accommodate throughput and availability
- OpenStack is considered more of a **framework**, that relies on **drivers** and **plugins**
- Largely written in Python and is heavily dependent on Linux

The OpenStack Community

- OpenStack community releases a new major version every **six months**
- Previous **upstream** version codenamed **Juno** (released **October 2014**)
- Current **upstream** version codenamed **Kilo** (released **April 2015**)
- Contributions to the project come from both corporate and non-corporate entities
- Ships source-code (**trunk code**) - no certifications, support, or packages
- In terms of scale, the **Kilo** release had:
 - **21,000+** code commits
 - **1,500+** individual contributors
 - **150+** organisations contributing



Red Hat's OpenStack Contributions


- Red Hat is a **Platinum Founding** member of the **OpenStack Foundation**
- Red Hat has been a leading corporate contributor to **Grizzly, Havana, Icehouse, Juno, and Kilo**
 - Commitment is **broad** across **all components**
 - Consistent leadership across individual **commits, code reviews, lines of code, and bugs fixed**
- Red Hat is also leading development in underlying **dependencies**
 - **Linux** kernel
 - **KVM** and **QEMU** hypervisor components
 - **Libvirt** - virtualization library
 - **Open vSwitch** - multilayer virtual switch
 - **Ceph** - since acquisition of Inktank

Why does this matter?

- Proves that Red Hat has the skills and resources to-
 - **Support** customers through the **entire** OpenStack platform and its dependencies
 - **Engineer** hot-fixes for customers, back-port security/stability fixes
 - Drive new customer **features** and requests
 - Influence the **strategy** and **direction** of the OpenStack project
- Red Hat is heavily focused on “upstream first”
 - All patches are contributed to the community for **adoption, integration**, and **testing**, first
 - Red Hat will not fragment their OpenStack distribution and support **forks**

RHEL OpenStack Platform Life-cycle

RHEL OpenStack Platform Life-Cycle



<http://www.redhat.com/en/technologies/linux-platforms/openstack-platform>

Life-Cycle Definitions

Description	1-year Production 1	2-year Production 2
Major bug fixes	yes	yes
Minor bug fixes	yes	yes*
Security errata	yes	yes
Bug fix errata	yes	yes
New features (backport potential)	yes*	no
Updated install image	yes	no
Partner enablements additions	yes	no

* need full product + engineering review prior to commitment

RHEL OpenStack Platform Roadmap

RHEL OpenStack Platform 6

IaaS+

MONITORING



CEILOMETER

DATA PROCESSING




SAHARA

ORCHESTRATION



HEAT

DBAAS
TECH PREVIEW



TROVE

DEPLOYMENT
TECH PREVIEW



TRIPLEO


SHARED
SERVICES

IDENTITY



KEYSTONE


DASHBOARD



HORIZON

IaaS

COMPUTE



NOVA

STORAGE

Block storage



CINDER

Image storage



GLANCE

Object storage



SWIFT

NETWORKING



NEUTRON

BARE METAL DRIVER
TECH PREVIEW



IRONIC

RED HAT ENTERPRISE LINUX

RHEL OpenStack Platform 7

RHEL OpenStack Platform Director:

- **Deployment (“TripleO”)** - Fully supported
- **Baremetal (“Ironic”)** - Fully Supported

Technology Previews:

- **Database-as-a-Service (“Trove”)** - Technology Preview
- **DNS-as-a-Service (“Designate”)** - Technology Preview
- **File Share Service (“Manila”)** - Technology Preview
- **Operational tools** (logging, monitoring, etc) - Technology Preview
 - - Centralized Logging: fluentd + Elasticsearch + Kibana
 - - Availability Monitoring: sensu + rabbitmq + redis + uchiwa

Roadmap Focal Areas

- Deployment Management
- Containerization
- High Availability
- Core Infrastructure Updates

Deployment Management

RHEL OpenStack Platform director

- Deployment management tool for RHEL OpenStack Platform
- Based on **TripleO**
 - Upstream friendly, upstream first
- 3 focus areas:
 - **Deployment Consistency**
 - **Lifecycle Management**
 - **Operational Visibility**

Deployment Consistency

- Best practices & reference architectures
 - Automation and repeatability
 - RHEL OpenStack platform director was created based on expertise from the field
 - Lessons learned from previous deployment tools
 - Reference architectures with certified hardware make deployment easier
- Hardware performance and validation testing
 - AHC (Automated Health Checks)
 - Hardware detection
 - Performance information
 - Black sheep detection
- Automated tests during and after deployment
 - Find and troubleshoot problems faster

Lifecycle Management

- Deployment is just the first step, most of the action happens afterwards
 - Add and remove capacity
 - Deploy critical updates
 - Upgrade to new OpenStack versions
- Automation is a must
 - API first, used by both the CLI and GUI and allow for better integration with external tools
 - Automated hardware detection and performance tests
 - Automated functional tests to validate the deployment as early as possible
 - Orchestrated upgrades, easier to keep closer to the newest features

Operational Visibility

- Important for troubleshooting and system status
 - Are my nodes behaving correctly?
 - Do I have enough resources?
 - What caused an operational failure?
- Operational tools
 - Log aggregation and search
 - Core service and infrastructure availability monitoring
 - Performance monitoring

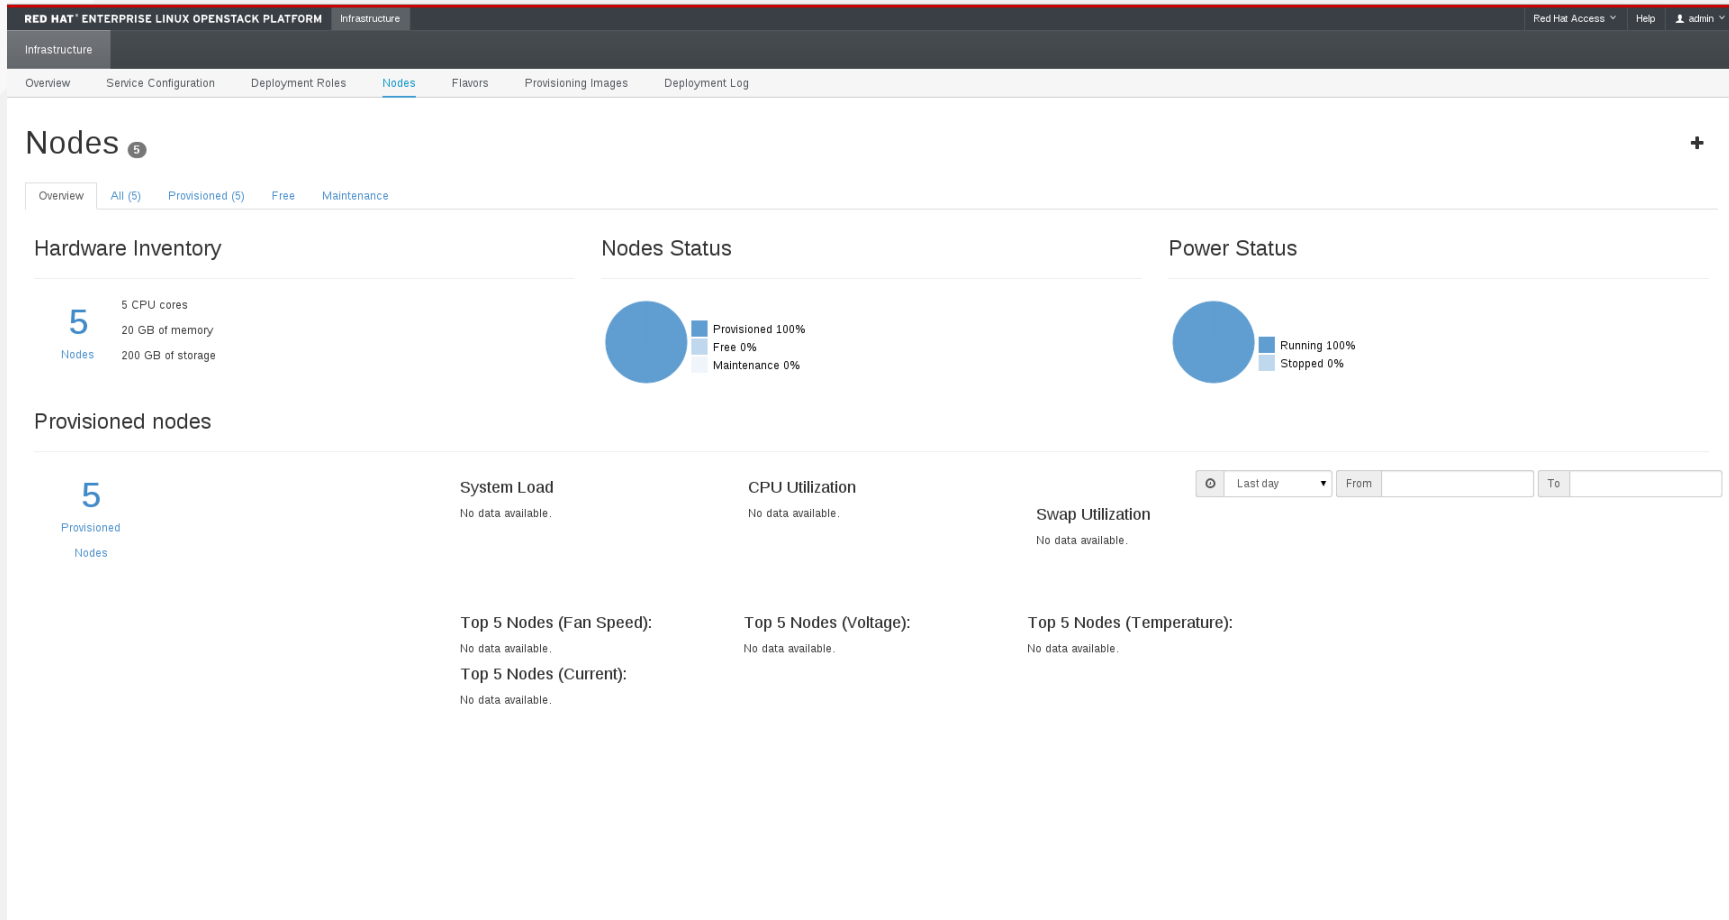
RHEL OpenStack Platform director demo

```
stack@instack ~$
-rwxr-xr-x. 1 stack stack 5027392 Jun 17 17:23 deploy-ramdisk-ironic.kernel
-rw-r--r--. 1 stack stack 150136570 Jun 17 17:39 discovery-ramdisk.initramfs
-rwxr-xr-x. 1 stack stack 5027392 Jun 17 17:39 discovery-ramdisk.kernel
drwxrwxr-x. 2 stack stack 4096 Jun 21 05:16 .instack
-rw-r--r--. 1 stack stack 15643 Jun 21 02:27 instackenv.json
drwxr-xr-x. 3 stack stack 4096 Jun 21 05:26 .novaclient
-rw-r--r--. 1 stack stack 36760178 Jun 17 17:45 overcloud-full.initrd
-rw-r--r--. 1 stack stack 908509184 Jun 17 17:48 overcloud-full.qcow2
-rwxr-xr-x. 1 stack stack 5027392 Jun 17 17:45 overcloud-full.vmlinuz
drwxr-xr-x. 2 stack stack 4096 Jun 21 05:26 .ssh
-rw-r--r--. 1 stack stack 221 Jun 21 05:26 stackrc
-rw-r--r--. 1 stack stack 4579 Jun 21 02:27 undercloud.conf
-rw-r--r--. 1 stack stack 1328 Jun 21 05:16 undercloud-passwords.conf
[stack@instack ~]$ openstack overcloud image upload
ERROR: openstack Missing parameter(s):
Set a username with --os-username or OS_USERNAME
Set an authentication URL, with --os-auth-url or OS_AUTH_URL
Set a scope, such as a project or domain, with --os-project-name or OS_PROJECT_NAME
[stack@instack ~]$ source stackrc
-bash: stackrc: No such file or directory
[stack@instack ~]$ source stackrc
[stack@instack ~]$ openstack overcloud image upload
[stack@instack ~]$ openstack image list
+-----+-----+
| ID | Name |
+-----+-----+
| 500959f5-51ef-4a4d-a50a-b10d0d69410c | bm-deploy-ramdisk |
| a90e7fbb-36d0-4203-86ec-a80e134b5bac | bm-deploy-kernel |
| 62b6ec4a-4b52-4502-9ea4-22de41806396 | overcloud-full-initrd |
| 09f7bd53-7743-46ae-b16e-f2d70a23b8ae | overcloud-full |
| e92098ff-54da-4d9b-9659-c1e41a1ec939 | overcloud-full-vmlinuz |
+-----+-----+

[stack@instack ~]$
[stack@instack ~]$
[stack@instack ~]$
[stack@instack ~]$ vi instackenv.json
[stack@instack ~]$ jq '.nodes[].memory=null | .nodes[].disk=null | .nodes[].arch=null | .nodes[].cpu=null | .nodes' instackenv.json > nodes.json
[stack@instack ~]$ vi nodes.json
[stack@instack ~]$
[stack@instack ~]$
[stack@instack ~]$ openstack baremetal import --json nodes.json
/usr/lib/python2.7/site-packages/novaclient/v1_1/_init_.py:30: UserWarning: Module novaclient.v1_1 is deprecated (taken as a basis for novaclient.v2). The preferable way
to get client class or object you can find in novaclient.client module.
warnings.warn("Module novaclient.v1_1 is deprecated (taken as a basis for "
```

Registering Nodes

Soon on <https://www.youtube.com/redhat> or right now at our booth



RHEL OpenStack Platform director 7.0

- APIs for deployment and management
 - Ironi, Nova, Glance, Heat, Tuskar
- CLI and GUI operation
- Per node Automatic Health Check (AHC)
 - aka black sheep detection
- Automated sanity checks for the deployment
 - Tempest run at the end to validate cloud as whole
 - Full and custom test suite
- Ready state configuration for selected hardware
 - RAID
 - Network
 - BIOS
- Automatic HA configuration using Pacemaker

RHEL OpenStack Platform director 7.x

- Accelerated release cadence
 - Released more often than core components
 - Reduced cycle time between field feedback and improvements
- Certified partner integration
- Automated orchestrated upgrades from 7.0 onwards
 - Upgrades should be a non-event, making it easier to use the newest features
- Automated HA fencing setup
- Distributed Automatic Health Check (DAHC)
 - Troubleshoot networking issues before deployment
- Automated tests in between deployment steps
 - Detect deployment issues earlier and make troubleshooting easier
- Images and templates customization
 - Increased deployment flexibility, adapting the director to customers' needs

RHEL OpenStack Platform director 8

- Ready state configuration for any hardware supporting IPMI
 - Increased options of deployment
 - Easier to repurpose hardware
- Alternative automatic HA configuration using Keepalived
 - More HA options, not everybody uses Pacemaker
- Container based deployment
 - Lots of innovation on the container front
 - Introducing the technology early in incremental steps

Containers

In the works

Supporting two use cases:

- Deploying OpenStack itself
- Deploying workloads on top of OpenStack

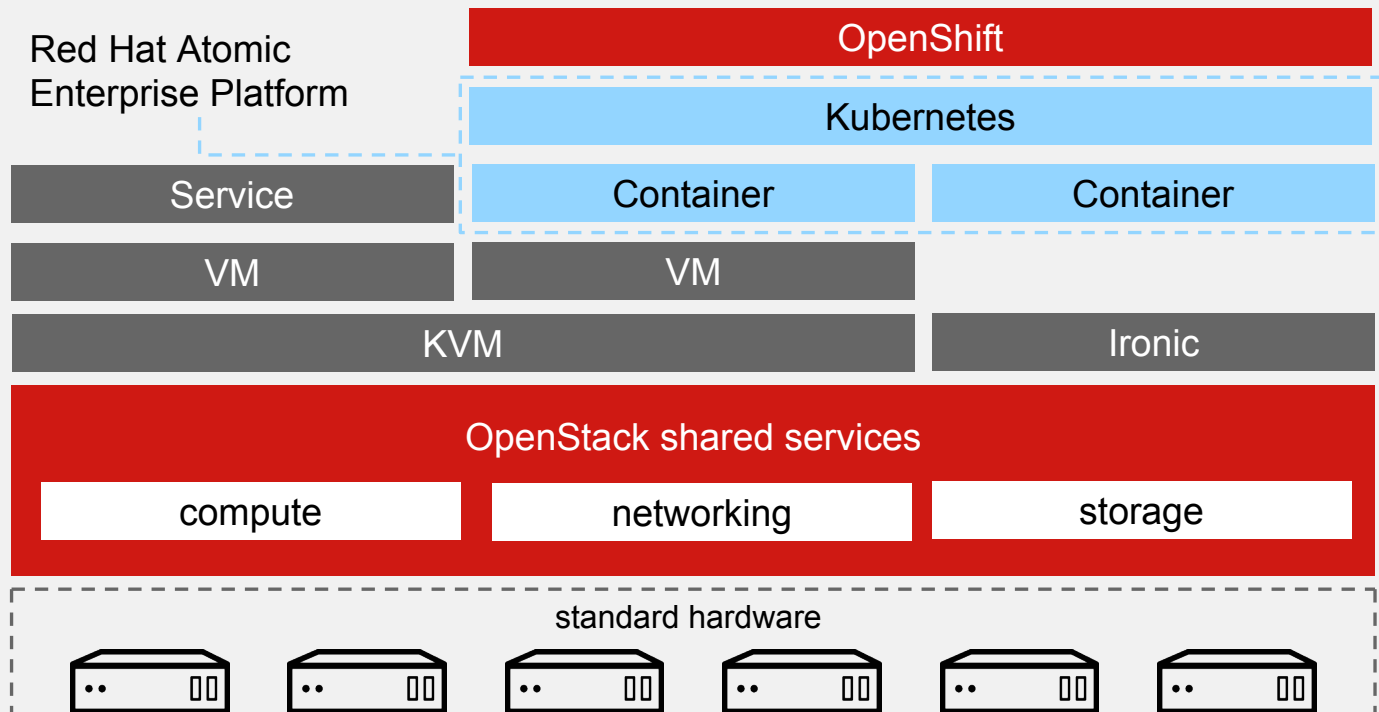
Focused on:

- **Kolla** for OpenStack deployment
- **heat-kubernetes** for deploying Atomic on OpenStack and configuring a Kubernetes cluster for user workloads
- **Kubernetes integration** with OpenStack infrastructure services (e.g. Networking and Storage).
- Kubernetes **cloud provider plug-in** for OpenStack

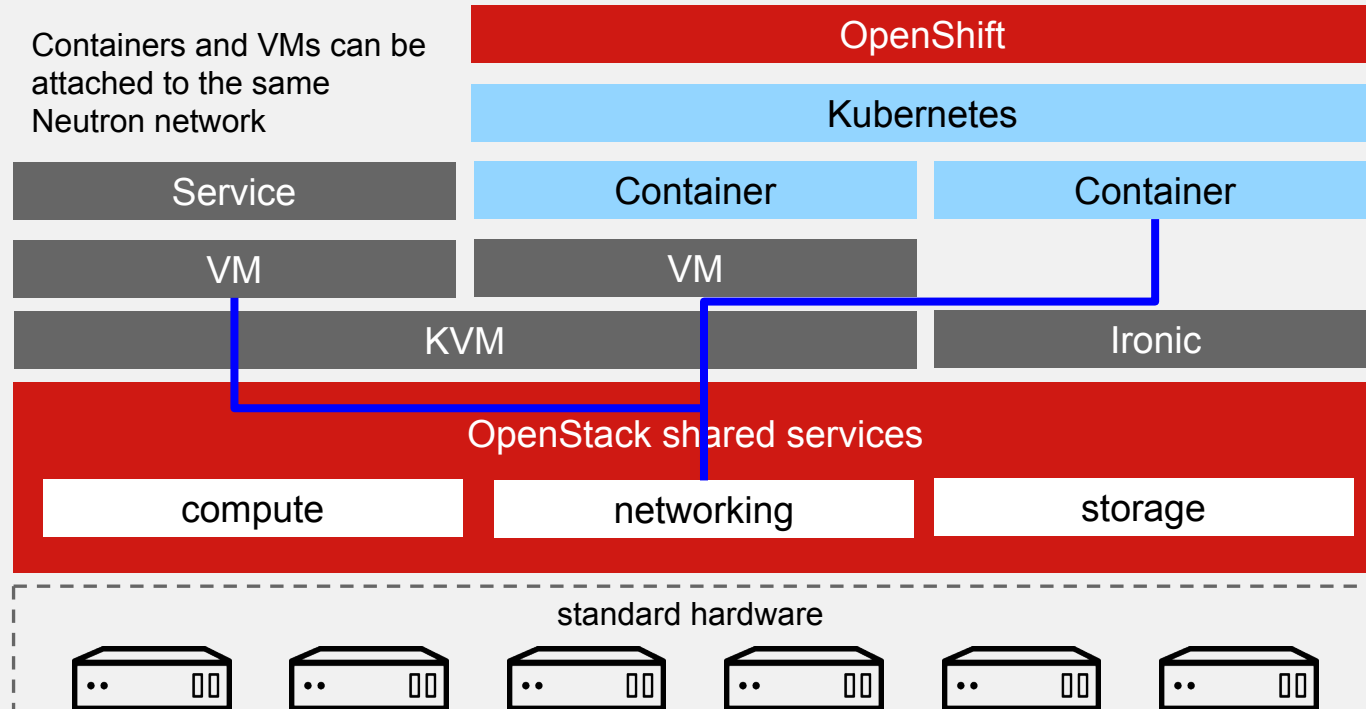
Architectural tenets

- 1) **Technical independence:** Ensure that containers are defined such that they remain independent of the underlying infrastructure. Containers must continue to be portable across host environments.
- 2) **Avoid redundancy:** Limit redundancies where possible to minimize performance and other resource hits. This includes limiting the number of layers between the container and the hardware.
- 3) **Contextual awareness:** Allow containers to easily take advantage of OpenStack shared services beyond compute (i.e. networking and storage). To do this, Red Hat Atomic Enterprise (and other Red Hat container offerings) must be context aware.
- 4) **Simplified management:** Simplify management by delivering a holistic, integrated view across platforms.

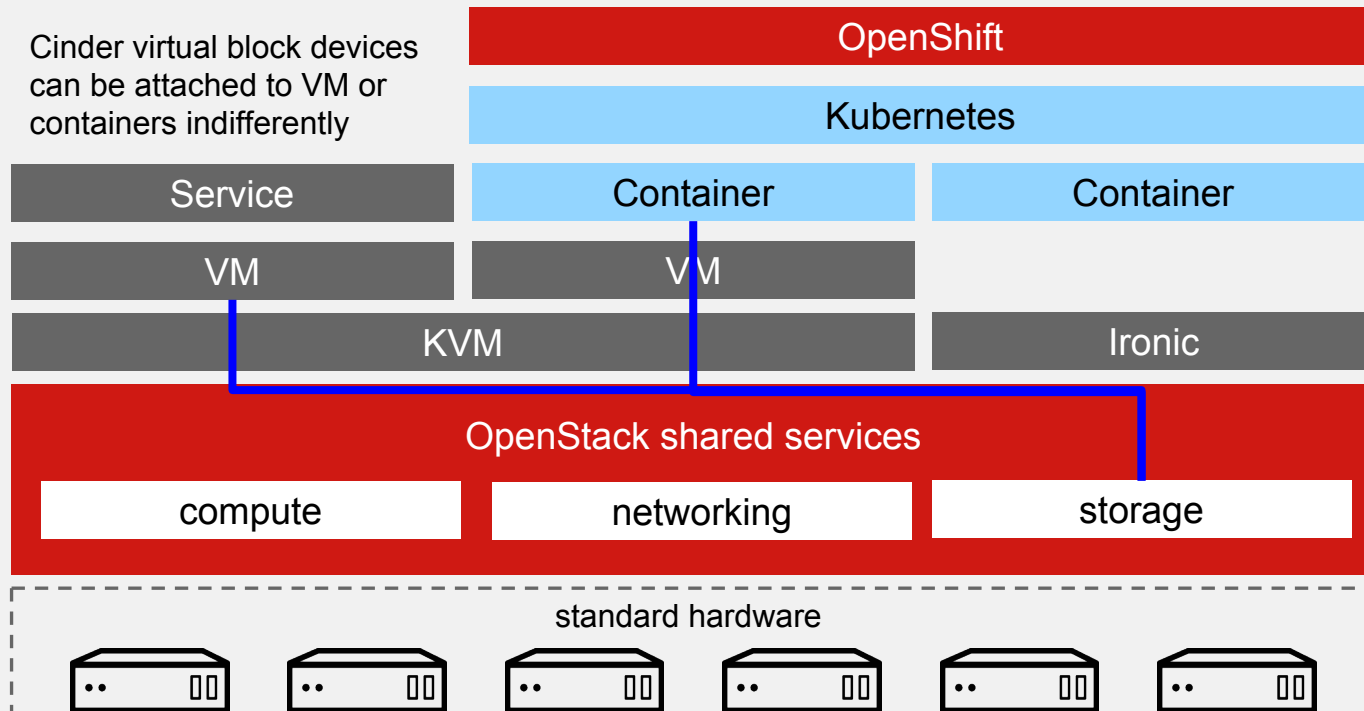
Convergence - Physical and application abstraction



Convergence - Physical and application abstraction

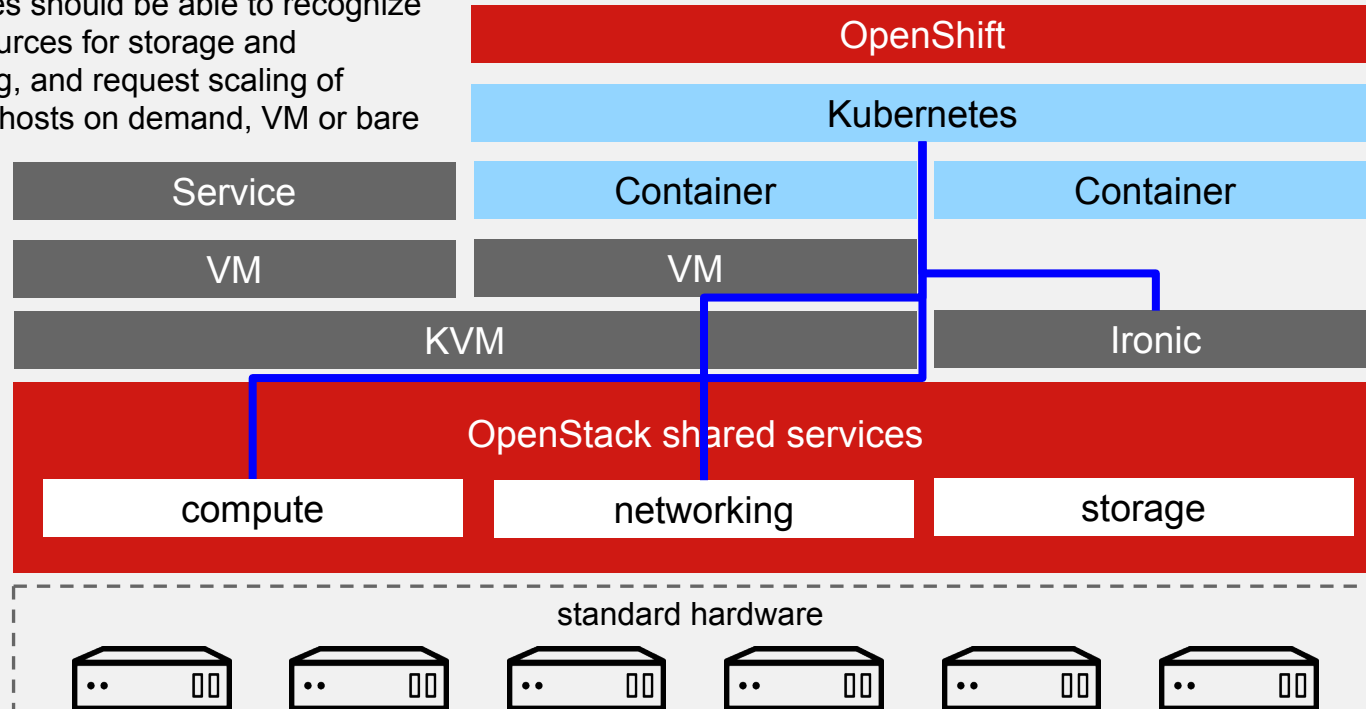


Convergence - Physical and application abstraction



Convergence - Physical and application abstraction

Kubernetes should be able to recognize local resources for storage and networking, and request scaling of container hosts on demand, VM or bare metal



High Availability

High Availability - Current State

- Have supported integration with RHEL High Availability Add-On since RHEL OpenStack Platform 5
 - **Pacemaker** - resource mgr for active/active & active/passive services
 - **HAProxy** - Load balancing active/active services
 - **Keepalived** - Load balancer availability
- Two architectures: **Pacemaker/HAProxy** & **Keepalived/HAProxy**
- All core OpenStack services **active/active** except **cinder-volume**
 - Red Hat developers working upstream on enhancing cinder-volume to support active/active safely

High Availability - Roadmap

Compute Node & Instance High Availability:

- Uses Pacemaker to monitor all Compute Nodes, detect faults, and recovers instances to other running Nodes
- Uses vanilla Nova functionality for node disable and evacuation
- Can detect failures from kernel panics, loss of connectivity or other failure
- Works with most baseboard management controllers

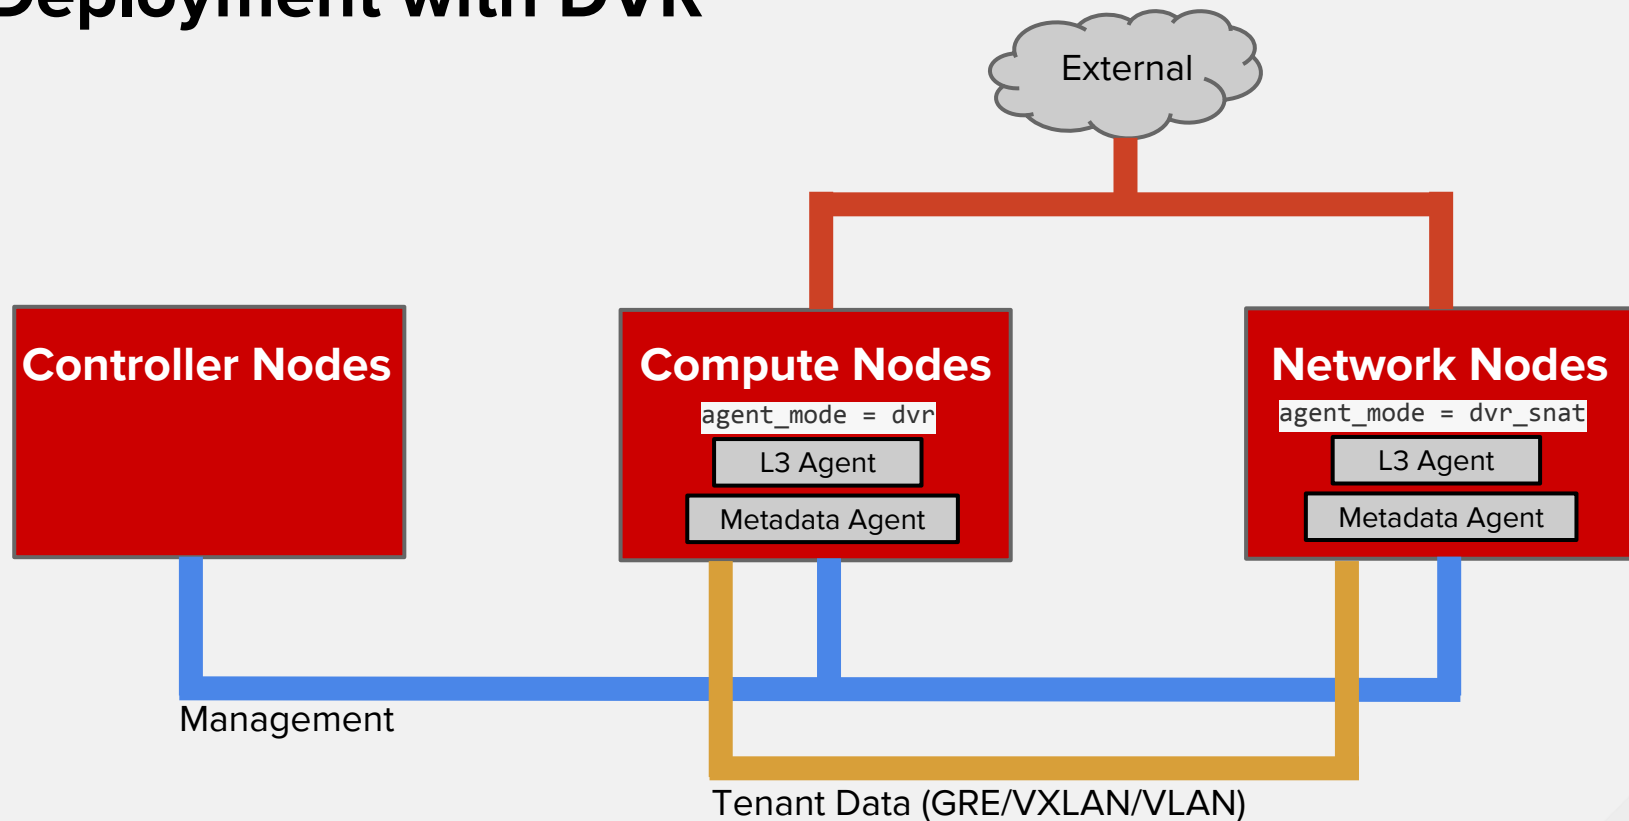
Core Infrastructure Updates: Networking

RHEL OpenStack Platform 7 - IPv6 enhancements

- Multiple IPv6 prefixes on a network
 - It is now possible to attach several IPv6 subnets to a network
 - When the subnet type is either SLAAC or DHCPv6 stateless, one IPv6 address from each subnet will be assigned to the Neutron port
- Extra DHCP options
 - It is now possible to specify extra DHCP options (e.g., DNS) for both DHCPv4 and DHCPv6
 - Important for dual-stack designs, where a VM is assigned with both IPv4 and IPv6 addresses on the same port
- IPv6 external router support
 - The virtual router can automatically learn its default gateway information via SLAAC; or
 - The default route can be manually set by the operator using a new `ipv6-gateway` option

Note: there is no NAT or floating IP support with IPv6. Tenants are expected (and trusted) to configure globally routable prefixes

Deployment with DVR



RHEL OpenStack Platform 7

- VLAN networks with DVR (Technology Preview)
 - Distributed Virtual Router can now be configured to serve VLAN networks, in addition to GRE or VXLAN
 - DVR allows to distribute east/west routing and floating IPs to the Compute nodes
- LBaaS v2 API - RHEL OpenStack Platform 7
 - A new version of the API, including support for TLS/SSL termination
 - Customers can deploy with the HAProxy service plugin/agent (default in RHEL OpenStack Platform) or with certified plugins for some of our partners

RHEL OpenStack Platform 8 and beyond

- Role-based Access Control (RBAC) for networks
- Neutron quality of service (QoS)
- Pluggable IPAM
- IPv6 Prefix Delegation
- L3 HA + L2 Population support
- L3 HA support for IPv6
- Stateful OVS firewall
- VLAN trunking into a VM

Core Infrastructure Updates:

Storage

RHEL OpenStack Platform 7

- Volume Backup API was extended to support snapshot based (“incremental”) backups, where the volume can remain online and in-use for the duration of the operation.
- Image conversion during image import process
 - Useful to unify stored image types and use a type that works better with the hypervisor and the storage backend. (supported formats are: raw <-> qcow2)
- Introspection of images
 - A new task has been added to Glance’s v2, which makes it possible for introspecting image’s metadata and populate it.
- Ability to deactivate or “hold” images
 - The admin is now able to put the image “on hold” preventing instances from being built with it until it has be properly examined.
- Object Storage Composite Tokens & Service Accounts
 - Allow other OpenStack services to store data in Swift on behalf of a client so that neither the client nor the service can update the data without both parties consent.

RHEL OpenStack Platform 8

- Generic Volume Migration
 - Add the ability to migrate volumes of drivers that don't support iSCSI such as Ceph RBD.
- Import/Export Volume snapshots
 - Allows to import volumes snapshot from one Cinder to another.
 - Allows to import "non" openstack snapshots already on a backend-device. Where, export snapshots should work the same way as export volumes.
- Image Artifact Repository (ongoing)
 - Extend Glance functionality to store not only the VM images
 - Glance to become a catalog of such artifacts, providing capabilities to store, search and retrieve their artifacts, their metadata and associated objects.
- Support snapshotting ephemeral disks stored in Ceph RBD
 - By clean up unused parents of clones on deletion.
- Object Storage
 - Erasure Coding (tech-preview in RHEL OpenStack Platform 7)

Core Infrastructure Updates:

Compute

Compute - RHEL OpenStack Platform 7

- Iteration on **Enhanced Platform Awareness (EPA)** features developed in “Kilo” and backported to RHEL OpenStack Platform 6 release:
 - vCPU pinning
 - Large pages (2M and 1G)
 - vCPU, RAM, and I/O device NUMA awareness
- Baremetal (“**Ironic**”) driver supported, with caveats!
- Support for quiescing file-systems during image snapshot using QEMU guest agent

Compute - RHEL OpenStack Platform 8

- RHEL for Real-time integration (incl. **real-time KVM**)
- **virtio** and **vhost-user** performance enhancements
- Hardening of instance live-migration
 - **migration monitoring** capabilities
 - support evacuate/migrate on instances with enhanced platform awareness
- Enhancements to Nova API to better support virtual machine HA
 - Allow external tools (e.g. Pacemaker) to tell Nova that a node is being fenced.

Partner Ecosystem

Red Hat Connect for Technology Partners

- RHC4TP is an open, collaborative ecosystem connecting technology companies with the tools and resources to build certified solutions for the Red Hat family of products
- The program provides:
 - Early access to pre-release products
 - Best practices, architectural guidance, SMEs
 - Certification

Comprehensive Partner Ecosystem

Red Hat Connect for Technology Partners




OEMs and IHVs



ISVs



System Integrators



Cloud Service Providers



Managed Service
Providers



Channel Partners

- Over 350+ members
- Over 900 certified solutions
- Over 4,000 Red Hat Enterprise Linux certified compute servers
- Over 13,000 applications available on Red Hat Enterprise Linux
- Large catalog of Windows certified applications

Red Hat Enterprise Linux OpenStack Platform Certification

- Certification Value Proposition
 - Partner testing and validation on Red Hat products
 - Collaborative technical support issue resolution
 - Certifications for:
 - Nova compute
 - Cinder storage
 - Neutron networking

OpenStack Key Strategic Alliances



Cisco and Red Hat



- Cisco Unified Computing System Integrated Infrastructure for Red Hat Enterprise Linux OpenStack Platform (USCO)
- Cisco Validated Designs
- Application Centric Infrastructure and Group Based Policy
- Cisco Intercloud

Dell and Red Hat



- Jointly developed OpenStack cloud solutions with tested and proven reference architectures
- Co-engineered to be quality engineered, secure and validated
- Accelerated private cloud deployment shortens time to business value

NetApp and Red Hat



- Pre-integrated / validated OpenStack solutions
- Start small and expand over time
- Streamlined support more quickly resolves technical support issues
 - Experience
 - Collaboration
 - Cross-training

Intel and Red Hat



- Extending the OpenStack On-Ramp to the Enterprise program
- Focused on ease of deployment
- Rolling upgrades with zero downtime
- High availability of services / tenants
 - Evacuate hosts impacted by failure
 - Move workloads to consolidate host usage
 - Enables traditional and cloud workloads to co-exist

RED HAT SUMMIT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.

Questions?

Don't forget to submit feedback using the Red Hat Summit App!

Stephen Gordon

sgordon@redhat.com / [@xsgordon](https://twitter.com/xsgordon)

Bill Mason

bimason@redhat.com

Coming up!

- OpenStack Engineering Roadmap - **Ballroom A @ 3:40 PM.**