




# Hadoop on OpenStack Cloud



Dmitry Mescheryakov  
Software Engineer, @MirantisIT

# Agenda

- OpenStack
- Sahara
- Demo
- Hadoop Performance on Cloud
- Conclusion

# OpenStack

Open source cloud computing platform

17,209 commits by 1202 people for Icehouse release\* (6 month dev cycle)

Top 20 contributing companies include Red Hat, IBM, HP, Rackspace, VMWare, Intel, Samsung and others\*



# OpenStack

- OpenSource from the very beginning (Apache 2.0)
- All pythonic, services exposed via REST API
- Is split into a number of projects
  
- Scalable
- Supports various deployment modes
- Flexibility in choice of underlying technologies
  - There is always an open source choice available

# OpenStack Identity Service

- Authenticates / Authorizes users
- Provides multi-tenancy
- Provides interface for managing users & tenants
- Single entry point for OpenStack users. To use OpenStack you need to know:
  - username
  - password
  - tenant name
  - Identity API URL

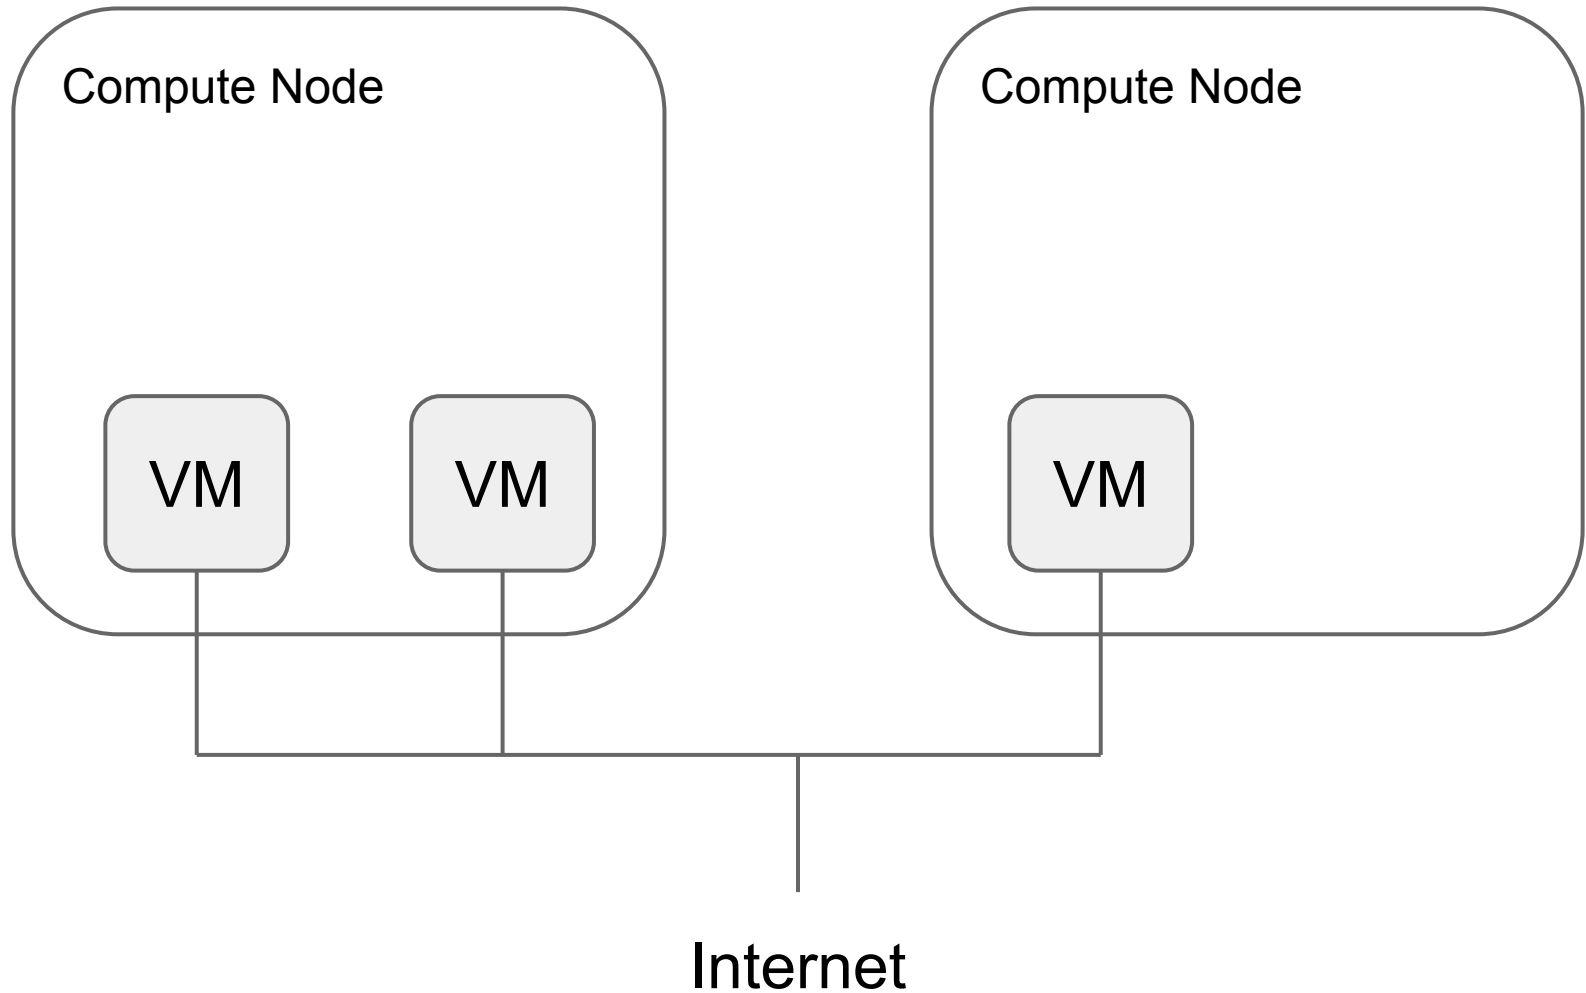
# OpenStack Compute



# OpenStack Compute

- Virtual Machines lifecycle management
- Supported hypervisors:
  - QEMU/KVM
  - Xen
  - LXC
  - Hyper-V
  - VMWare

# OpenStack Networking





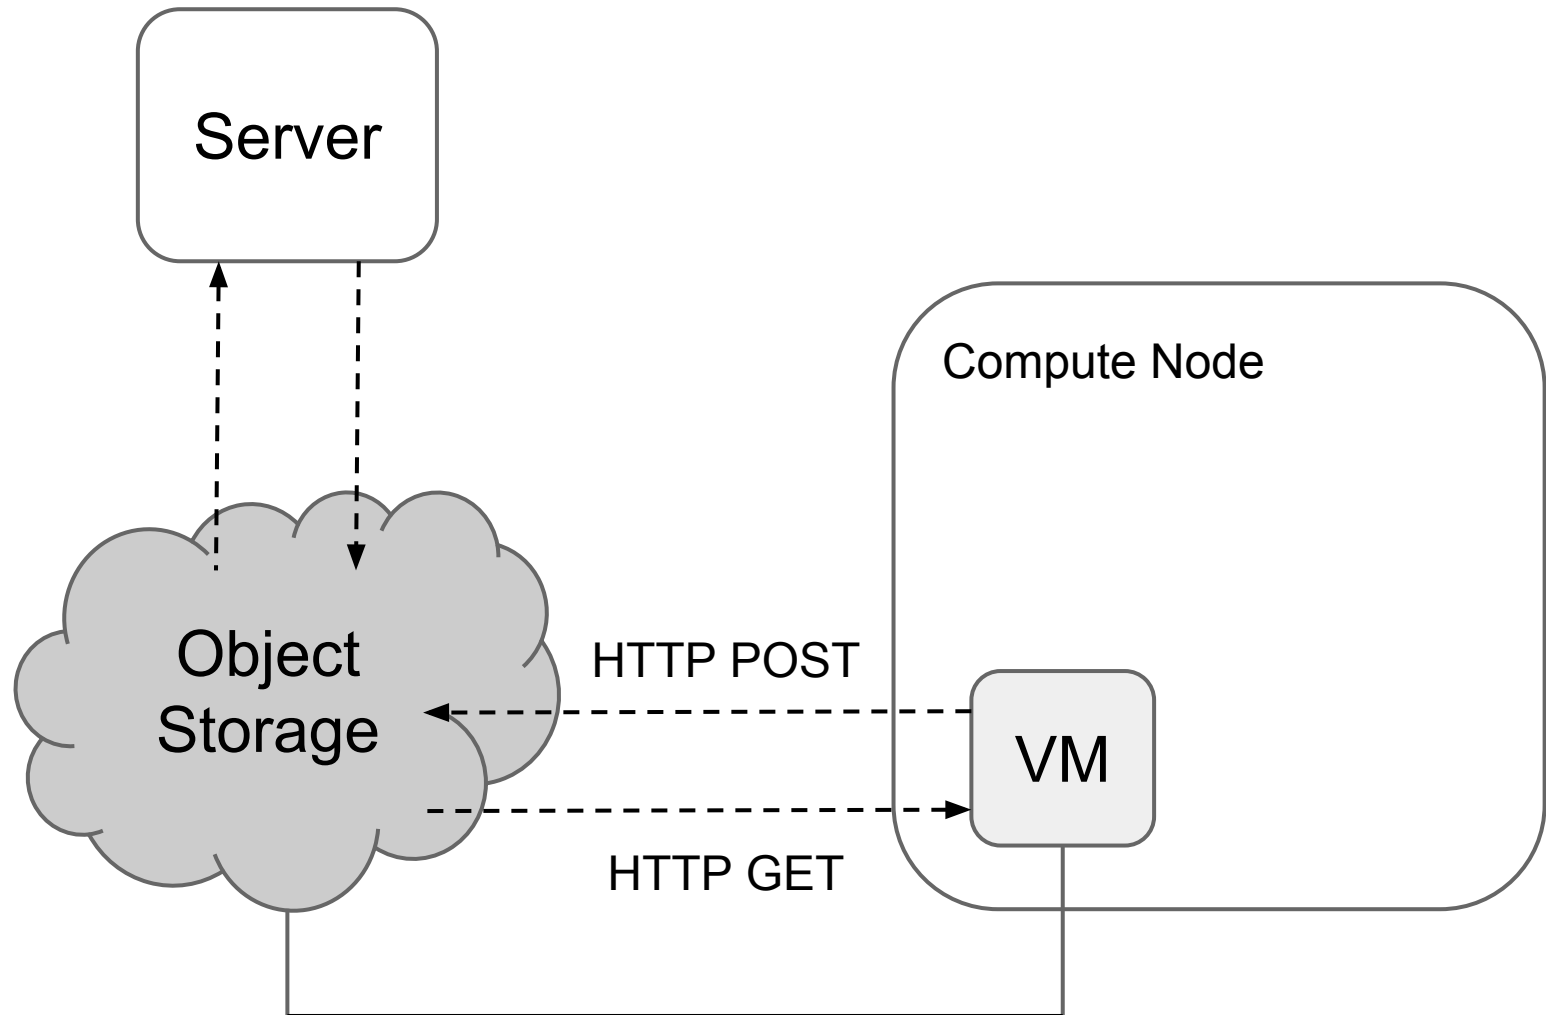
# OpenStack Networking

- Provides networking for VMs using two concepts:
  - virtual network
  - virtual router
- Networking plugin:
  - Open vSwitch
  - Cisco
  - Brocade
  - BigSwitch
  - And many more...

# OpenStack Image Service

- Image catalog for Compute
- Supported backends:
  - Local FS
  - OpenStack Object Storage
  - GridFS
  - Ceph RBD
  - And some more...

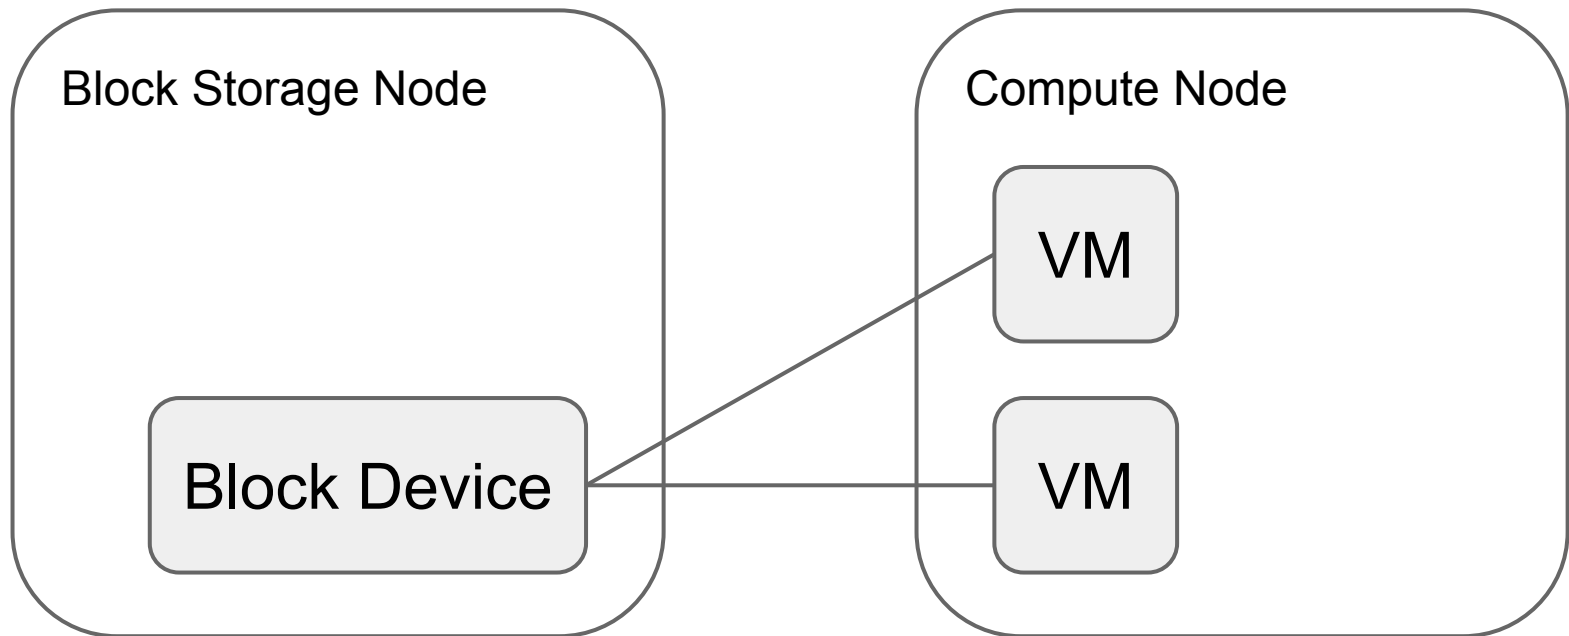
# OpenStack Object Storage



# OpenStack Object Storage

- Storage of unstructured data
- Swift, could be replaced with Ceph

# OpenStack Block Storage



# OpenStack Block Storage

- Provides persistent block storage (plug your SAN here)
- Storage plugins:
  - LVM
  - Ceph RADOS
  - Coraid AoE
  - Dell EqualLogic
  - And many more ...

# OpenStack Dashboard

- OpenStack Web UI

# AWS vs OpenStack

Amazon	OpenStack
--------	-----------

EC2	Compute
-----	---------

	Networking
	Image Service

Identity & Access Manager	Identity Service
---------------------------	------------------

S3	Object Storage
----	----------------

Elastic Block Storage	Block Storage
-----------------------	---------------

Web UI	Dashboard
--------	-----------



# OpenStack API

`http://127.0.0.1:8774/v2/1e2afda0.../servers`

`X-Auth-Token: 2c1ecf5...`

```
{"server": {"name": "my-instance",  
            "imageRef": "fe35ee17-...",  
            "key_name": "my-keypair",  
            "flavorRef": "2",  
            "networks": [{"uuid": "bb80cc75-..."}]}}
```

# OpenStack CLI Clients

```
nova boot \  
  --image ubuntu-14.04 \  
  --key-name my-keypair \  
  --flavor m1.small \  
  --nic net-id=bb80cc75-... \  
my-instance
```

OS\_USERNAME, OS\_PASSWORD, OS\_TENANT\_NAME, OS\_AUTH\_URL environment variables must be defined

# OpenStack Python Bindings

```
from novaclient import client
```

```
nova = client.Client('2', 'admin', 'nova', 'admin', 'http://127.0.0.1:5000/v2.0/')
```

```
nova.servers.create('my-instance',  
                    image='fefbee17-...',  
                    flavor='2',  
                    key_name='my-keypair',  
                    nics=[{"net-id": 'bb80cc75-...'}])
```

# DevStack

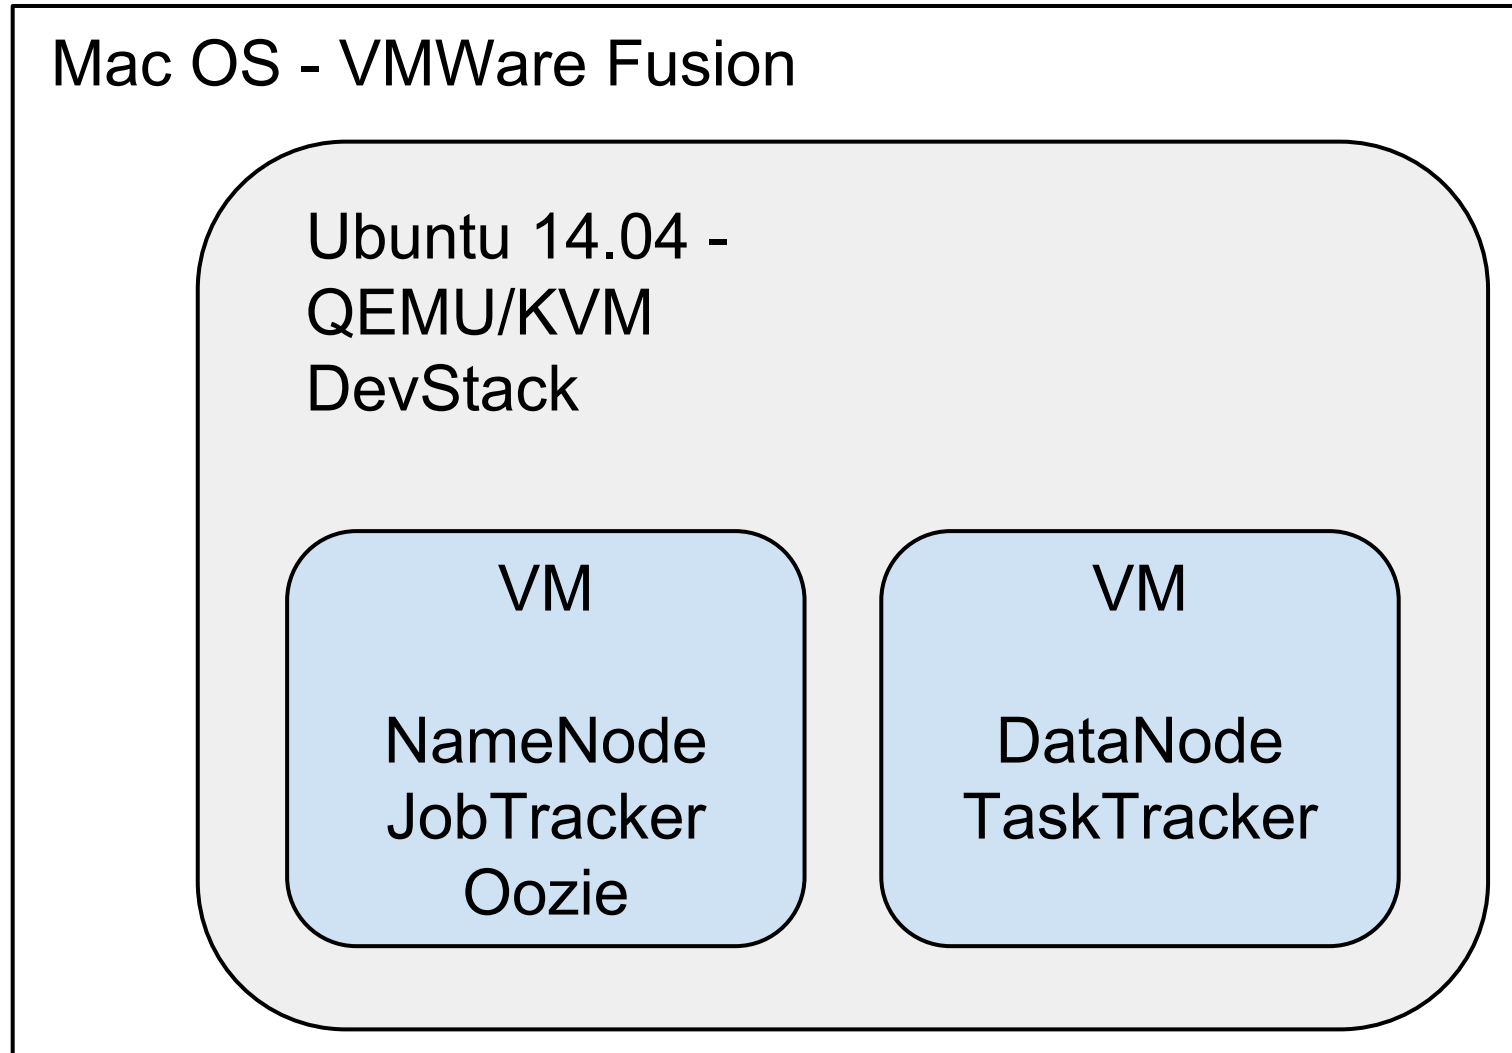
All-in-one OpenStack installation for dev and demo purposes

<http://devstack.org>, and follow instructions

To enable Sahara:

<http://docs.openstack.org/developer/sahara/devref/devstack.html>

# DevStack Demo Environment



# DevStack on VM: a tip

Host hypervisor should pass through hardware virtualization: QEMU/KVM for Linux, VMWare Fusion for Mac OS X.

VT-x (vmx) for Intel, AMD-V (svm) for AMD

Without it, nested VMs will be very slow. To check:

```
cat /proc/cpuinfo | grep --color "vmx\|svm"
```

# Sahara (ex. Savanna): OpenStack Data Processing

Simplify running Hadoop on OpenStack

Started a year ago and currently major contributors include Mirantis, Red Hat and Hortonworks

Will be integrated project in OpenStack Juno release (October 2014)

# Sahara Overview

- template based cluster provisioning
- different distributions via plugins:
  - Vanilla Hadoop
  - HDP
  - CDH (in progress)
- Each plugin supports several versions



# Supported Hadoop Ecosystem Projects

- HDFS
- MapReduce
- YARN
- Oozie
- Hive

# Sahara Functionality

- Bringing up cluster
- Configure it along the way
- Scale cluster
- Terminate cluster
- Job execution (Elastic Data Processing)

# Integration with Object Storage

Work with Object Storage like with HDFS

- `swift://test-container.sahara/my_file`
- username
- password
- tenant name

# Prepared Images

- Take cloud image (Ubuntu, Fedora, CentOS) as a base
- Install Hadoop, Java and other stuff on it
- Enjoy much faster cluster provisioning

# Data Locality

Sahara can provide data locality info, if configured properly

Works for both HDFS and Object Storage

VMs running on the same hardware machines are 'close', and Sahara knows that

# Other Stuff

- REST API
- CLI client
- Python bindings
- UI

# Hadoop in the Cloud: Performance

- Mirantis OpenStack Express cluster
- 20 nodes
- CPU: 24 x 2.10 GHz (2 x Intel Xeon CPU E5-2620)
- Memory: 8 x 4.0 GB, 32.0 GB total
- Disk: 1 drive, 0.9 TB (WDC WD1003FBYX-0)
- Network: 2 x 1 GbE

# Performance tests

- disk read/write
- network throughput
- cpu
- composite test

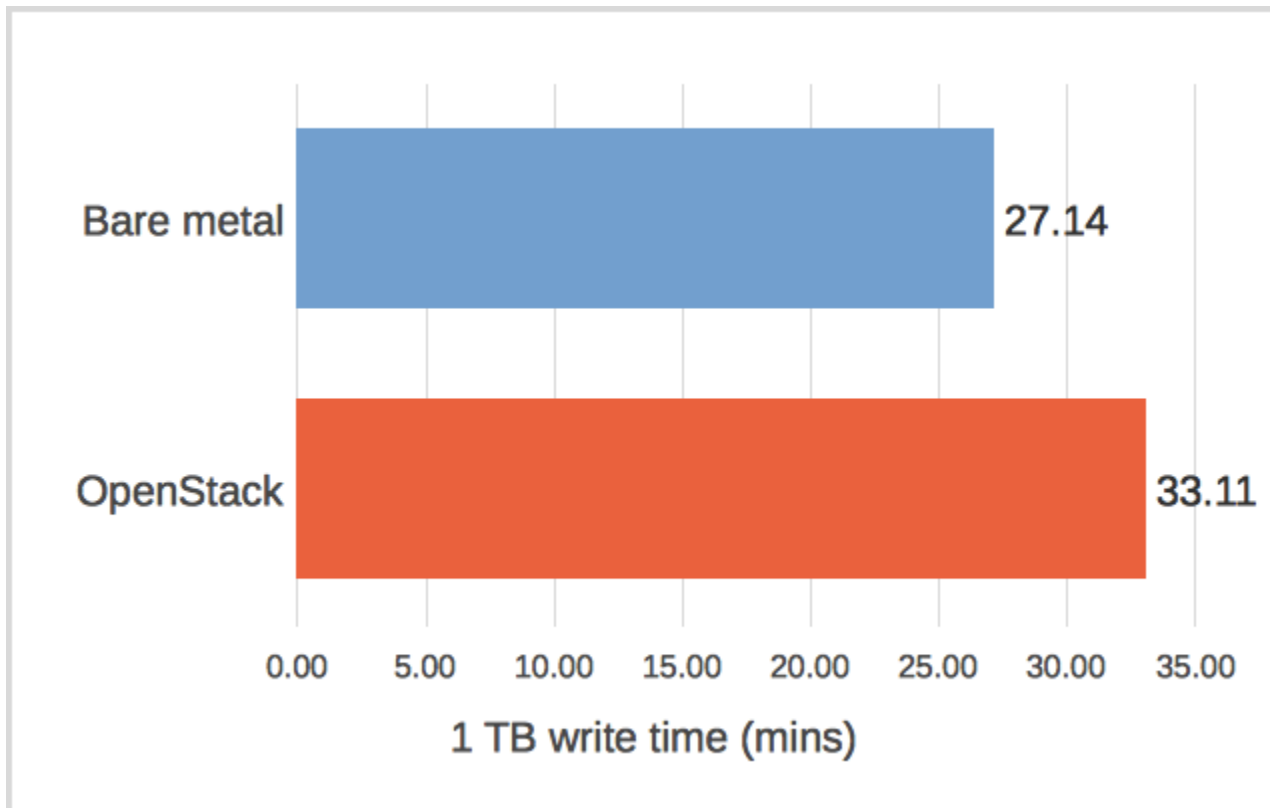


# Disk Read/Write

TestDFSIO - built-in hadoop I/O test


- 1000 files of 1GB (1 TB total)

# Disk Write



\*less is better

# Disk Read




\*less is better

# Network

time + nc

# Network



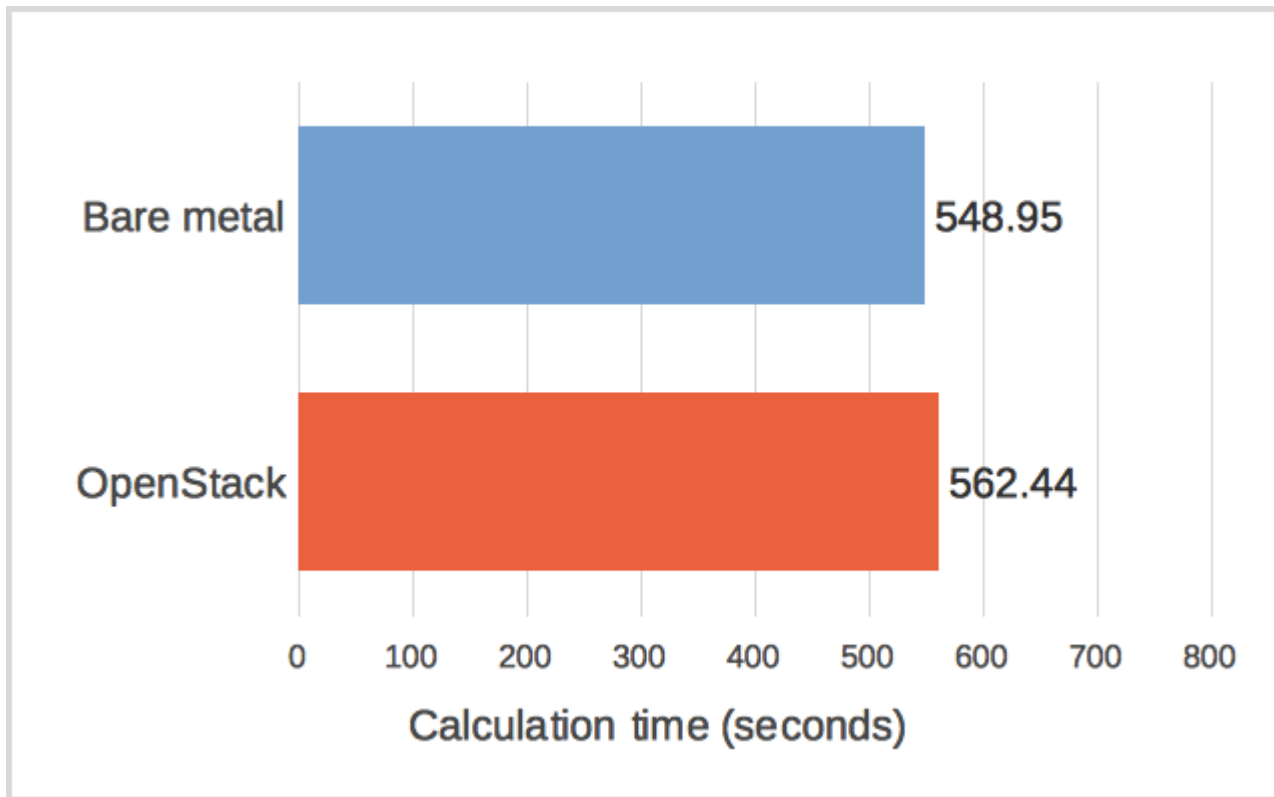
\*greater is better

# CPU

PI - built-in hadoop test, depends mostly on CPU

- 50 series of 10,000,000,000 probes

# CPU



\*less is better

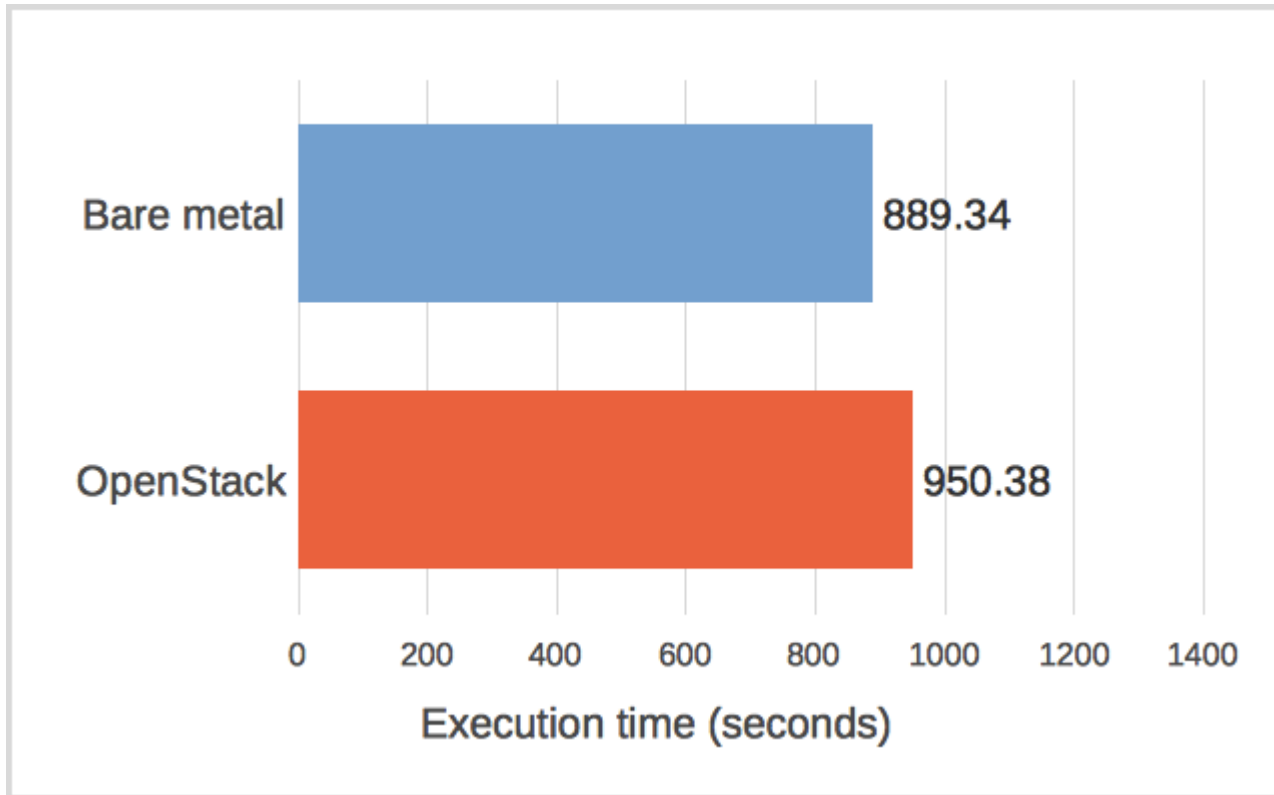
# Composite Text

Terasort - built-in hadoop test

- 200,000,000 of 100-byte entries (20 GB)



# Terasort



\*less is better

# Performance Testing Results

Virtualized Hadoop 24% slower than Bare Metal one in the worst case (disk read)

It is only 6% slower with the composite test (Terasort)

More details in talk “Performance of Hadoop on OpenStack” by Andrew Lazarev (find it on youtube)

# Why Sahara

- agility
- self-service
- multi-tenancy
- pay as you go

# Q&A