Shifting the Data Center

Transitioning Red Hat IT to hybrid cloud infrastructure using OpenStack and Ceph Storage

Brian J. Atkisson, RHCA
Principal Architect

Matt Carpenter, RHCE
Principal Systems Administrator

May 9, 2019

Our approach to digital transformation

Roadmap illustrates conceptual "waypoints" that *guide* the work rather than specifying a detailed plan.

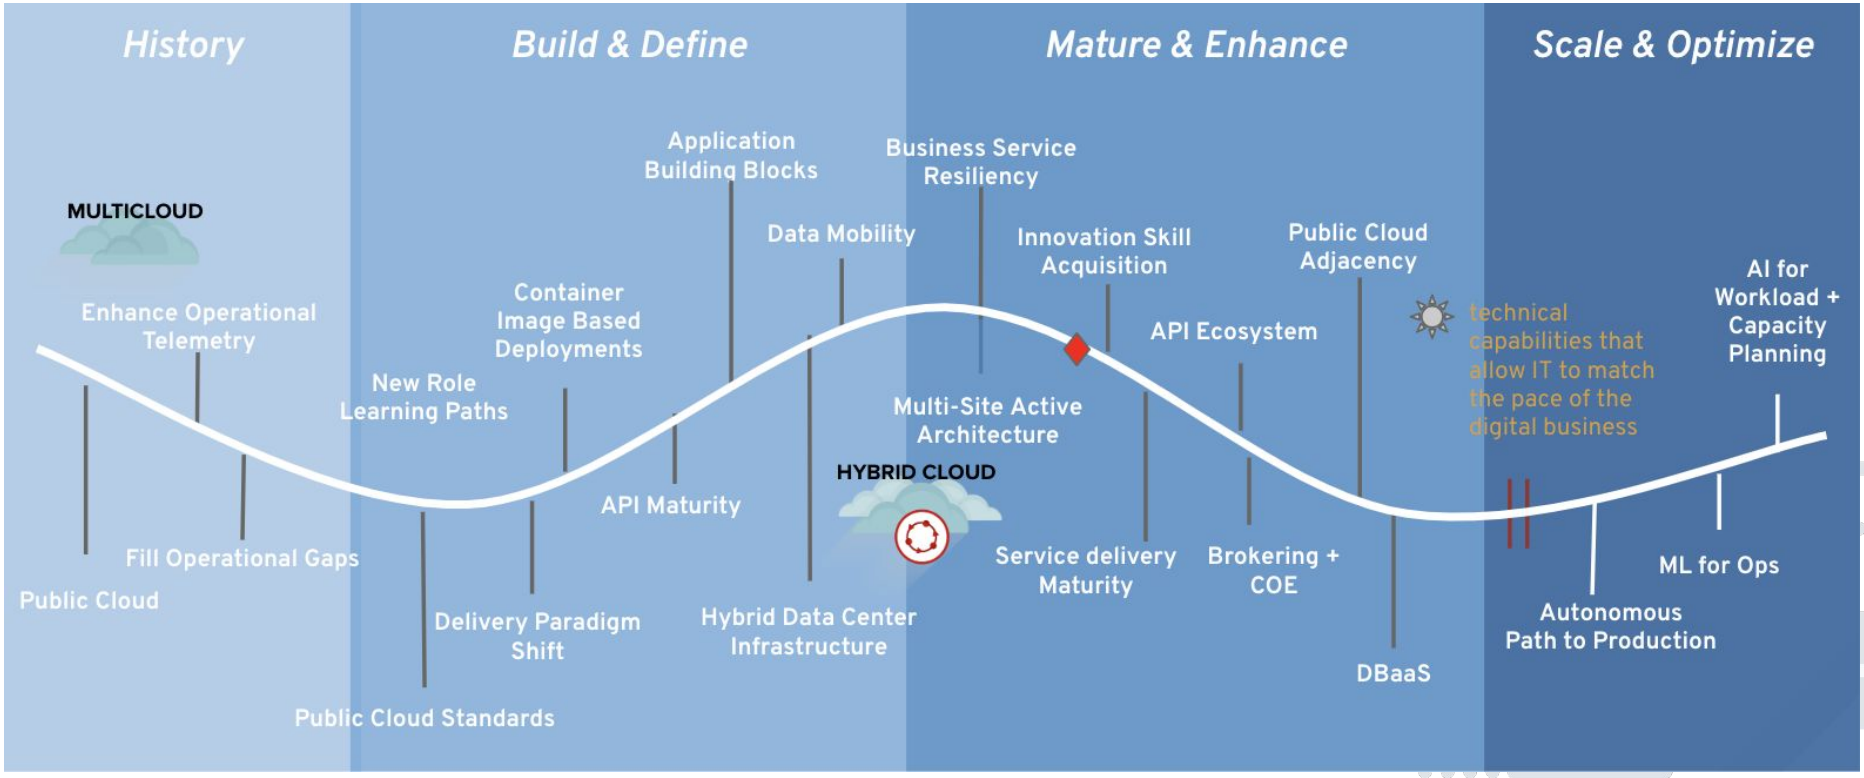

RED HAT ON RED HAT - SOLUTION + PRODUCT USE

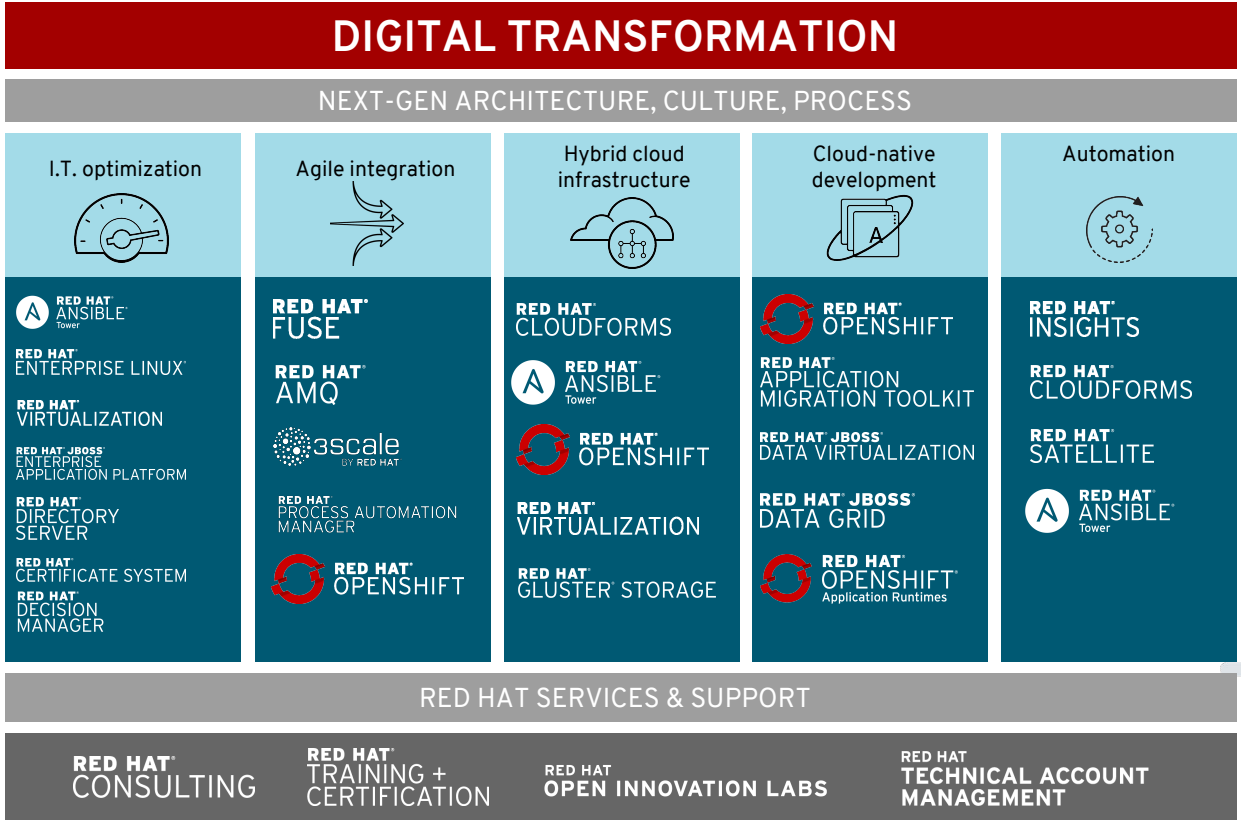

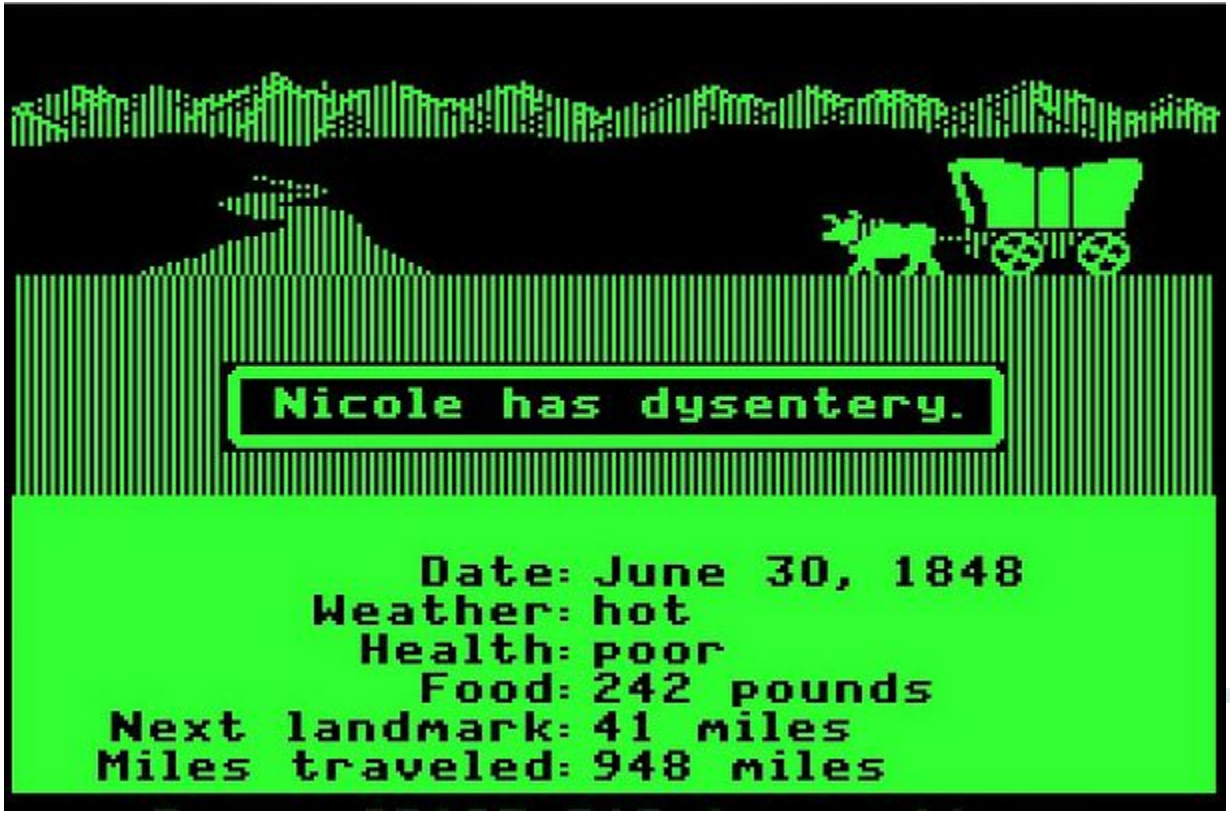

Photo credit: [The_Pug_Father]

Fast growing company

- Focus on business enablement and customer support
- Rapidly changing environment and organic growth
  - New applications
  - New technologies
  - Tech debt
- Portfolio of applications often built on shoehorned systems
  - Modern app layers on legacy infrastructure

Legacy Environment
(this might sound familiar)

One data center to rule them all

- VMs and bare metal
- Config Management
- Appliances
  - Storage
  - Load balancers
- All eggs in one basket

Automation Difficulties

- Many individual systems cobbled together
- Integration was the exception
- Automation meant screen-scraping and chaining APIs

Excess Redundancy

From the Department of Redundancy Department

- Infrastructure uptime was key
- Double/triple built
  - Active/standby Networking
  - Multiple RHV clusters
  - Active/standby Storage heads
- Cold DR site

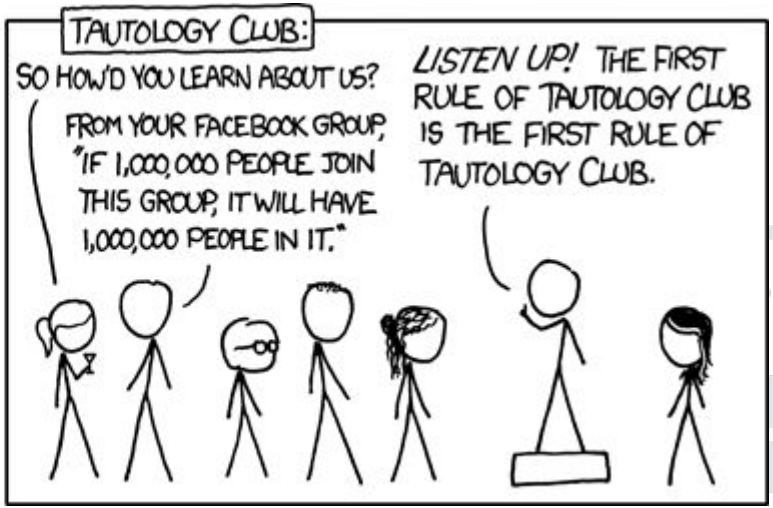

https://xkcd.com/703/

Application Architecture

- Applications assumed 100% infrastructure uptime
- Some would not tolerate device failovers
- Active/passive DBs
  - Manual failback

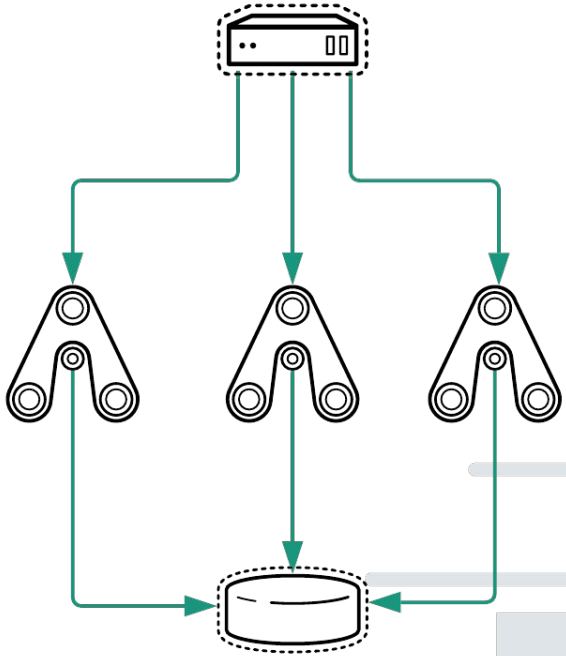

So this happened...

- HA services still resided in the same physical site
- DR failover unrealistic
- Application recovery took ages

Public Cloud Options

- Some apps moved to public cloud vendors
- Extension of data center
- Worst of both worlds
  - Named pets on someone else's hardware

We could do better,
let’s fix.

Kill the DB Monolith

- Effort to move data away from massive SQL databases
- Broke tight integration with DB
- SOA work gave way to stateless NoSQL and RESTful API data patterns

Single Sign On

- Deployed Red Hat SSO
  - Remove authentication responsibility from applications
  - Handles session management
- Multisite
  - MariaDB Galera
  - Jboss Data Grid
  - External User store
    - Mongo
    - IdM

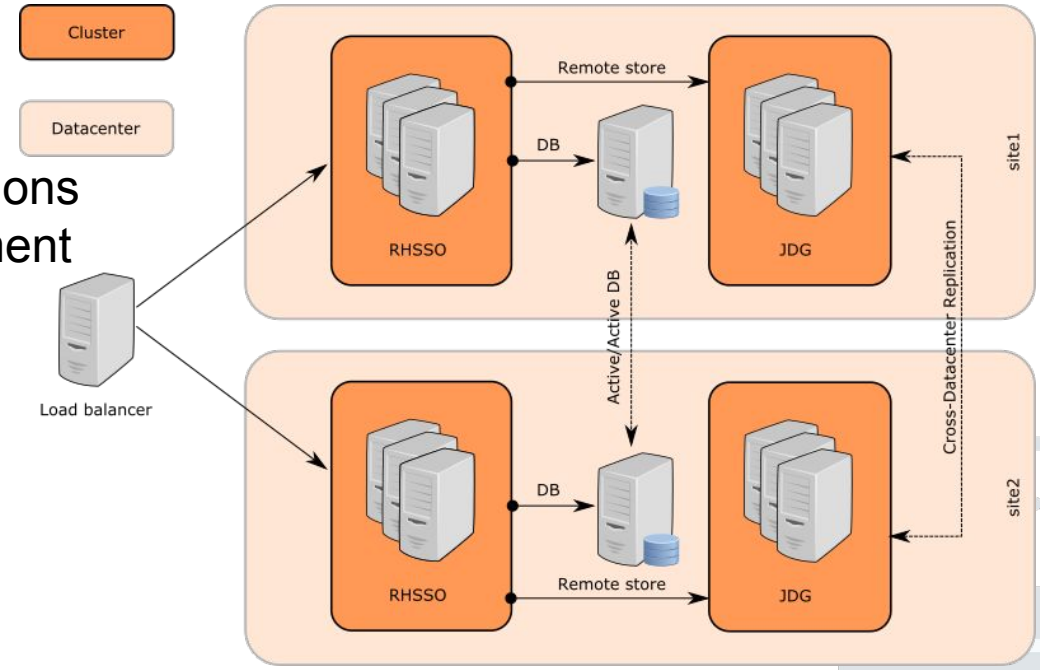

OpenShift

- OpenShift changed deployment practices in a profound way
- Container-based design principles
- Migrating app tiers
  - Decouple application from state and data tiers
- Built-in Updates

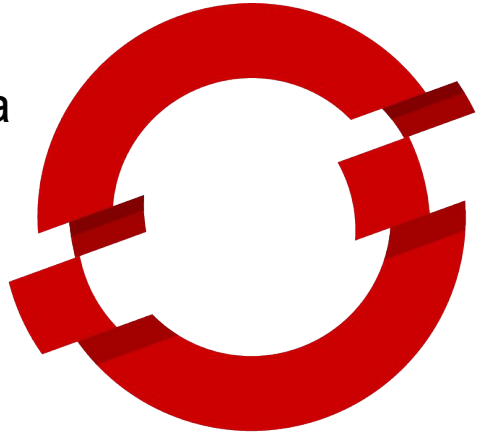

RED HAT®
OPENSHIFT®

Multisite Active

Development Pattern Prereqs

- Guard rails to ensure success
- Applications deployed on OpenShift
  - Stateless
  - Changes expected
- Session management delegated to SSO
- Loose coupling and/or microservices
- Graceful fault tolerance

Multisite Routing Strategies

- CDN handles Global Server Load Balancing
- Geographic Proximity
- Edge servers allow for advanced logic

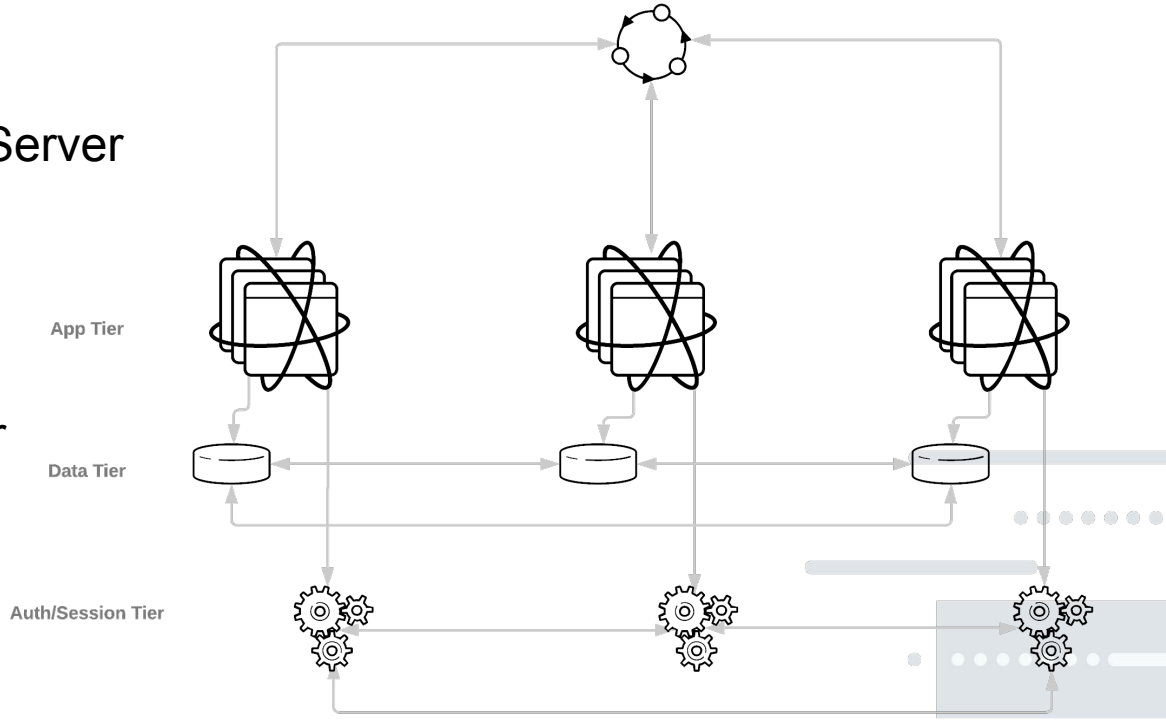

Multisite Data Strategies

- Synchronous data storage
  - MariaDB Galera
  - JBoss Data Grid
- Eventual consistency
  - MongoDB
  - Directory Server / IdM
- Storage file and block replication

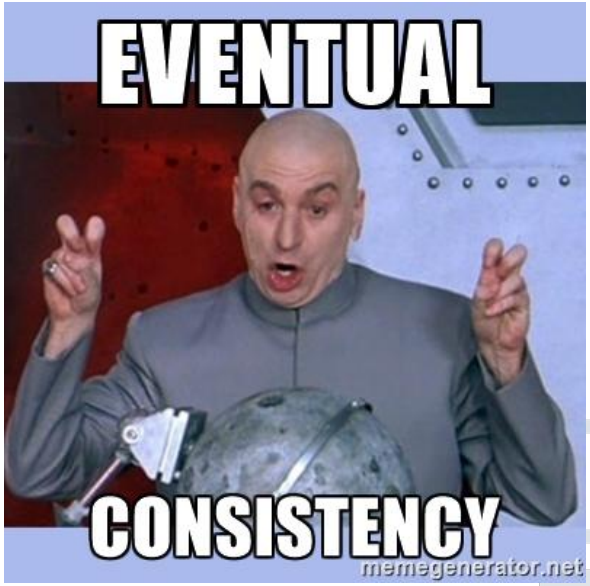

Multisite PaaS

- OpenShift deployed in three sites
- PaaS abstracted underlying infrastructure
- Front-end with a CDN

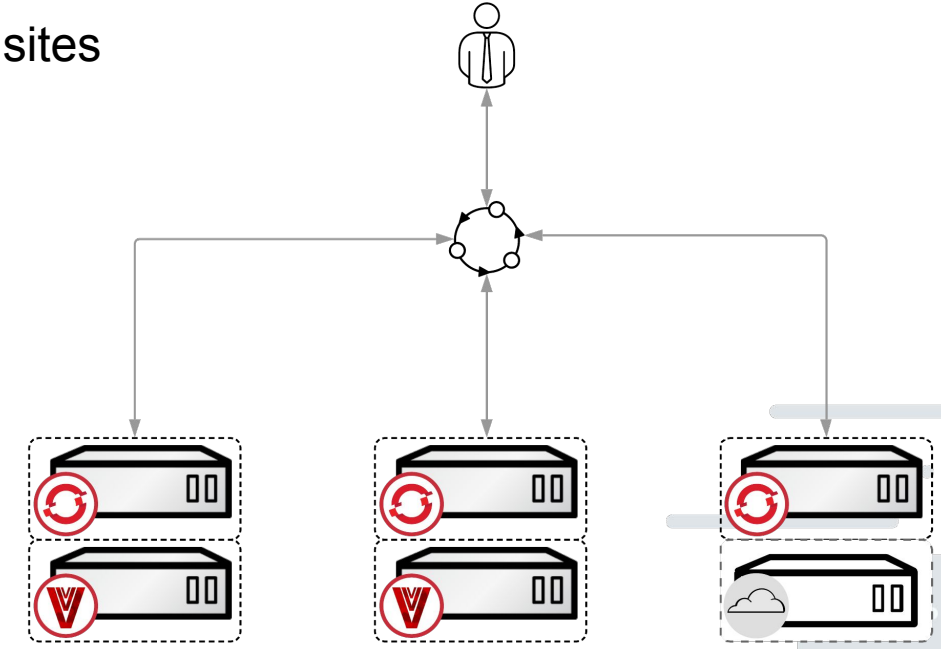

OpenStack and Ceph

Why OpenStack

- Faster Delivery of applications and services
- Vast automation potentials
- Empower development teams
- Infrastructure as code
- Lower cost through software-defined services

One Leg on OpenStack

- New data center
- Replace RHV with OpenStack
- Minimal appliances
  - Load Balancer -> Octavia + proxy layer
  - Storage -> Ceph
- RHHI-V cluster
  - Bootstrapping
  - Utility services

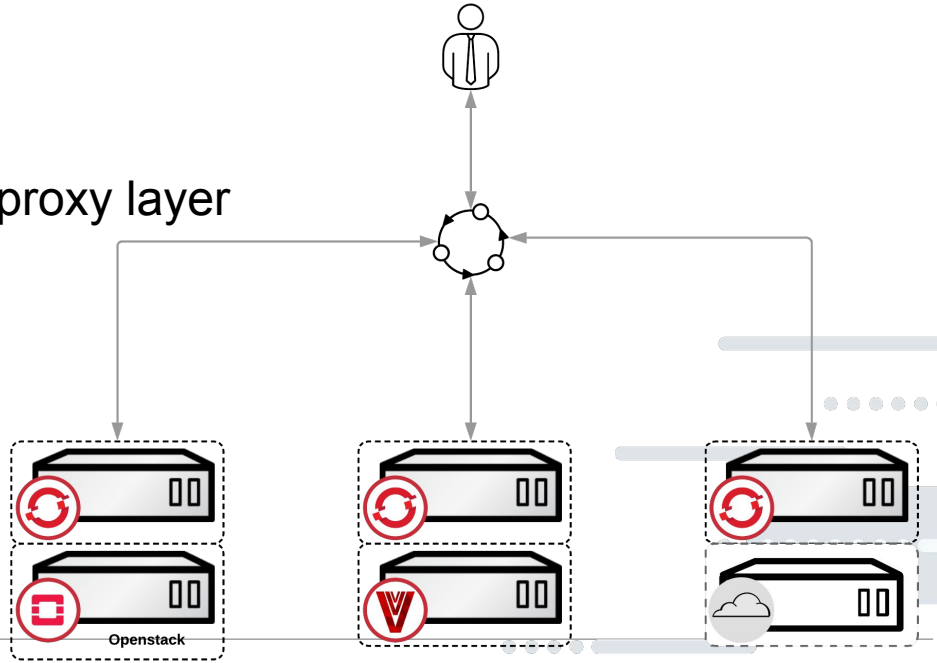

OpenStack

- Not a virtualization solution
- Cloud ecosystem
  - Including virtualization
- Software-Defined Data Center
- Integrated APIs and CLI tools

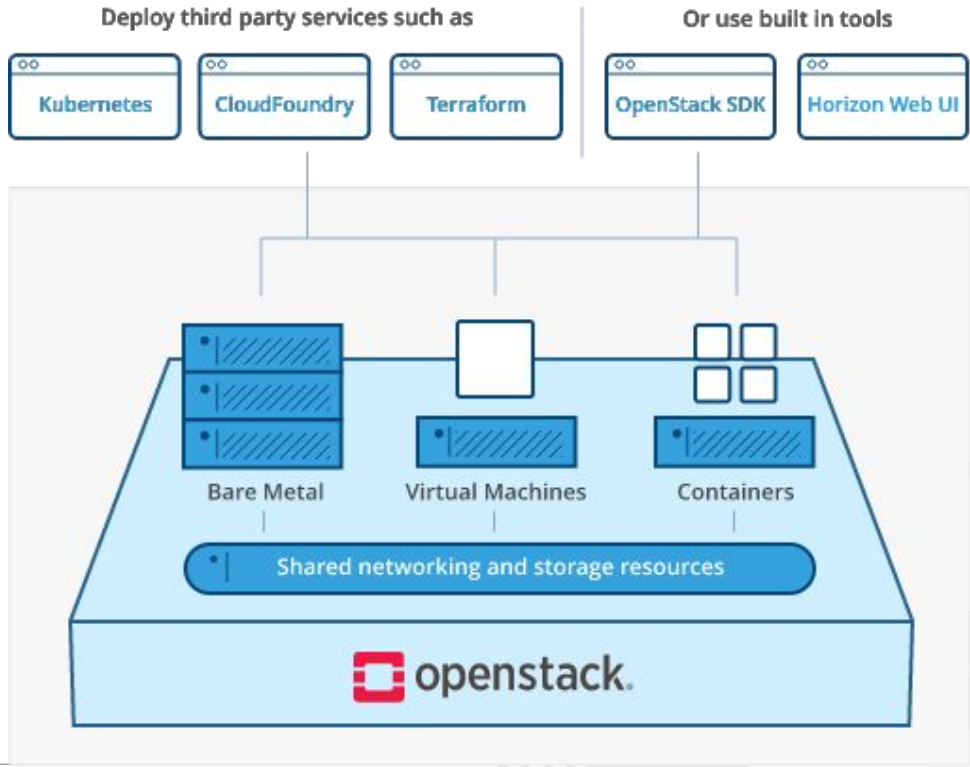

https://www.openstack.org/software/

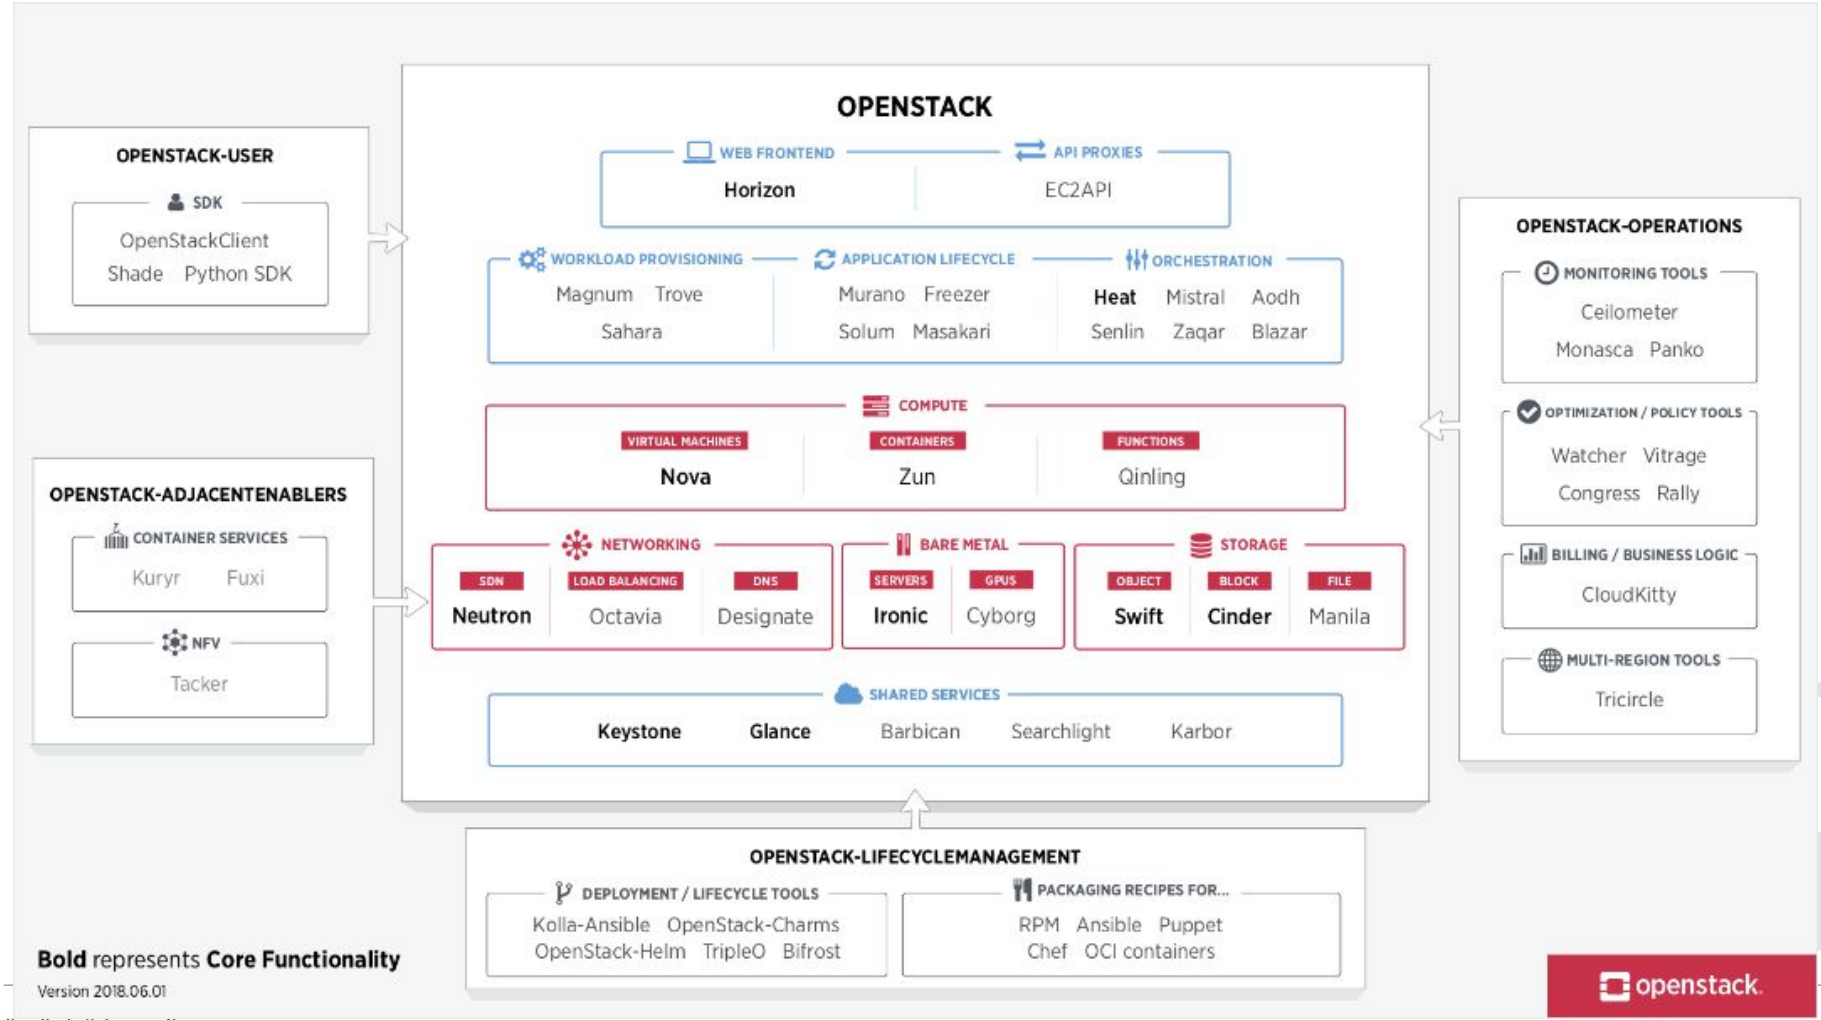

https://www.openstack.org/software/

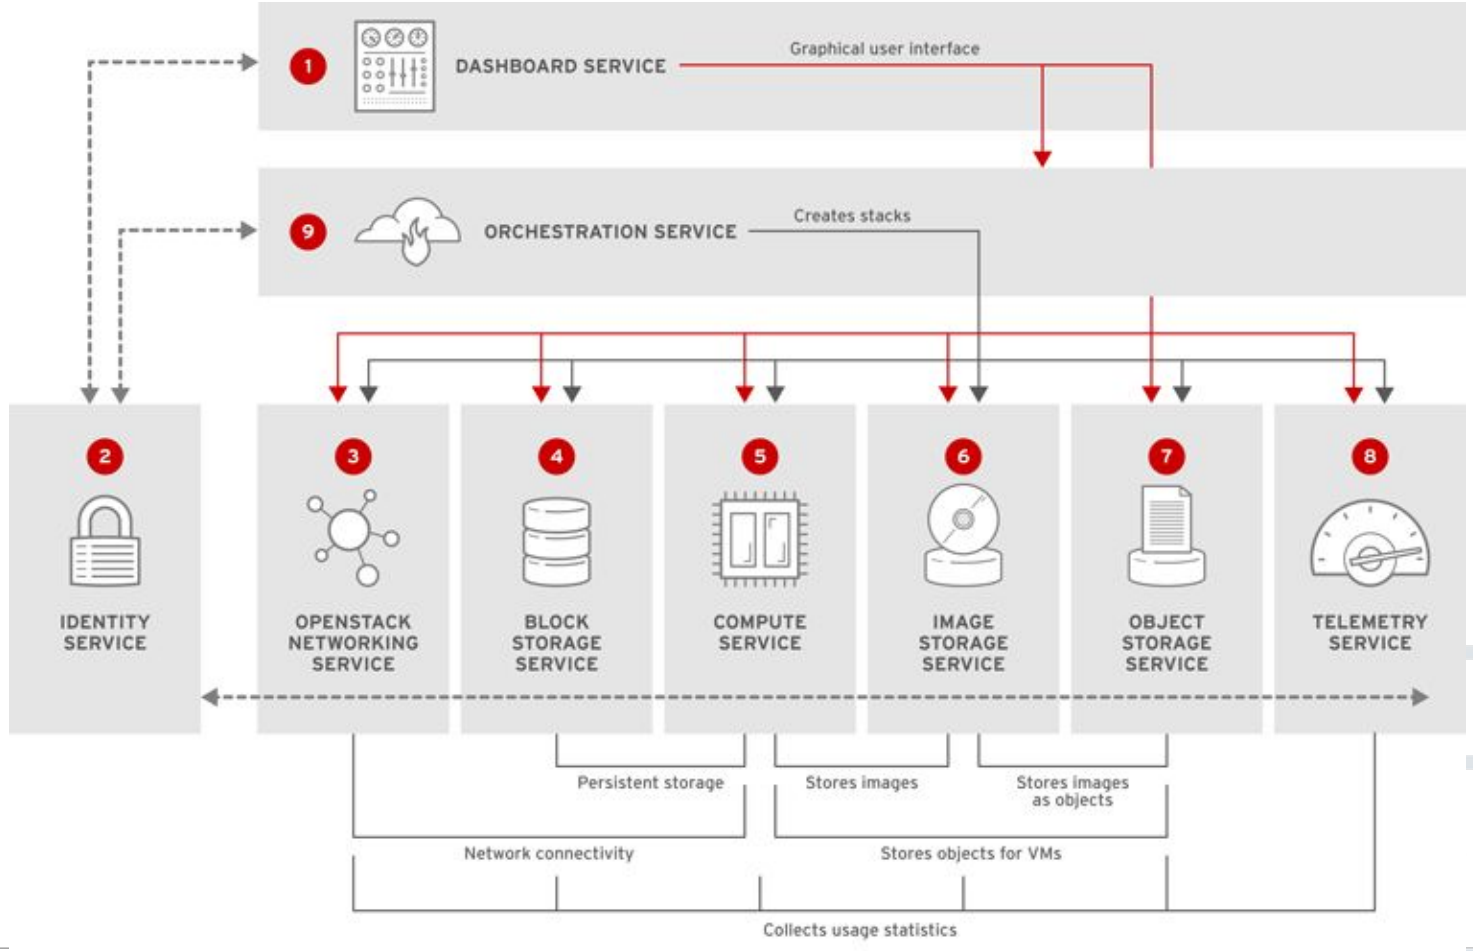

RHELOSP_347192_0615

Legacy Stack

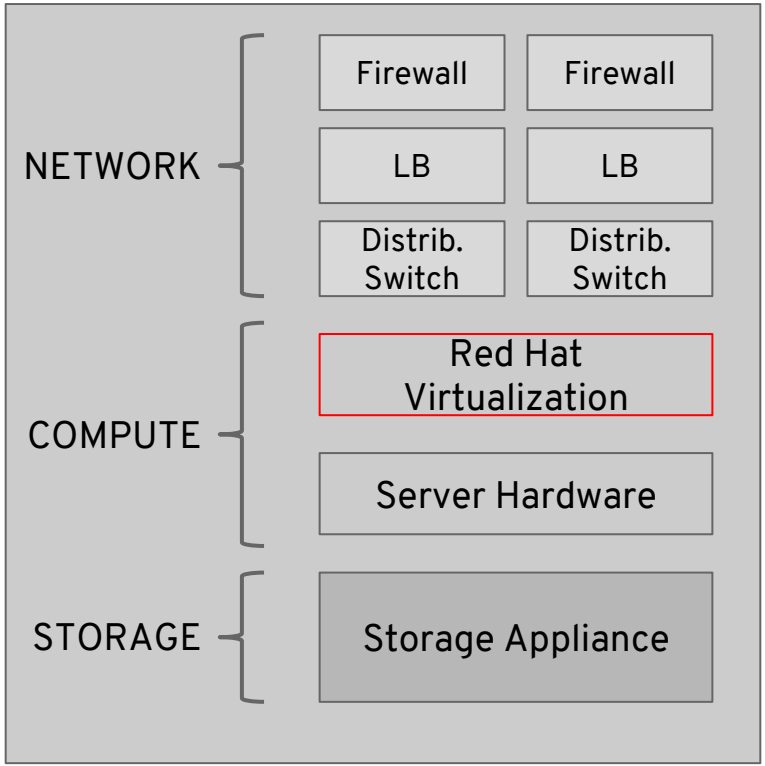

OpenStack

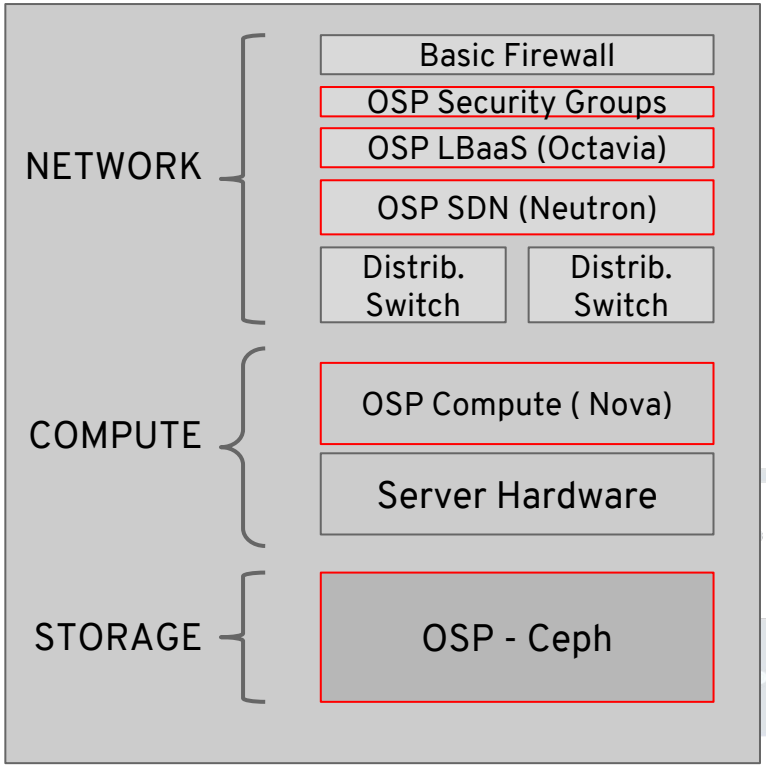

OpenStack Director

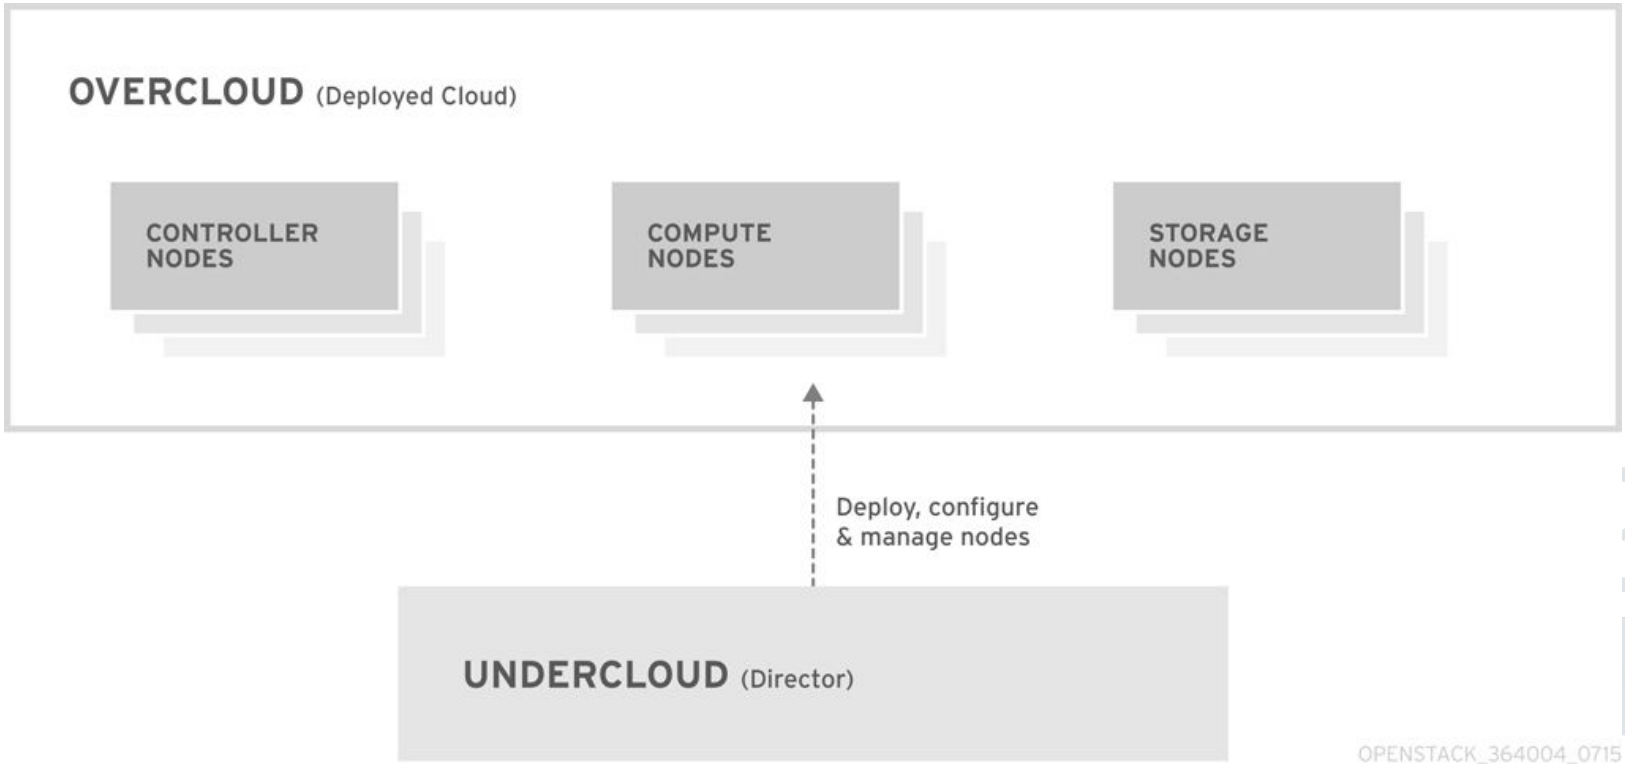

RHHI-V

- RHHI-V
  - Integrated RHV + Gluster
- Hosts
  - Director
  - Red Hat IdM, DHCP, monitoring
- HA VMs and Snapshots

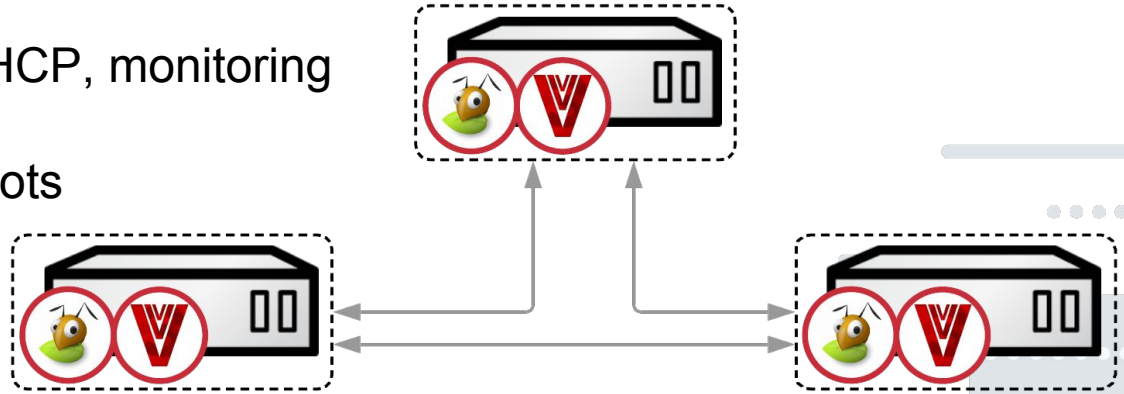

OpenStack Overcloud

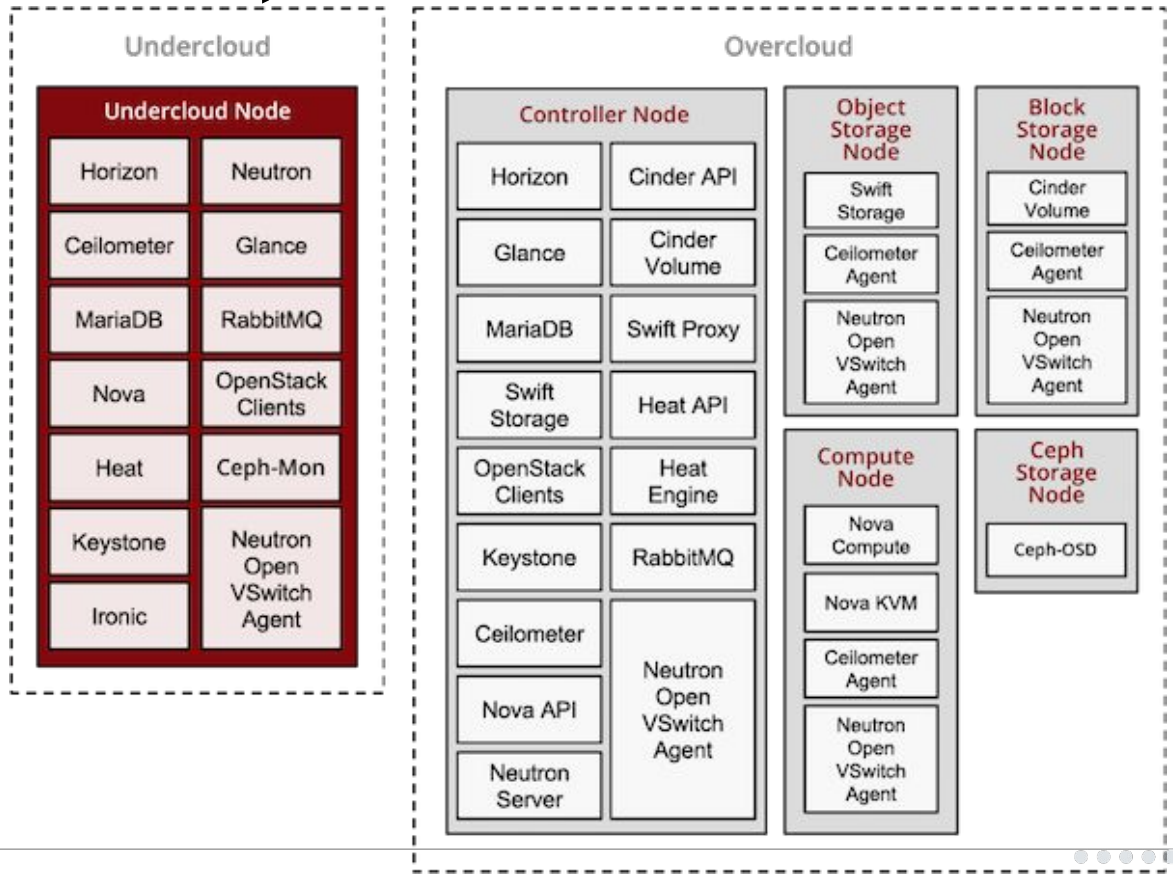

https://docs.openstack.org/tripleo-docs/latest/install/introduction/architecture.html

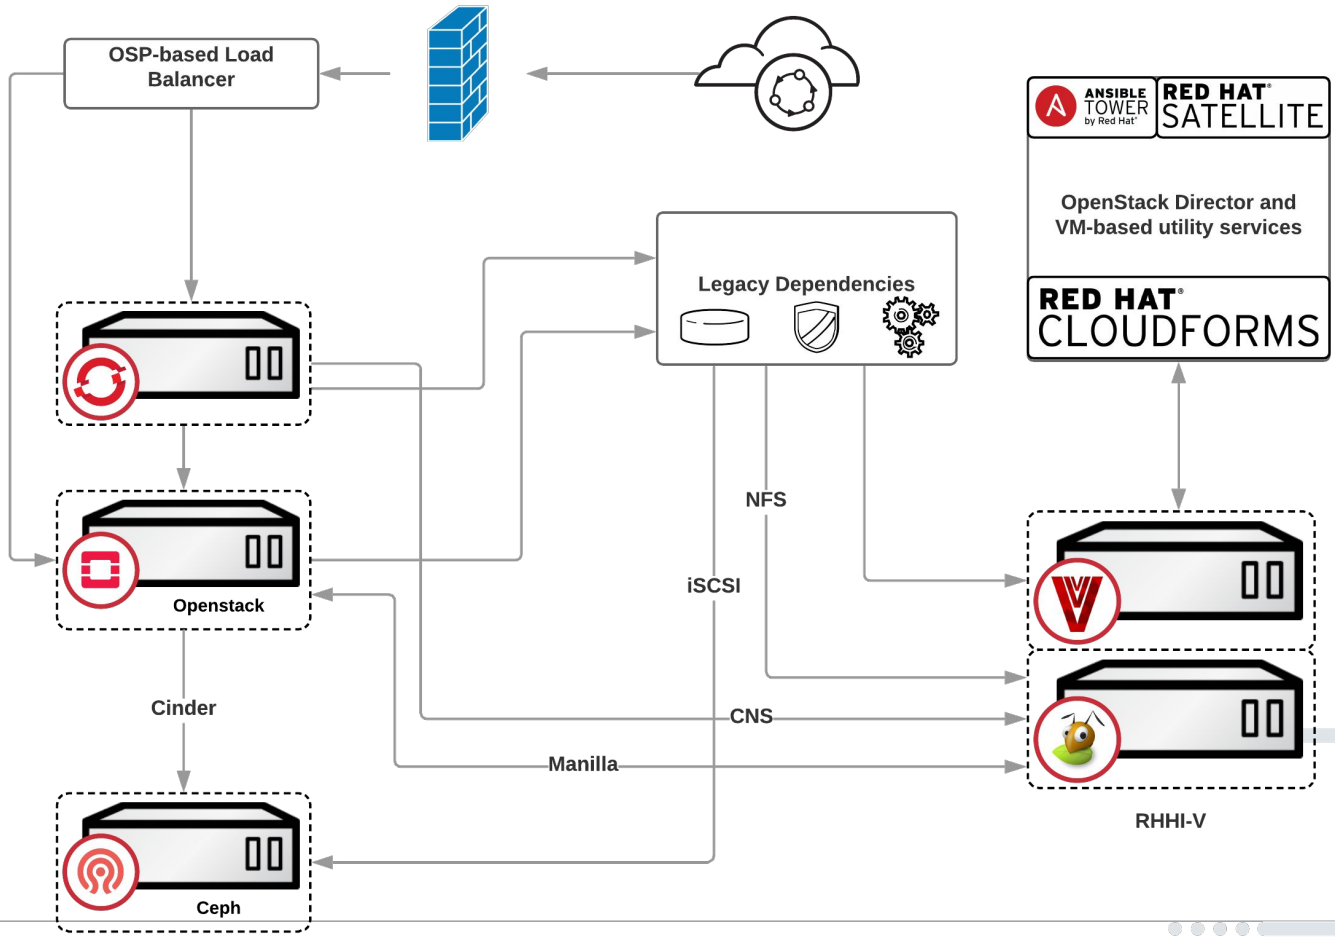
OSP-based Load Balancer
ANSIBLE TOWER by Red Hat
RED HAT SATELLITE
OpenStack Director and VM-based utility services
RED HAT CLOUDFORMS
Legacy Dependencies
Openstack
NFS
iSCSI
Cinder
CNS
Manilla
RHHI-V
Ceph

Solution Implementation

Availability Zones

- Zones
  - Failure domains
  - Resource islands
- Separate racks, server hardware, upstream network gear, power circuits, etc.
- Taking down one AZ has no impact on others
- Three zones

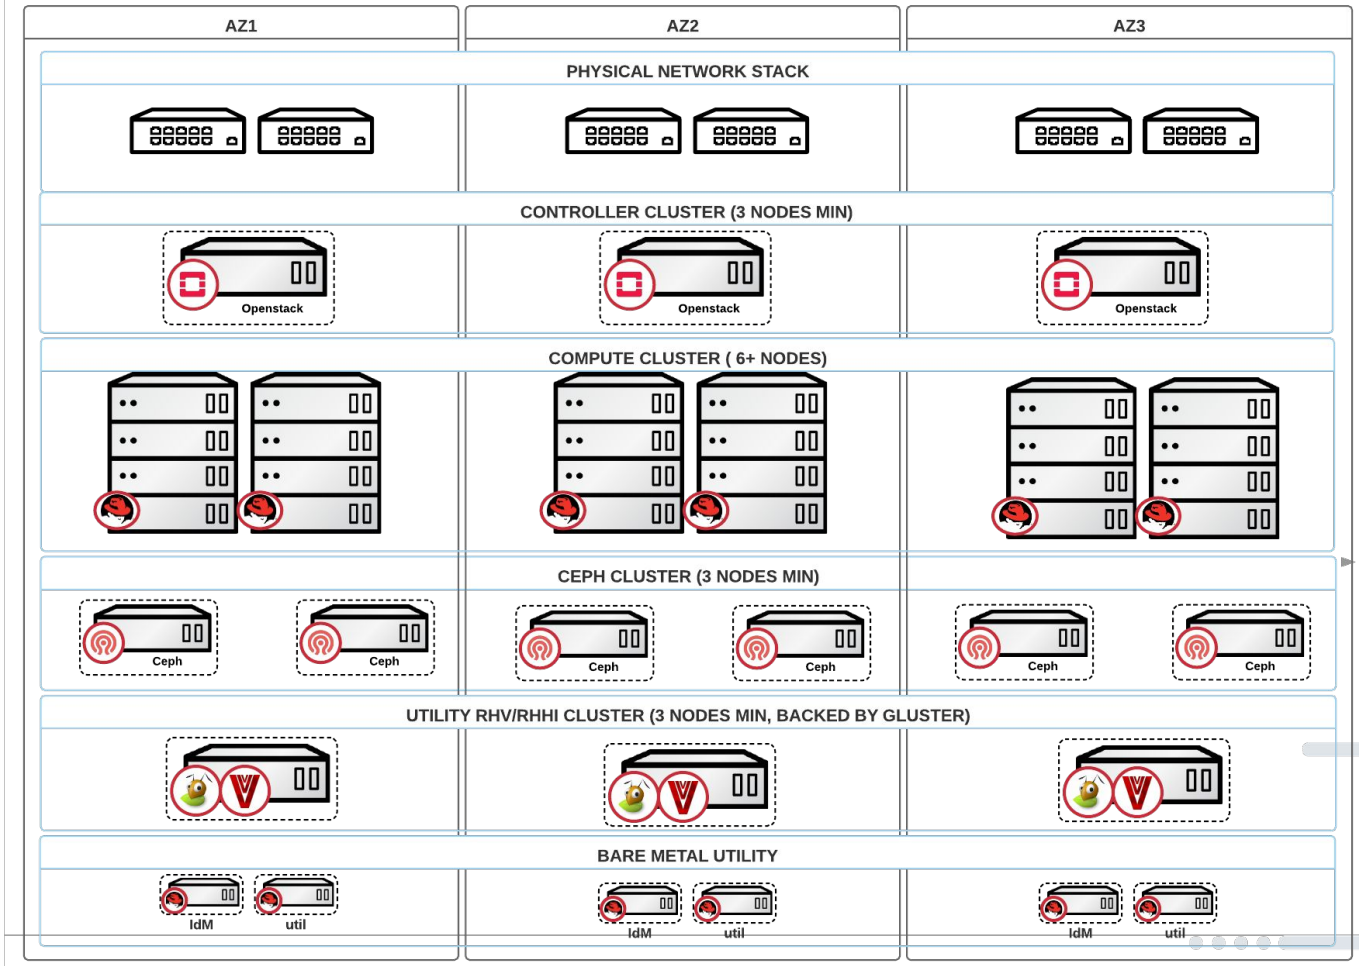
AZ1
AZ2
AZ3
PHYSICAL NETWORK STACK
CONTROLLER CLUSTER (3 NODES MIN)
Openstack
Openstack
Openstack
COMPUTE CLUSTER (6+ NODES)
CEPH CLUSTER (3 NODES MIN)
Ceph
Ceph
Ceph
Ceph
Ceph
Ceph
UTILITY RHV/RHHI CLUSTER (3 NODES MIN, BACKED BY GLUSTER)
BARE METAL UTILITY
IdM
util
IdM
util
IdM
util

Network Integration

- Provider Networks
  - Stretched VLANs spanning availability zones
  - Faster, shared resource
- Tenant Networks
  - Private overlay networks
  - Slower, more secure
- Security Groups

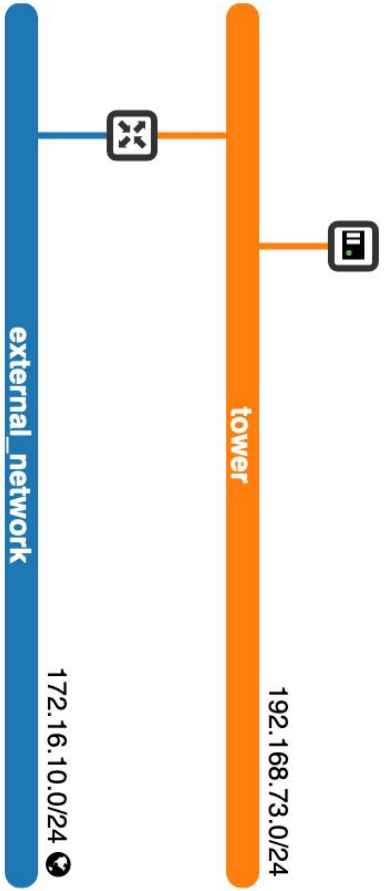

Load Balancing - Octavia

- Quick and simple OpenStack native load balancing
- Advanced logic in application tiers
- Individual Amphora provides failure domain separation

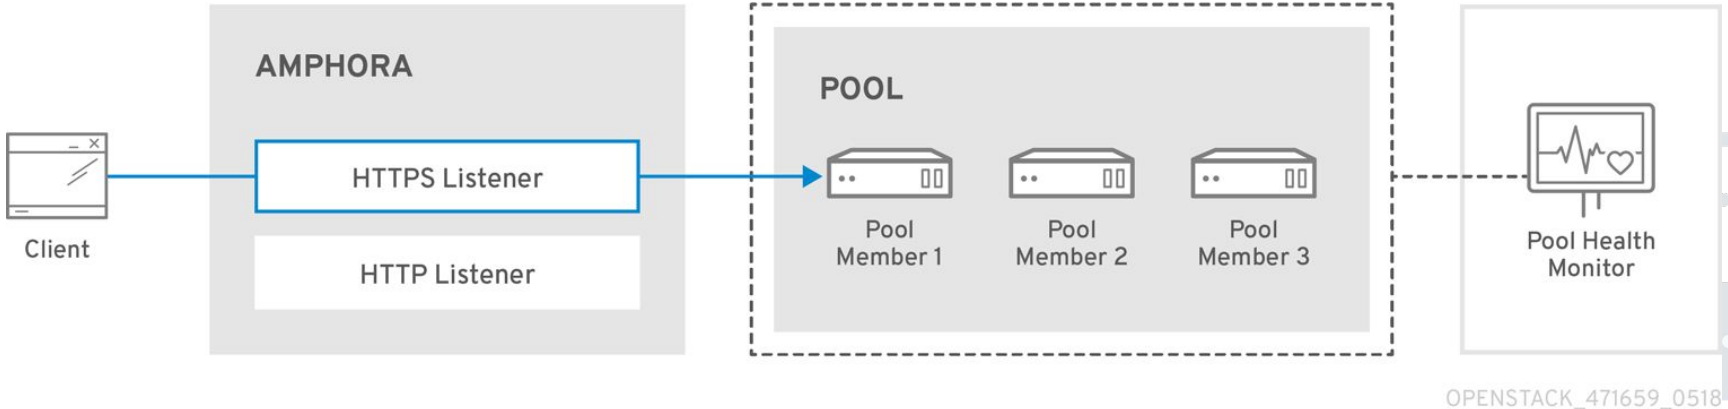

Ceph Storage

- Benefits
  - Software-defined
  - Cheaper, reliable
  - Scale-out
  - Block, Object and File storage

- Use-cases
  - VM disk storage
  - Glance image storage

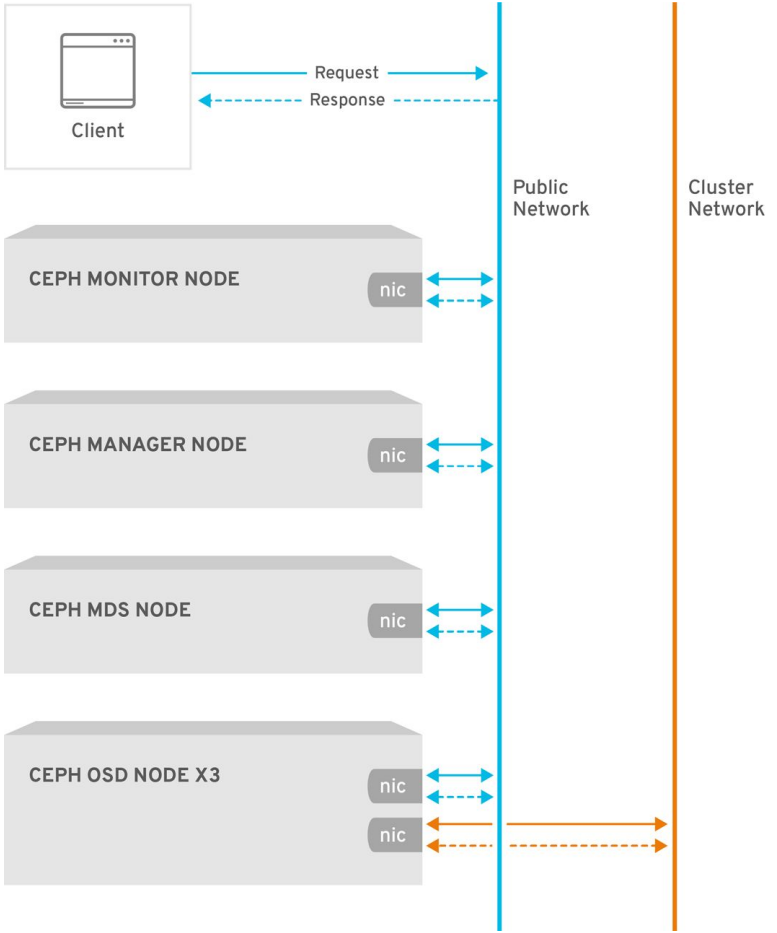

Lessons Learned

OSP Long Life Releases

- Historical cadence
  - 1 - 2 major releases per year
- Long Life releases
  - 5 years of support
    - 3 years production
    - 2 years extended life
- Short Life releases
  - 1 year of support
- We settled on OSP 13

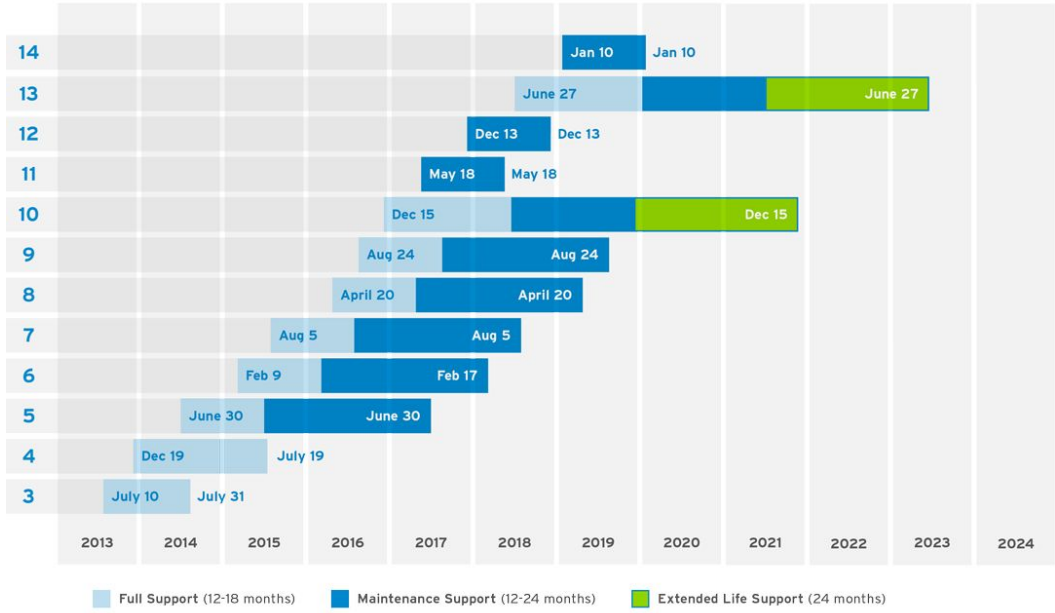

Use Cases

- OpenShift
- Non-container workloads
- Red Hat SSO
  - Multicast
- Databases
- IdM
- LucidWorks
- Legacy workloads

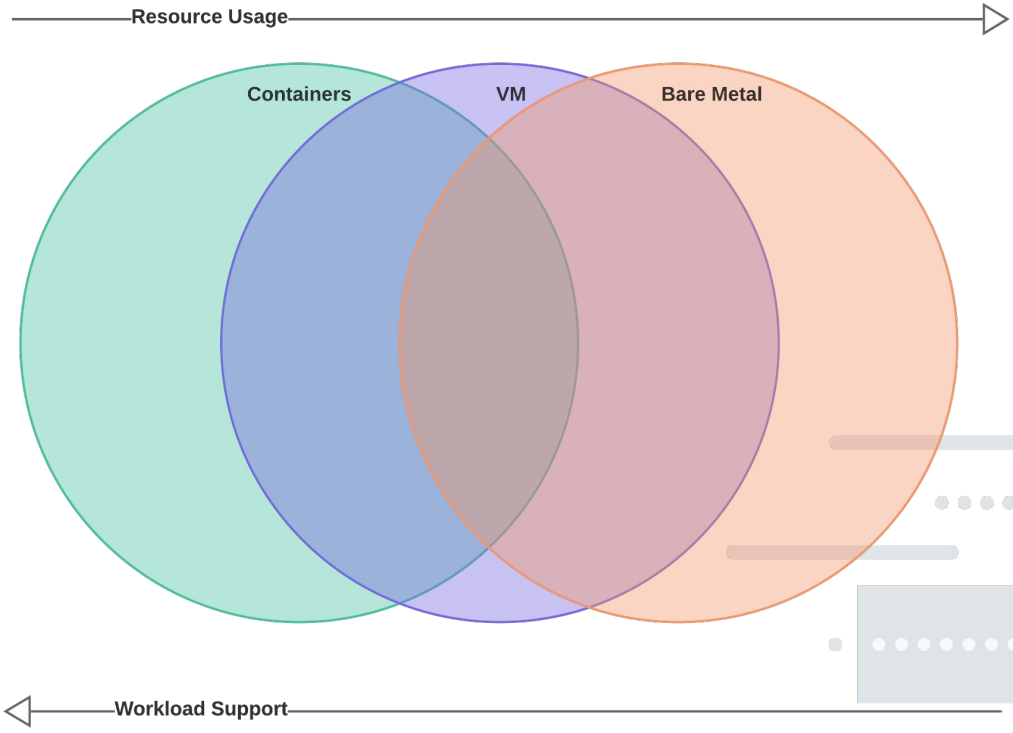

IdM and Ansible Integration

- OpenStack Novajoin
  - Native IdM integration
  - Secures all endpoints automatically (~100 certs)

- Ansible for project provisioning
  - Creates IdM host
  - Provision VM, load balancer, storage
  - Scale-up events
  - Tower API endpoint for playbook

Load Balancing

- TLS terminate and re-encrypt not supported in Octavia
  - Drove some design, ie CDN & App tier
- Stretch VLANs between the AZs for Octavia
  - One virtual IP(VIP) on all three AZs
- Embrace simple

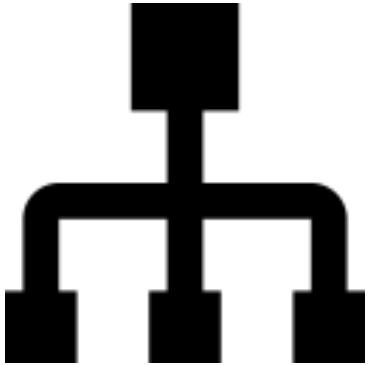

Autoscaling

- Heat-based templates, harder to implement auto scaling than hoped
- Tooling is there for your own solution
- Template-based provisioning accelerates scaling

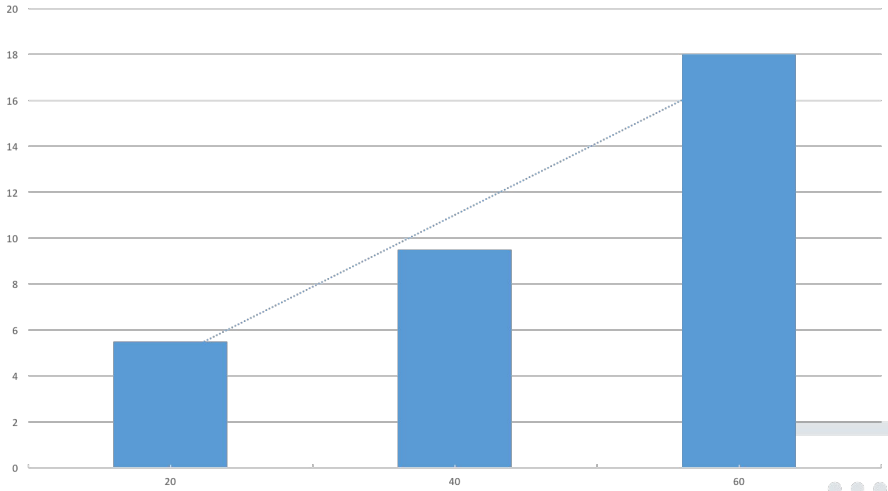

Questions?

RED HAT ON RED HAT

RED HAT I.T.
Powered by Red Hat Products

Ask us how it's done

Visit us at the "Ask Me Anything" booth and ask us about how we implement and deploy Red Hat products!

redhat.com/redhat-on-redhat

RED HAT ON RED

Thurs, May 9 3:15 - 4:00pm	Developing and running cloud-native apps on OpenShift in Red Hat's IT organization
Thurs, May 9 3:15 - 4:00pm	Developing and deploying applications in a multisite hybrid cloud

THANK YOU


plus.google.com/+RedHat


facebook.com/redhatinc


linkedin.com/company/red-hat


twitter.com/redhat


youtube.com/user/RedHatVideos