

RED HAT ENTERPRISE LINUX OPENSTACK PLATFORM: OpenStack Development Roadmap

Tim Burke – Vice President of Linux, Cloud & Storage Engineering Perry Myers – Senior Director of OpenStack Engineering Hugh Brock – Senior Manager of OpenStack Installation and Management



Agenda

- Deployment Success Factors Tim Burke
- Feature Enhancements Perry Myers
- Management and Usability Hugh Brock



Red Hat Enterprise Linux OpenStack Platform: Themes and Priorities

Tim Burke Vice President of Linux, Cloud, Storage Product Development





Appropriate expectations – Success factors:

- Agile workloads (cattle)
- Subset of data center
- Team culture receptive dev + ops
 collaboration for agility
- Start realistic organic growth
- Standardize





Is OpenStack a science experiment? [{]



Appropriate expectations – Misfit scenarios currently:

- Traditional IT. Legacy apps i.e. database, ERP, static environment (pets)
- Static performance tuned, determinism focused
 - Better suited for RHEV
- Mandates
- Expectation of seamless rolling upgrade
- Roll your own OpenStack parts



Example Success Deployments

 X1 DVR – large # of incremental apps – i.e. weather, sports, events

Walmart 🔆

- provisioning 11,000 stores in 27 countries
- Covered "black Friday"

D Bank

- Transactional infrastructure self-service, automation
- Consolidation
 Source: OpenStack Summit 2015



Red Hat Customer Types – Success Stories

- Telco/Network net / messaging platforms, hosted cloud providers
- Big box retailers
- Financial services
- Government
- Healthcare / Insurance
- Video production

So yes, RHEL-OSP is production ready today, for appropriate workloads



Do you need to be a rocket scientist to deploy OpenStack?

Usability is a key focus of Red Hat's OpenStack efforts

- RHEL-OSP installer rewrite
- Partner ecosystem connect.redhat.com
- Red Hat's product portfolio integration product solution stack
- Reference architectures:
 - http://red.ht/1Nd64ha
 - RHCI RHEL, RHEL-OSP, RHEV, CloudForms
- Training & certification
- Cloud services consulting





RHEL-OSP Partners Alliances

·IIIII CISCO

InterCloud

Delivers to multiple large Telco providers



UCSO / FlexPod UCS reference architectures *Telefonica* NFV Reference Lab

NetApp[®]

Dell Cloud Solutions Reference Architecture





Enterprise Readiness Focus
Upstream aligned innovation





RHEL-OSP Co-engineered







Red Hat Cloud Solution Stack





Solution Stack Success Story



Background

- FICO score consumer credit risk
- Transitioning to analytics services into reusable components
- Transforming from VMW based to services
- Proprietary cloud didn't meet

Red Hat Solution

- OpenShift PAAS
- Ceph based storage
- RHEL-OSP infrastructure



Red Hat Enterprise Linux OpenStack Platform: Roadmap and Feature Enhancements

Perry Myers Senior Director of OpenStack Engineering

Evolving for the Enterprise

- 11 releases of upstream OpenStack
- 7 releases of Red Hat Enterprise Linux OpenStack Platform

OpenStack is being deployed, but has traditionally required a large organization with deep skills to deploy and operate.

The theme is now around preparing OpenStack for broad enterprise adoption.

Going beyond just large scale public cloud providers and into more traditional enterprise IT environments.

The next few slides detail the innovations that Red Hat engineers are driving in the OpenStack community, with help from our strategic partners.



"Always On" OpenStack

- RHEL HA Add On support since RHEL OpenStack Platform 5
 - Pacemaker resource mgr for active/active & active/passive
 - HAProxy load balancing active/active services
 - Keepalived load balancer HA
- Two architectures: Pacemaker/HAProxy & Keepalived/HAProxy
- Noteable HA enhancements in RHEL OpenStack Platform 6 and 7:
 - RabbitMQ heartbeat, tcp keepalives and oslo.messaging
 - L3 HA with IPv6 & VRRP (Technology Preview)
 - Galera initialization and failure recovery
- All core OpenStack services active/active except cinder-volume



Have some "non-cloudy" workloads?

- OpenStack == horizontally scalable, cloud-ready workloads
- But, what about non-cloud ready workloads?

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
- Detects failures from kernel panics, network or other failure
- Works with most baseboard management controllers



Getting to Online Upgrades

- Minor updates can be done without downtime, but...
- Major version updates require database schema updates (offline) and message-bus RPC versions to be synchronized
- Nova pioneered techniques in OpenStack Kilo to support online major upgrades (to OpenStack Liberty)
 - oslo.versionedobjects adoption provides tolerance for mixed awareness of old & new schemas during an upgrade
 - Online schema migration replace disruptive big-bang migration with ongoing expand-transform-contract cycle



High Performance Workloads

- Use case: Virtual Network Functions (VNF) for Telco/Network space (NFV)
- CPU-pinning
 - Allows driver to strictly pin instance vCPUs to host pCPUs
- libvirt hugepages support
 - Driver can now use large pages for backing instance RAM
- SR-IOV support with NUMA-aware scheduling
 - Intelligent NUMA node placement for instances assigned a host PCI device
- vhost-user awareness
 - Leverages a new QEMU feature to support efficient virtio-net I/O between instance and user-space vswitch



Securing OpenStack

- sVirt & SELinux enabled by default
 - Isolates guest instances using Mandatory Access Control (MAC)
 - Policies maintained and updated with each release and new projects
- Federated Identity
 - SAML based Single Sign-On (SSO) for CLI and Dashboard
 - Cross-cloud federation using K2K (Keystone to Keystone federation)
- Identity Management in RHEL provides enterprise level Identity capabilities
 - Centrally managed authentication and authorization
 - Kerberos Single Sign-On
 - One Time Passwords (OTP)
 - X.509 certificate issuance for TLS protection
 - Host-based access control (HBAC) and centrally managed sudo for securing access to your underlying cloud infrastructure
 - Integration with Active Directory using cross-realm trusts





Making the Operator's life easier

Shipped with RHEL OpenStack Platform 7 (Technology Preview)

Centralized Logging (EFK stack)

- Elasticsearch Distributed search engine
- Fluentd Unified log collection layer
- Kibana Dashboard/Visualization layer for Elasticsearch

Performance & Systems Monitoring

- Collectd System statistics collection
- Graphite Realtime graphing
- Grafana Metrics dashboard
- Sensu Systems monitoring framework
- Uchiwa Dashboard for Sensu





RHEL OpenStack Platform 7: New Components

- Fully Supported
 - Sahara Data Processing as a Service (Hadoop)

- Technology Previews
 - Trove Database as a Service
 - Manila Filesystem as a Service (NFS, glusterfs)
 - Designate DNS as a Service



But it's not just about features...





Testing, Testing, Testing

- Continuous testing of packaging for distributions:
 - Delorean/RDO Project
- Upstream Puppet CI
- Neutron:
 - Agents functional tests
 - Integration Testing Framework
- Voting Ceph CI in upstream gate for Nova/Cinder/Glance



Building an Ecosystem Violin Cloudian **PLUM**grid ARISTA kura Mellanox ERC ONA Networks MariaDB **RED HAT** NEC NetApp[®] **TERPRISE LINUX X**big switch **OPENSTACK PLATFORM** .1 CISCO. BROCADE (intel) D&LL) Symantec. One Convergence INFINIDAT CPLANE ETWORKS **nuage**networks FIRF

Red Hat Enterprise Linux OpenStack Platform Roadmap





Innovation Leaders



#redhat #rhsummit

Red Hat Enterprise Linux OpenStack Platform Roadmap



Red Hat Enterprise Linux OpenStack Platform 8: To Liberty and beyond ...



To Liberty and beyond...

- Hardening and moving new components to fully supported:
 - Trove
 - Manilla
 - Designate
- Hardening of instance live-migration
 - Monitoring of migration progress and avoidance of stuck migrations
- Completion of work to support Online Upgrades
 - Expanding oslo.versionedobject and online schema upgrade support across all projects beyond Nova
- RHEL Realtime with KVM integration



To Liberty and beyond: Net/Storage

- Neutron / Networking
 - OVS Security Groups
 - QoS
 - openvswitch + DPDK
 - SR-IOV support for live attach and physical function (PF) assignment
- Cinder
 - Volume migration between different storage types
 - Ceph integration enhancements



To Liberty and beyond: Availability

- Active/Active support for cinder-volume
- Layer 3 HA: Layer 2 Population & DVR Integration
- Enhancements to Nova APIs to better support Instance HA
- Integration of alternative HA architectures into Red Hat deployment tools (Keepalived/HAProxy)
- Added robustness for Compute Node evacuations
- Emit notifications charting progress of a host evacuation
- Explicit recording of evacuation in-flight to avoid guess-work on the status of left behind instances when the compute node is restarted
- External orchestration of long-running evacuations



To Liberty and beyond: Security

- TLS/SSL security configured and deployed out of the box
- Encrypted live migration
- No passwords! (Kerberos or X.509 client-certificate auth)
- Centralized OpenStack policy management
- Barbican for key management with RHEL Identity Management for PKI



References

- Blogs:
 - http://planet.rdoproject.org
 - http://redhatstackblog.redhat.com
- HA Arch: <u>https://github.com/beekhof/osp-ha-deploy/blob/master/ha-openstack.md</u>
- RDO Project: <u>http://rdoproject.org</u>



Red Hat Enterprise Linux OpenStack Platform: Management and Usability

Hugh Brock Senior Manager of OpenStack Installation and Management



Installation, Management, and Usability

- New for RHEL OpenStack Platform 7: "director," our deployment management tool for RHEL OpenStack Platform
- Also, new features in
 - Ironic, OpenStack's bare metal provisioning project
 - Heat, OpenStack's orchestration project
 - OpenStack Puppet Modules, now an integrated
 "big tent" OpenStack project



The Problem: Installing OpenStack

- Nothing fits "just right." Current OpenStack installers are either too proscribed (Mirantis Fuel, Red Hat "StayPuft"), too limited (PackStack, DIY approaches), or too difficult to use
- Here today, gone tomorrow. Current installers don't offer lifecycle management help, like update support
- Lone wolf: APIs for integration with higher level management tools are weak or not there at all



RHEL OpenStack Platform Director

Our mission:

- Provide a single tool that covers small (two-node) to large (100+ node) OpenStack installations, with real world flexibility, scriptable APIs, and wide integration
- Be around for Day 2 management with support for updates, upgrades, and scaling
- Make life easier for the operator deploying, managing, and upgrading OpenStack





My OpenStack Deployment & Edit Global Configuration + Register Nodes

Available Deployment Roles

? Compute	? Ceph-Storage	? Cinder-Storage
? Controller	? Swift-Storage	

Deployment Checklist

4 of 7 Steps Completed

Yerify and Deploy

Hardware

No flavor			







Nodes 👩

All (7)

Overview

Provisioned

Maintenance

Free (7)

Free

Filter Power On Nodes Power Off Nodes Delete Nodes								
	Node Name	CPU (cores)	Memory (MB)	Disk (GB)	Power Status	Actions		
	e42e1727-bdcb-4862-b844-ed158be3f9ca	1	4096	40	power off	Power On Node 👻		
	a732ade8-35f8-4900-941d-d24a02c6bf29	1	4096	40	power off	Power On Node 👻		
	6f57d927-2201-4127-b434-f1606082676c	1	4096	40	power off	Power On Node 👻		
	67f1a674-aaf2-4012-9399-704c5a3f5632	1	4096	40	power off	Power On Node 👻		
	1f7c44ca-eb6a-4592-b232-232bc70fae8e	1	4096	40	power off	Power On Node 👻		
	faf35f55-8338-425c-980f-36bf707db85b	1	4096	40	power off	Power On Node 👻		
	ce3fac78-365a-4422-95cc-8e5a7d5a0e50	1	1024	40	power off	Power On Node 👻		
Displa	ying 7 items							

#redhat #rhsummit



+

RED HAT	ENTERPRISE LINUX OPENS	TACK PLATFORM Infrastruc	ture	Crea	ating Flav	Red Hat Access ~ Help	1 admin
Infrastruct	ture				0		
Overview	Service Configuration	Deployment Roles	Nodes	Flavors	Provisioning Images	Deployment Log	
Flav	ors 🕤						+
1× Sugge	ested Flavor						*
Filter	۲	Delete Flavors					
D Na	ame Ar	chitecture	CPUs		Memory	Disk	Actions
CO	mpute x8	6_64	1		4GB	40GB	
С	ntrol x8	6_64	1		4GB	40GB	
sto	orage x8	6_64	1		4GB	40GB	
Displaying	3 iteras						
at #rhsumm	iit	Red Hat Enterprise	Linux OpenS	Stack Platf	orm Roadmap	39	🤗 r







My OpenStack Deployment

Register Nodes

Initialization Needed

Deployment Roles







What's in the box?

Real-world production deployments at scale, with third-party certified plugins including:

- Automated configuration of multiple networks with tight control over on-machine configuration (more below)
- Five standard roles out of the box: compute, controllers, Ceph storage, block storage, object storage
- Scale out compute, scale out Ceph storage
- DRAC/iLO/AMT support for hardware provisioning, NetApp, Dell, Cisco UCS, more



Plus...





What's in the box?

Integration with Red Hat CFME and Satellite:

- View tenant workloads in CFME and correlate them with load levels and events on the underlying cloud
- Migrate tenant workloads from VMWare, other platforms
- Inspect configuration and health of individual machines
- Scale out compute and Ceph storage nodes
- View and manage entitlement status and drift from Red Hat Satellite 6
- Pull updated OpenStack Platform content from Satellite



OpenStack Platform Integration

- RHCI drives installation via OSP director (Heat, Ironic)
- CloudForms presents operational view across clouds
- Satellite presents entitlement info for RHEL-OSP nodes and RHEL guests, provides content security via errata











What's in the box?

Help with common OpenStack deployment problems:

- Pre-deployment machine discovery and health checks
- Director applies Puppet config in five stages; you can add breakpoints between each stage to check everything is working as expected
- Post-deployment test run on the deployed cloud (uses Tempest)









What's in the box?

Complete, scriptable CLI with plain-text config:

- Deploy your small or large cloud from the CLI
- Store your deployment config in version control it's all YAML. Repeat your deployment as needed
- For advanced users: edit director Heat templates to add custom orchestration or apply additional Hiera data to the deployment



PLUS...



What's in the box?

Update and upgrade support:

- Use a Heat stack-update to apply selected RPM updates to your deployment security patches, hot fixes, etc.
- Starting on RHEL OpenStack Platform 7, we support automated upgrades to future versions with minimum downtime



THERE'S STILL MORE!



What's in the box?

Automated network configuration:

- Configure bonds, bridges, or individual NICs on your nodes automatically on deployment
- Connect the right services to the right networks with a straightforward YAML mapping template
- Use the default setup of five separate VLANs over a trunked bond on each node, or modify your config to suit
- Compatible with any Neutron plugin



What's next?

- Automated testing between Puppet deployment stages
- Distributed health checks we already check the health of each machine, now add checks for network and storage setup, bad cables, etc. before deployment
- Deploy OpenStack in containers
- Try it out today at:

http://www.rdoproject.org/RDO-Manager



What's next in Heat?

- Disconnected and nested stacks dramatically improve Heat scaling
- Memory management improvements increase limit on stack size and complexity
- "Convergence" self-healing stacks are on the schedule for Liberty and RHEL OpenStack Platform 8
- Lots of usability and error reporting fixes quickly inspect your stack, pull up actual deployment errors from nested operations, retry once the error is rectified



What's next in Ironic?

- "Ironic-Inspector" ramdisk runs on unprovisioned hardware, returns comprehensive hardware metadata.
 Optionally, runs performance checks on the hardware and returns those results as well
- Post-Inspector tagging can classify hardware for particular uses based on Inspector data
- DRAC driver, Intel AMT driver, BIOS and RAID config for DRAC
- Ironic state machine launched. When complete, will enable safe metal-to-tenant in your OpenStack cloud



Summary

RHEL-OpenStack director builds on upstream tools to provide a flexible, repeatable, long-lived solution for OpenStack deployers and operators.

For more information:

- RDO-Manager is our upstream project, you can find more information here: <u>https://www.rdoproject.org/RDO-Manager</u>
- RDO-Manager install documentation:

https://repos.fedorapeople.org/repos/openstack-m/docs/master/

• Email: rdo-list@redhat.com; IRC: irc.freenode.net/#rdo





LEARN. NETWORK. EXPERIENCE OPEN SOURCE.

. **red**hat

#redhat #rhsummit