



# MAKING INFRASTRUCTURE BORING

Tony Scully  
Senior Solution Architect  
3rd October 2017

HYBRID SERVICES

RED HAT  
OPEN INNOVATION LABS

RED HAT  
CONSULTING

RED HAT  
TRAINING

RED HAT  
SERVICES

APPLICATIONS AND BUSINESS PROCESSES

MIDDLEWARE AND APPLICATION SERVICES

RED HAT JBOSS  
BPM SUITE

RED HAT JBOSS  
FUSE

RED HAT JBOSS  
DATA GRID

RED HAT JBOSS  
DATA VIRTUALIZATION



RED HAT JBOSS  
BRMS

RED HAT JBOSS  
AMQ

RED HAT  
MOBILE APPLICATION  
PLATFORM

RED HAT JBOSS  
ENTERPRISE  
APPLICATION PLATFORM

DEV

OPS

SOFTWARE-DEFINED DATACENTER

CLOUD-NATIVE INFRASTRUCTURE

RED HAT JBOSS  
DEVELOPER STUDIO

RED HAT  
VIRTUALIZATION

RED HAT  
OPENSTACK  
PLATFORM



RED HAT  
OPENSIFT  
Container Platform

RED HAT  
REGISTRY

RED HAT  
INSIGHTS



RED HAT  
SATELLITE

RED HAT  
CLOUDFORMS

RED HAT  
CONTAINER  
DEVELOPMENT KIT

RED HAT  
ENTERPRISE LINUX  
ATOMIC HOST

RED HAT  
ENTERPRISE LINUX

RED HAT  
STORAGE

PHYSICAL AND CLOUD INFRASTRUCTURE

RED HAT CERTIFIED CLOUD AND SERVICE PROVIDER PROGRAM

RED HAT  
APPLICATION  
LIFECYCLE TOOLS

RED HAT  
OPEN INNOVATION LABS

RED HAT  
CONSULTING

RED HAT  
TRAINING

RED HAT  
SERVICES

APPLICATIONS AND BUSINESS PROCESSES

MIDDLEWARE AND APPLICATION SERVICES

RED HAT JBOSS  
BPM SUITE

RED HAT JBOSS  
FUZE

RED HAT JBOSS  
DATA GRID

RED HAT JBOSS  
DATA VIRTUALIZATION



RED HAT JBOSS  
BRMS

RED HAT JBOSS  
AMQ

RED HAT  
MOBILE APPLICATION  
PLATFORM

RED HAT JBOSS  
ENTERPRISE  
APPLICATION PLATFORM

SOFTWARE-DEFINED DATACENTER

CLOUD-NATIVE INFRASTRUCTURE

RED HAT  
VIRTUALIZATION

RED HAT  
OPENSTACK  
PLATFORM



RED HAT  
ENTERPRISE LINUX  
ATOMIC HOST

RED HAT  
ENTERPRISE LINUX

RED HAT  
STORAGE

PHYSICAL AND CLOUD INFRASTRUCTURE

RED HAT CERTIFIED CLOUD AND SERVICE PROVIDER PROGRAM


RED HAT  
REGISTRY

RED HAT  
INSIGHTS



RED HAT  
SATELLITE

RED HAT  
CLOUDFORMS



DEV



OPS

RED HAT JBOSS  
DEVELOPER STUDIO

RED HAT  
CONTAINER  
DEVELOPMENT KIT

RED HAT  
APPLICATION  
LIFECYCLE TOOLS

# I.T. OPERATIONS IS BEING CHALLENGED




**INTERNAL  
BUSINESS**


**I.T. OPERATION  
CHALLENGES**

**DEVELOPERS**


# MODERN APPS & IAAS ENABLE THE DIGITAL BUSINESS



Streamlined  
and automated




Orchestration



Elastic  
and scalable



Massive, fault-tolerant  
infrastructure




Agile  
and responsive



Rapid deployment



Utility-like



Managed, policy driven  
& always on

# THE VIRTUE OF BORING

# The virtues of boring

Spontaneous  
Unpredictable  
Puzzling  
Suspenseful  
Complex  
Unique

Constant  
Dependable  
Clear  
Routine  
Simple  
Repetitious

# Principles

Everything can be rebuilt

Everything is disposable

Everything is consistent

Avoid Fragility, Embrace Simplicity

There is no 'end state'



# Methods

Automate all the things

Version Control all the things

Test all the things

CI all the things

Practice all the things

# When things go wrong...

“Ways in which things go right are special cases of the ways in which things go wrong”

John Allspaw

# When things go wrong...

Know things have gone wrong:

- White box monitoring
- Black box monitoring

Know your system:

- Have a method
- Be open and transparent

Count things that matter:

- Service MTTR and MTBF

# Things do go wrong...

## Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region

We'd like to give you some additional information about the service disruption that occurred in the Northern Virginia (US-EAST-1) Region on the morning of February 28th, 2017. The Amazon Simple Storage Service (S3) team was debugging an issue causing the S3 billing system to progress more slowly than expected. At 9:07AM PST, an authorized S3 team member using an established playbook executed a command which was intended to remove a small number of servers for one of the S3 subsystems that is used by the S3 billing process. Unfortunately one of the inputs to the command was entered incorrectly and a larger set of servers was removed than intended. The servers that were inadvertently removed supported two other S3 subsystems. One of these subsystems, the index subsystem, manages the metadata and location information of all S3 objects in the region. This subsystem is necessary to serve all GET, LIST, PUT, and DELETE requests. The second subsystem, the placement subsystem, manages allocation of new storage and requires the index subsystem to be functioning properly to correctly operate. The placement subsystem is used during PUT requests to allocate storage for new objects. Removing a significant portion of the capacity caused each of these systems to require a full restart. While these subsystems were being restarted, S3 was unable to service requests. Other AWS services in the US-EAST-1 Region that rely on S3 for storage, including the S3 console, Amazon Elastic Compute Cloud (EC2) new instance launches, Amazon ElastiCache, Amazon ElastiBlock Store (EBS) volumes (when data was needed from a S3 snapshot), and AWS Lambda were also impacted while the S3 APIs were unavailable.

S3 subsystems are designed to support the removal or failure of significant capacity with little or no customer impact. We built our systems with the assumption that things will occasionally fail, and we rely on the ability to remove and replace capacity as one of our core operational processes. While this is an operation that we have relied on to maintain our systems above the launch of S3, we have not completely restarted the index subsystem or the placement subsystem in our larger regions for many years. S3 has experienced massive growth over the last several years and the process of restarting these services and running the necessary safety checks to validate the integrity of the metadata took longer than expected. The index subsystem was the first of the two affected subsystems that needed to be restarted. By 1:28PM PST, the index subsystem had activated enough capacity to begin servicing GET, LIST, and DELETE requests. By 1:48PM PST, the index subsystem was fully recovered and GET, LIST, and DELETE APIs were functioning normally. The S3 PUT API also required the placement subsystem. The placement subsystem began recovery after the index subsystem was functional and finished recovery at 1:54PM PST. At this point, S3 was operating normally. Other AWS services that were impacted by the event began recovering. Some of these services had accumulated a backlog of work during the S3 disruption and required additional time to fully recover.

We are making several changes as a result of the operational event. While removal of capacity is a key operational practice, in the future, the tool used allowed for much capacity to be removed too quickly. We have modified the tool to remove capacity more slowly and added safeguards to prevent capacity from being removed when it will take any subsystem below its minimum required capacity level. This will prevent an incident equal to or bigger than the incident in the future. We are also auditing our other operational tools to ensure we have similar safety checks. We will also make changes to improve the recovery time of key S3 subsystems. We employ multiple techniques to allow our services to recover from any future capacity. One of the most important involves breaking services into small partitions which we call cells. By factoring services into cells, engineering teams can assess and thoroughly test recovery processes of even the largest service or subsystem. As S3 has grown, the team has done considerable work to refactor parts of the service into smaller cells to reduce blast

# Things do go wrong...



On January 31st 2017, we experienced a major service outage for one of our products, the online service [GitLab.com](#). The outage was caused by an accidental removal of data from our primary database server.

This incident caused the [GitLab.com](#) service to be unavailable for many hours. We also lost some production data that we were eventually unable to recover. Specifically, we lost modifications to database data such as projects, comments, user accounts, issues and snippets, that took place between 17:20 and 00:00 UTC on January 31. Our

# Encapsulate all this in a 'platform'

Look at high performing IT organisations

Look at the way public cloud is architected

Look for commonality in your infrastructure tools

Aim to provide IaaS, IaaS+ and IaC

# What does Red hat bring?

Automation with Ansible and OpenStack Heat

IaaS and IaaS+ with Red Hat OpenStack Platform

Scale out distributed storage with Ceph

Hybrid Cloud management with CloudForms

(Plus lots of other products and services)


# FROM COMMUNITIES TO ENTERPRISE




© 2014



Release:  Project Type:  Module:  Company:  Contributor:  Metric:



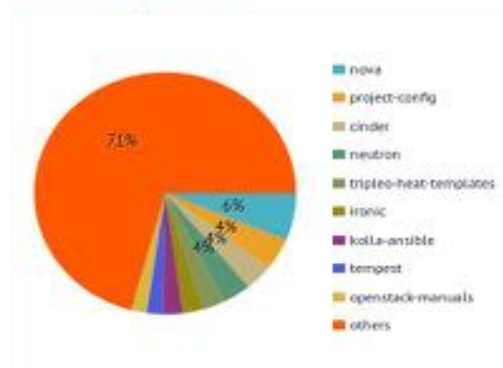
### Contribution by companies



Show  entries Search

#	Company	Reviews
1	Red Hat	33918
2	Huawei	10917
3	Mirantis	8945
4	Rackspace	8490

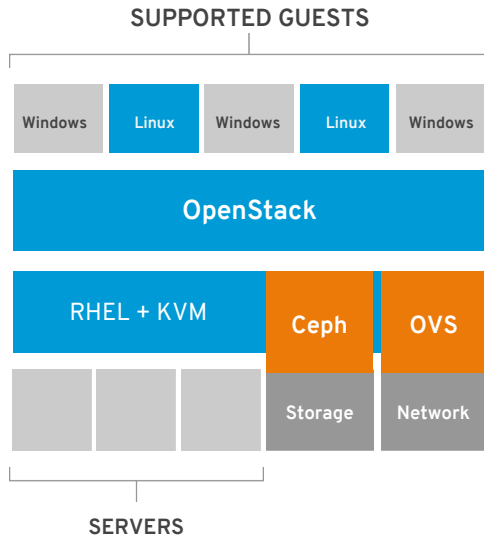
### Contribution by modules



Show  entries Search

#	Module	Reviews
1	nova	9234
2	project-config	6194
3	cinder	5147
4	neutron	5140

# Co-engineered with RHEL



Virtualization

Security

Ecosystem

Network

Storage

Security Enhanced Linux (SELinux)

KVM

Network Stack

Device Drivers

LINUX KERNEL

# LIFECYCLE CONSUMPTION OPTIONS



CUSTOMERS DESIRING LONGER LIFE VERSION	CUSTOMERS DESIRING LATEST FEATURES
Long life evrsions offered every 3rd release	Offered on each release
Offers standard 3-year lifecycle, with optional 1-2 years of ELS (extended lifecycle support)	Supported for 1 year
Will offer long life → long life tooling for migrations	Utilize director for automated upgrades and updates continuously
Customers staying on Red Hat OpenStack Platform 10	Customers moving to Red Hat OpenStack Platform 11

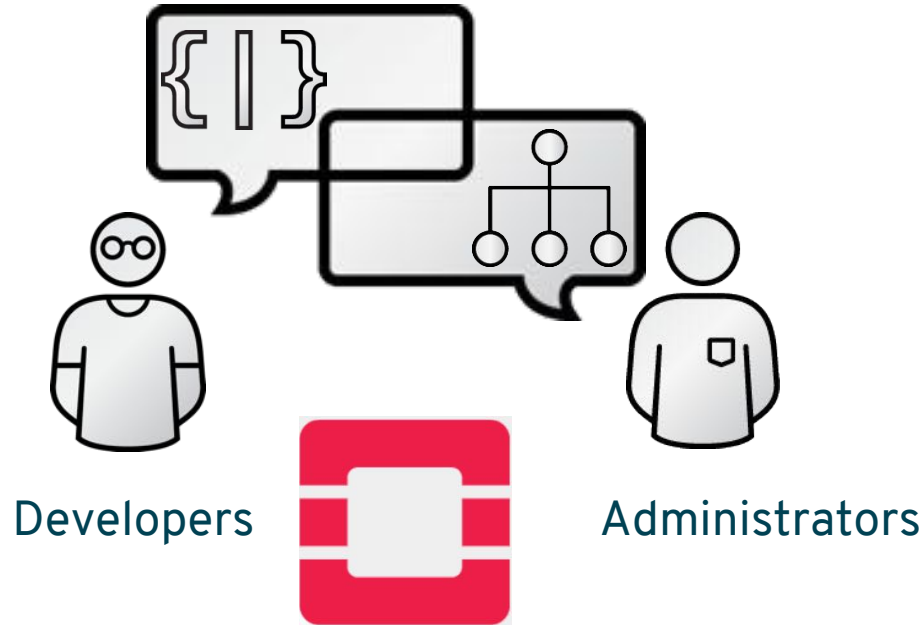
# LIFECYCLE SUPPORT

Every 6 months we release a version of Red Hat OpenStack Platform supported for 1 year

Every 18 months we produce a “long life” version, which customers can opt to have support for up to 5 years

		Long life			Long life		
RHOSP 8 Liberty	RHOSP 9 Mitaka	RHOSP 10 Newton	RHOSP 11 Ocata	RHOSP 12 Pike	RHOSP 13 Queens	RHOSP 14 R....	RHOSP 15 S....
3 years	3 years	3 years (+2 years)	1 year	1 year	3 years (+2 years)	1 year	1 year

# OpenStack connects two worlds

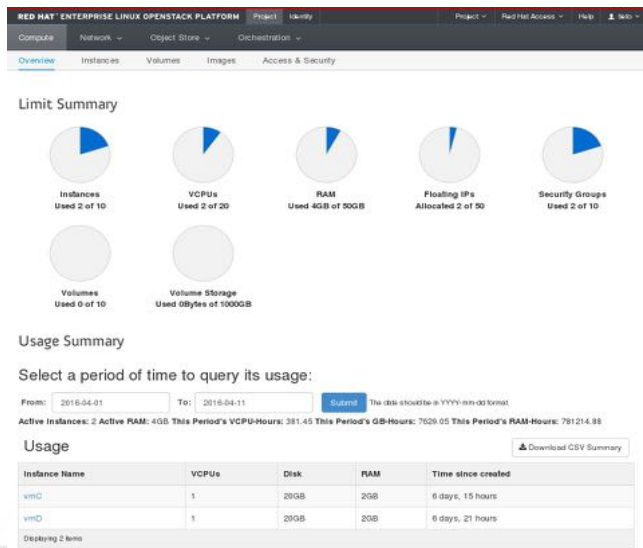


# OpenStack connects two worlds

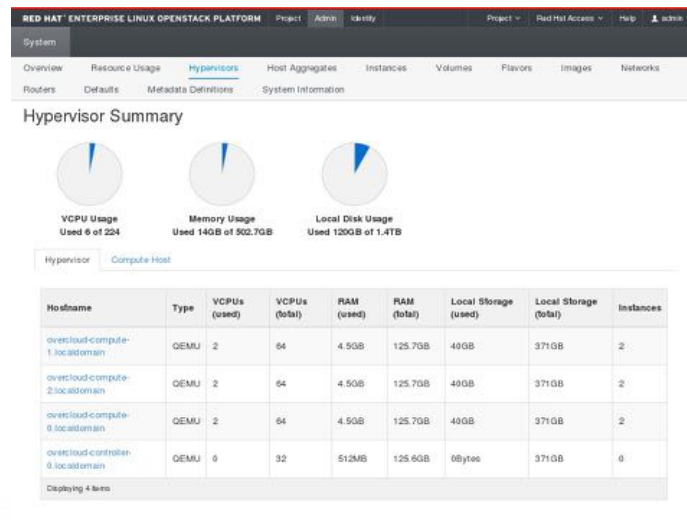
**Tenant view** – the actual OpenStack IaaS **user**  
Limited by what the Operator decides to offer in that cloud

**Operator view** – often the same role that has **root** access to the systems  
Combines configuration files and API actions to create a working environment for his tenants.

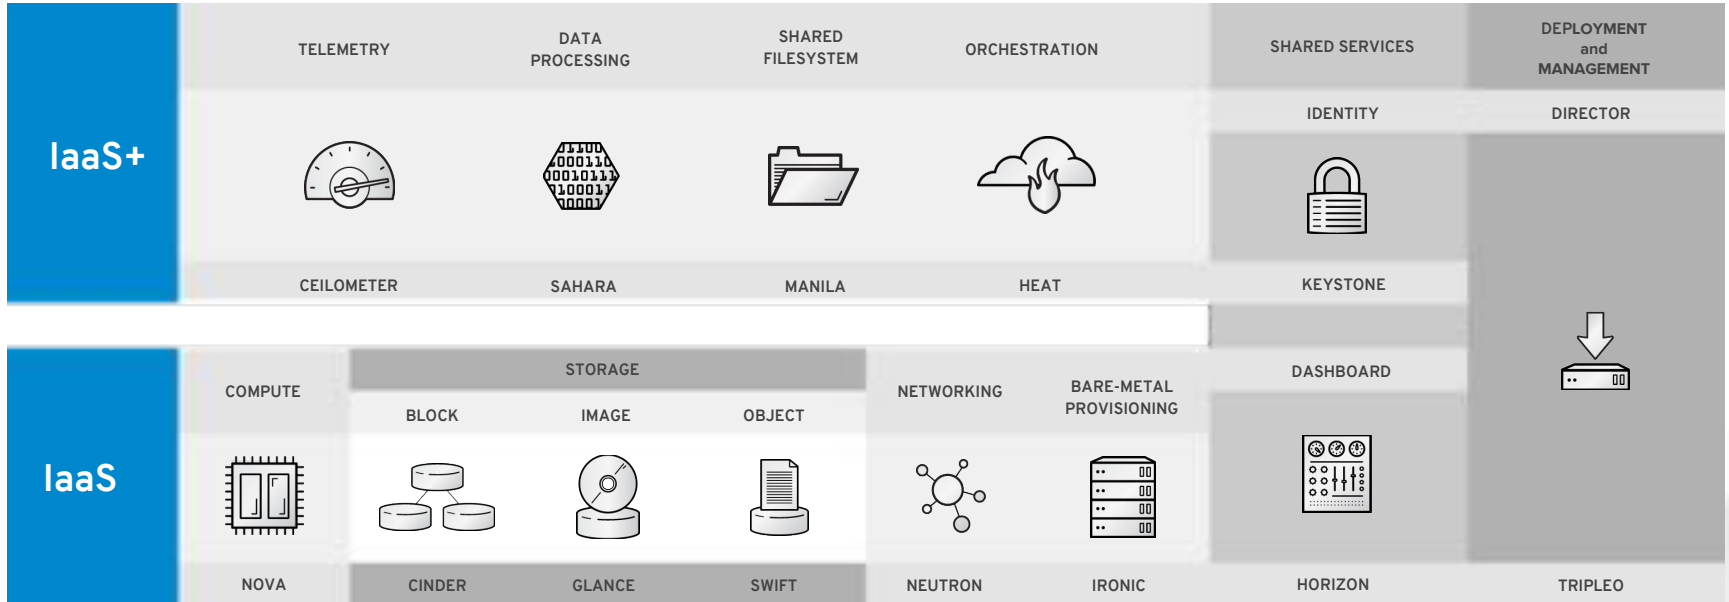
Tenant view



Operator view




# Core Components in version 11 (Ocata)



# Compute (Nova)


- I need VMs, anytime
- How many can I have?
- It must be secure
- SSH and VNC please?



- I have hardware capacity available
- This is how you consume it
- I set usage quotas
- I design for performance and scalability




# Networking (Neutron)



- I need my own network, isolated from others
- Some private IPs, some public IPs
- These are my QoS specs
- Let me share networks with others

- I design a network overlay and provide external access
- I have very few Public IPs
- I set rules, policies, quotas
- With SDN, I can centrally manage and monitor it all

# Block Storage (Cinder)




- Too much data in my VMs!
- I need permanent storage
- Can I snapshot and backup/rollback?
- Encrypted, please!



- I constantly buy storage
- I must allocate space to tenants
- I can combine different tiers of technologies (NAS, SAN, NFS)
- I set rules, policies, quotas

# Object Storage (Swift)




- My application needs object storage (files, media)
- I can use HTTP(s)
- Stateless please! No time for mounting filesystems



- I will offer a private S3-like experience
- I must scale without limits
- I want advanced features

# VM Image Storage (Glance)




- What operating systems can I use?
- This is my own version, store it just for me
- Is the OS image genuine?
- Take this VMWare template and import it



- Only approved OS can be used in my cloud
- Centrally offer updated OS
- Leverage storage integration to reduce network usage

# Shared File System (Manila)




- I need a network folder to share files between VMs
- Sometimes I'll share it with other users in my team
- I don't want to manage the folder (permissions, quotas)



- I don't have the time to create temporary shares and enable network security
- I prefer to automatically leverage OpenStack users and groups

# Identity and Access Control (Keystone)




- I am not a hacker, believe me!
- My boss just gave me permission to ask for VMs
- Where are all the services?
- I am a project lead, I must be admin of my project



- Who are you?
- Let me validate with LDAP
- I must integrate with my company's SSO
- I must secure entry points with TLS Certificates

# Dashboard (Horizon)

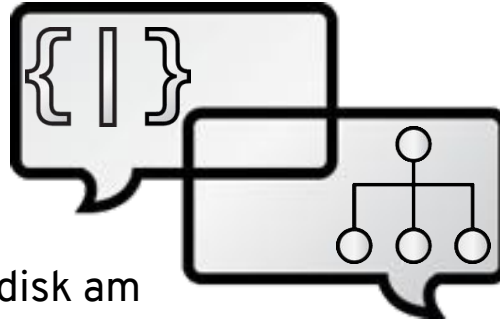


- I need a UI to manage my workloads or troubleshoot
- I don't like the CLI
- I want to see my Heat topologies
- Quickly display my quota usage and default options




- I want an admin panel
- I want a quick access to my Red Hat Access account
- I want to see all Neutron networks and routers

# Telemetry (Ceilometer)




- How much CPU, RAM, and disk am I using, i.e. per hour, per week?
- Allow me to set alarms and use my own infra to react



- I wish I could charge back / show back how much every user is consuming
- This is useful for my own internal usage!



# Orchestration engine (Heat)

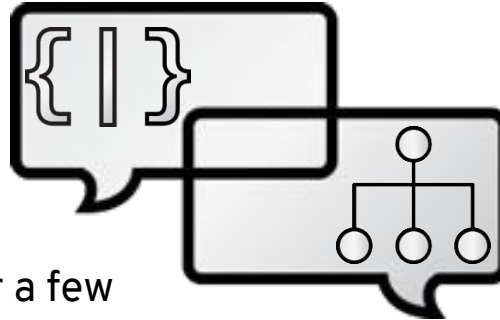


- This is the blueprint of my application deployment: dependencies, config, etc
- Can you run this for me?
- [Scale it out](#) when this threshold is reached




- I want to automate as much as I can and offer public-cloud-like efficiency
- Auto-scaling, load balancers and quotas allow me to monitor and predict demand

# Data Processing (Sahara)



- I need a Hadoop cluster for a few hours
- I need to try different Big Data platforms
- I want my clusters to scale automatically



- I don't have the manpower to customize big data platforms to all my tenants
- I will get 3<sup>rd</sup> party providers and deliver their stack as a service

- <https://betsandbits.wordpress.com/>
- <https://betsandbits.files.wordpress.com/2016/10/openstackreferencearchitecturewhitepaper.pdf>
- <https://www.openstack.org/videos/summits/boston-2017>

“

“We now have the capability to alter our infrastructure based purely on code. This ability allows our developers to make changes and use infrastructure at a pace that suits them. We can deliver products to market more quickly than our competitors. It’s a game-changer.”

”

- Paul Cutter, CTO, Betfair

- World’s largest online betting exchange
- 135M+ daily transactions, 3.7B+ daily API calls
- Simplified development to help developers focus on improving product quality and customer experience
- Automated testing and infrastructure provisioning, reducing time to deployment for new products from days or weeks to hours or minutes
- Hjjggh
- Ghjggh
-



RED HAT

**FORUM**

Europe, Middle East & Africa