

BOSTON, MA JUNE 23-26, 2015

OpenStack Nirvana: Big Data & Elastic Infrastructure Together at Last

Keith Basil

Principal Product Manager, Red Hat

Red Hat Summit Boston June 2015





Keith Basil

• personal Virginia hare scrambler, plays chess.

• professional



Cloudscaling, Time Warner Cable, FederalCloud.com, Cisco and a couple of startups

Image: blended skype/twitter/github/irc, life: noslzzp





Agenda

- So.. What Are We Talking About Here?
 - OpenStack, Sahara & Ironic
- Who?
- Implementation and Performance
 - Network

• VMs vs. Bare Metal

- Getting to the End State
- Questions
- Resources



(slide 3)

So. What Are We Talking About Here?







	-
	1.
	Real Property lies
	1000
	100
The second secon	
THE REPORT OF THE PARTY OF	-
A REAL PROPERTY AND A REAL	
a line a superior	
and the second se	
100 100 100 100 100 100 100 100 100 100	
NAME OF TAXABLE PARTY.	
STATE OF THE OWNER	-
-	
	-

Data, Data, Data.

- We are generating a large amount of data...very fast.
 - Social media footprints
 - Financial transactions
 - System instrumentation (logs, metrics, etc)
 - Interview of the sense of th
- Across many industries..
 - Financial Services, Health Care, Human Sciences
 - Telecommunications, Travel, Energy, Retail

(slide 6)



Data, Data, Data: Why Elasticity Is Needed

• Our Data is too large

• We are past the point of single computers being able to efficiently handle our data

Service Requests are too large

Client devices are more plentiful than ever

• Or, BOTH..



We need more cowbell cloud!





OpenStack, Sahara & Ironic

#redhat #rhsummit

Openstack CLOUD SOFTWARE







Hadoop's MapReduce as a Unix Command Line

Map

#redhat #rhsummit

cat /my/log | grep '\.html' | sort | uniq -c > /my/outfile

Sort & Shuffle

Reduce





OpenStack Sahara: Logical Architecture



MAPR



Sahara Detail



(slide | 3)







Ironic Logical Architecture







Ironic Architecture Detail

Magic here!









Elastic Infrastructure .. meets (OpenStack)





Elastic Infrastructure (Hadoop)



Herefords and Angus!



No pets here - all cattle!

#redhat #rhsummit





(slide 19)





The Holy Grail of Elasticity

- One set of APIs to drive both laaS and Big Data
- Both love east-west network optimization
- Both expect failure



aS and Big Data otimization



(slide 20)



Who?

#redhat #rhsummit

Openstack CLOUD SOFTWARE

••••

•



Red Hat, HP, Mirantis and Rackspace







Implementation and Performance

#redhat #rhsummit

Openstack CLOUD SOFTWARE

•••



Networking for Elastic Clouds





Solving for the Network Latency





Enterprise vs Cloud Fabric



Traditional Enterprise Topology

Network diagrams referenced from http://cto.vmware.com/is-your-cloud-ready-for-big-data/ (slide 26)



Modern Cloud Friendly Topology





Network Elasticity is Required.



Elastic Cloud Resource Map





Each unit here could be a server or rack of servers



Core Fabric Requirements

- OpenStack and Hadoop friendly networking features:
 - Availability and Resiliency (multi-path, per-flow routing)
 - Resource Node (compute/storage) Data Throughput
 - Low Network Latency
 - Congestion Management



Spine and LeafTopology

• Ask your friendly network vendor for guidance

Cisco, PlumGrid, Mellanox, Brocade, Juniper, Force I O, etc.



http://bradhedlund.com/2012/01/25/construct-a-leaf-spine-design-with-40g-or-10g-an-observation-in-scaling-the-fabric/

BRAD HEDLUND .com



40G Leaf/Spine



Virtual vs. Bare Metal Performance





Virt vs. BM: Ephemeral Disk Performance



- 60% due to I/O overhead
- 30% due to memory efficiency in virtualization

Heavy tuning is required.







Getting to the End State

#redhat #rhsummit



(slide 33)



OpenStack and Non-Ephemeral Hadoop



Traditional OpenStack laaS plus Hadoop on virtual machines deployed by Sahara

openstack

More permanent, performant Hadoop clusters also deployed by Sahara but on bare metal



(slide 34)





Future Work and Considerations

In General

- Best practices and reference architectures created
- Bare metal to tenant security architecture
- Heavy validation of the bare metal deployment use case
- Early benchmarking shows a huge virtual vs bare-metal performance gap.
- Strong customer demand for a unified solution!
- Sahara in Kilo is ready!



(slide 35)

Future Work and Considerations

Ironic

- Ironic topology awareness should be a thing by RHEL OSP 8. Increased hardware Ready State coverage Ironic is a supported component in RHEL OSP 7 • Top of Rack switch configurations via Neutron?

RHEL OpenStack Platform director - the new deployment and management tool for RHEL OSP

- Sahara deployment by RHEL OSP director will be a 7.x feature • RHEL OSP director should have a tenant facing Ironic story







Questions?





Resources

- Isn't it Ironic? The Bare Metal Cloud
- Benchmarking Sahara-based Big Data as a Service solutions
- Bare Metal Hadoop and OpenStack: Together at Last!

https://www.openstack.org/summit/vancouver-2015/summit-videos/presentation/isn-and-039t-it-ironic-the-bare-metal-cloud

https://www.openstack.org/summit/vancouver-2015/summit-videos/presentation/benchmarking-sahara-based-big-data-as-a-service-solutions

https://www.openstack.org/summit/vancouver-2015/summit-videos/presentation/bare-metal-hadoop-and-openstack-together-at-last







LEARN. NETWORK. **EXPERIENCE OPEN SOURCE.**

#redhat #rhsummit

RED HAT SUMMIT



• • •