



Red Hat Reference Architecture Series

Deploying OpenStack Sahara

on Red Hat Enterprise Linux OpenStack Platform 6

Jacob Liberman, Principal Software Engineer

RHCE

Version 2.0

April 2015



100 East Davie Street
Raleigh NC 27601 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park NC 27709 USA

Linux is a registered trademark of Linus Torvalds. Red Hat, Red Hat Enterprise Linux and the Red Hat "Shadowman" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

Hadoop and the Hadoop elephant logo are trademarks of the Apache Software Foundation.

OpenStack is a trademark of the OpenStack Foundation.

PowerEdge is a trademark of Dell, Inc.

Intel, the Intel logo and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

© 2015 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

The information contained herein is subject to change without notice. Red Hat, Inc. shall not be liable for technical or editorial errors or omissions contained herein.

Distribution of modified versions of this document is prohibited without the explicit permission of Red Hat Inc.

Distribution of this work or derivative of this work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from Red Hat Inc.

The GPG fingerprint of the security@redhat.com key is:
CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

Send feedback to refarch-feedback@redhat.com





Comments and Feedback

In the spirit of open source, we invite anyone to provide feedback and comments on any reference architectures. Although we review our papers internally, sometimes issues or typographical errors are encountered. Feedback allows us to not only improve the quality of the papers we produce, but allows the reader to provide their thoughts on potential improvements and topic expansion to the papers.

Feedback on the papers can be provided by emailing refarch-feedback@redhat.com. Please refer to the title within the email.

Staying In Touch

Join us on some of the popular social media sites where we keep our audience informed on new reference architectures as well as offer related information on things we find interesting.

Like us on Facebook:

<https://www.facebook.com/rhrefarch>

Follow us on Twitter:

<https://twitter.com/RedHatRefArch>

Plus us on Google+:

<https://plus.google.com/u/0/b/114152126783830728030/>

Table of Contents

1. Executive Summary.....	1
2. Introduction.....	2
3. Architecture Overview.....	3
3.1 OpenStack Sahara.....	3
3.2 Solution Overview.....	3
3.3 Software Component Overview.....	4
3.3.1 Red Hat Enterprise Linux 7.....	4
3.3.2 Apache Hadoop.....	4
3.3.2.1 Hadoop Common.....	4
3.3.2.2 Hadoop Distributed File System (HDFS).....	5
3.3.2.3 Hadoop YARN.....	5
3.3.2.4 Hadoop MapReduce.....	5
3.3.2.5 Apache Oozie.....	5
3.3.2.6 Apache Pig.....	5
3.3.3 RHEL OpenStack Platform 6.....	5
3.3.4 OpenStack Services.....	5
3.3.4.1 Identity Service (“Keystone”).....	5
3.3.4.2 Orchestration Service (“Heat”).....	6
3.3.4.3 Image Service (“Glance”).....	6
3.3.4.4 Compute Service (“Nova”).....	6
3.3.4.5 Network Service (“Neutron”).....	6
3.3.4.6 Block Storage Service (“Cinder”).....	6
3.3.4.7 Object Storage Service (“Swift”).....	6
3.3.4.8 Dashboard Service (“Horizon”).....	7
3.3.4.9 OpenStack Data Processing (“Sahara”).....	7
3.4 Server Roles.....	7
3.4.1 Cloud Controller.....	7
3.4.2 Compute Node.....	8
3.4.3 Network Node.....	8
3.4.4 Swift Servers.....	8
3.4.4.1 Swift Proxy Server.....	8
3.4.4.2 Swift Storage Server.....	8
3.5 Network Names.....	8
3.6 Conceptual Diagram of the Solution Stack.....	9
3.7 Sahara Terminology.....	9



3.8 Submitting Sahara Jobs.....	10
3.8.1 Sahara Workflow.....	10
3.8.2 Submitting Sahara Jobs.....	10
4. Configuration Details.....	12
4.1 Environment.....	12
4.1.1 Service Placement.....	12
4.1.2 Network Topology.....	13
4.1.3 Network Addresses.....	14
4.2 Software Configuration.....	16
4.2.1 Required Channels.....	16
4.2.2 Software Versions.....	16
4.3 Security Details.....	18
4.3.1 Firewall Configuration.....	18
4.3.2 SELinux Configuration.....	18
4.4 Hardware Details.....	18
4.4.1 Server Hardware Configuration.....	18
4.4.2 Network Hardware.....	19
5. Implementing Sahara on OpenStack.....	20
5.1 Prepare the Hosts.....	20
5.2 Deploy OpenStack.....	20
5.3 Prepare the Environment.....	21
5.4 Install Sahara.....	28
5.5 Configure Sahara.....	31
5.6 Import Glance Images for Hadoop.....	32
5.7 Define Templates.....	34
5.7.1 Define Master Group Template.....	34
5.7.2 Define the Worker Group Template.....	35
5.7.3 Define Cluster Templates.....	36
5.8 Launch a Cluster via the Command Line.....	38
6. Validate and Test Sahara.....	41
6.1 Explore the Cluster.....	41
6.2 Run a Test Hadoop Job from the Command Line.....	44
6.3 Run a Test EDP Job Through Sahara.....	46
6.4 Verify Dashboard.....	52

7. Conclusion.....	55
Appendix A: References.....	56
Appendix B: Revision History.....	57



1. Executive Summary

Big data and analytics infrastructure spending continues to grow as businesses transform to being data driven. Apache Hadoop and other open-source frameworks for large-scale distributed data storage and processing are attractive to enterprise customers because they can run on commodity clusters. However, the cost to design, deploy, test, and maintain Hadoop clusters can quickly erode their favorable economics.

The Sahara project provides a simple means to provision Hadoop clusters in OpenStack. Sahara users can deploy Hadoop clusters in minutes by specifying simple parameters such as Hadoop version, cluster topology, and node count. The pre-configured clusters are deployed with all required software and libraries. Clusters provisioned by Sahara can be scaled on demand or removed when no longer needed. The Sahara database stores reusable cluster and job templates.

Typical use cases for OpenStack Sahara include:

- Rapid provisioning of Hadoop clusters for Dev and QA
- Utilization of unused compute power from a general purpose OpenStack cloud
- “Analytics-as-a-Service” for bursty or ad-hoc workloads

This reference architecture describes how to install and configure Sahara on Red Hat Enterprise Linux OpenStack Platform 6. It focuses on the single user to small group use case of deploying a Hadoop cluster for Dev or QA. It also demonstrates two methods for running Hadoop jobs within Sahara. First, it shows how to run an interactive MapReduce job to validate cluster functionality. It concludes with an example Pig job submitted through Sahara's Elastic Data Processing (EDP) engine that uses OpenStack Swift for input and output data storage. Although Sahara natively benefits from many of the fault tolerant features of both Hadoop and OpenStack, deploying Sahara for production is beyond the scope of this document.



2. Introduction

The Sahara project provides a scalable data processing stack and associated management technologies for OpenStack. Apache Hadoop is an industry standard MapReduce implementation. Sahara users can easily provision and manage Hadoop clusters on OpenStack. This allows Hadoop users to leverage the favorable economics of scale-out computing while at the same time reducing the complexity of cluster installation and management. Sahara is designed as an OpenStack component whose key features include:

- Self-service cluster deployment
- A template-based framework for defining and managing clusters
- Support for various job types including MapReduce, Hive, and Pig
- Support for various data sources including Swift and Hadoop Distributed File System (HDFS)
- Integration with OpenStack services and management tools such as Horizon, Glance, Heat, Cinder, and Keystone

This paper describes how to deploy Sahara on Red Hat Enterprise Linux OpenStack Platform (RHEL OSP) 6. It includes a description of Sahara's core components, an overview of how Sahara fits into the OpenStack software ecosystem, and complete steps for installing Sahara. It also includes instructions for running two test jobs to verify Sahara functionality.

NOTE: RHEL OSP 6 is based on the OpenStack Juno release.



3. Architecture Overview

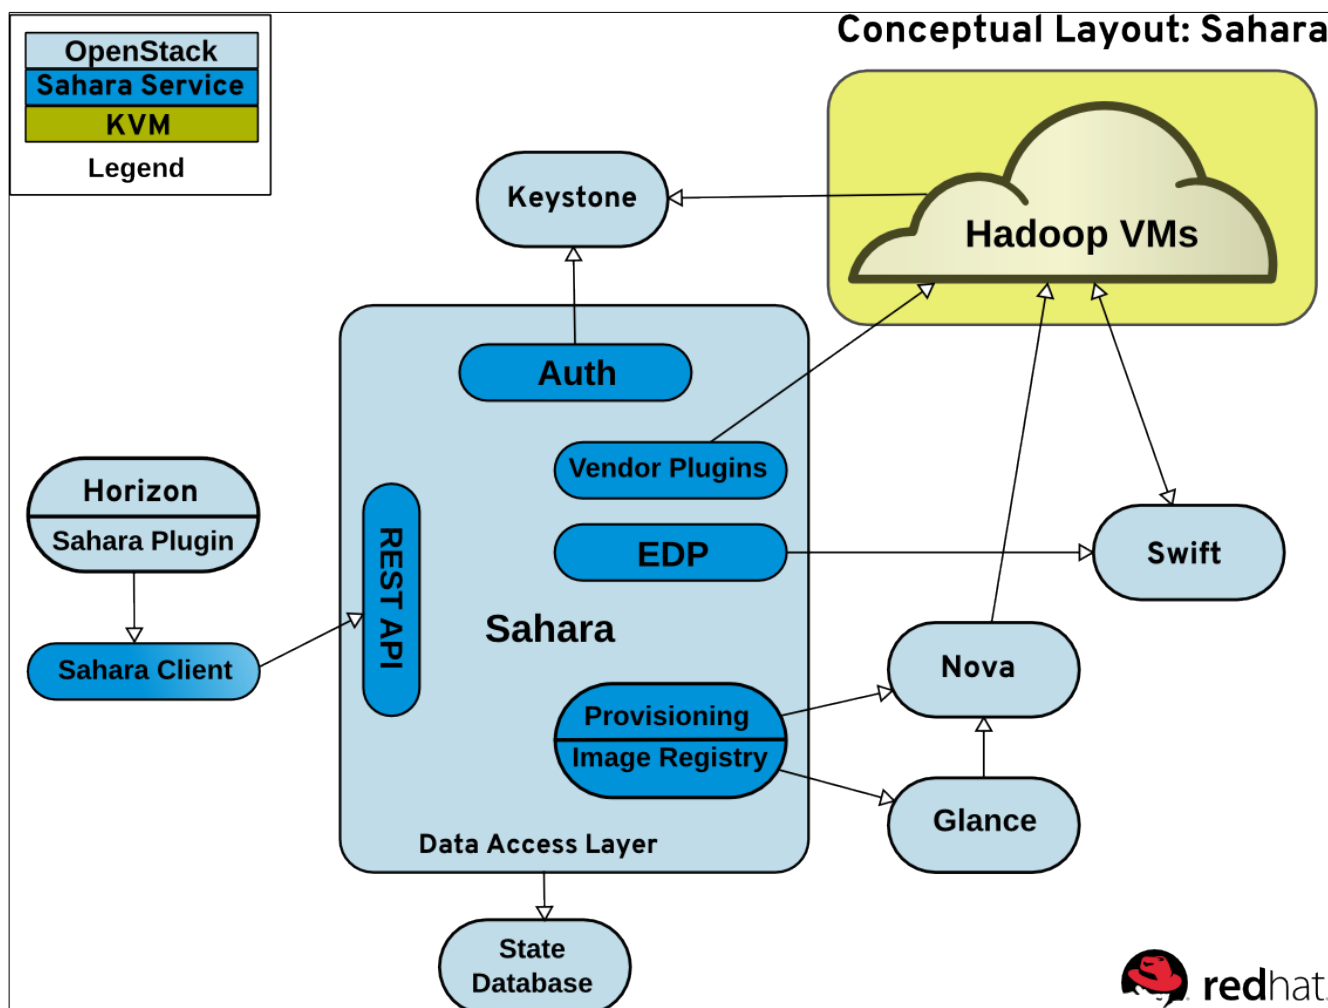
This section describes the reference architecture both physically and conceptually. It also defines terms used in the rest of the document.

3.1 OpenStack Sahara

The Sahara project provides users with a simple means to define and provision Hadoop clusters within OpenStack. It is designed as an OpenStack component and managed through a REST API. Sahara supports several Hadoop distributions and versions and integrates with vendor-specific management tools.

3.2 Solution Overview

This section introduces Sahara and summarizes how it interacts with OpenStack. **Graphic 3.2.1: Sahara Conceptual Layout** depicts the relationships between services.



Graphic 3.2.1: Sahara Conceptual Layout



- ✦ **Auth** -- The Sahara auth component is responsible for client authentication and authorization. It communicates with Keystone as a service endpoint.
- ✦ **Data Access Layer** – The Data Access Layer stores persistent models in an internal database.
- ✦ **EDP** – The Elastic Data Processing (EDP) component schedules and manages Hadoop jobs on clusters provisioned by Sahara. The EDP component can use Swift or HDFS as an input and/or output data store for Hadoop jobs.
- ✦ **Image Registry** – The Image Registry component interacts with Glance to ensure images meet the user's requirements when a cluster is formed. Image requirements for Sahara clusters depend on the Hadoop version and plugin.
- ✦ **Provisioning Engine** – The Sahara provisioning agent communicates with Nova, Neutron, Heat, Cinder and Glance to instantiate virtual machines and storage for Hadoop clusters.
- ✦ **REST API** – The Sahara REST API exposes Sahara functionality via REST to management tools such as the Python Sahara client.
- ✦ **Sahara Client** -- Sahara has a Python client similar to other OpenStack components.
- ✦ **Sahara Dashboard** – The Sahara dashboard allows users to manage Sahara via the web-based OpenStack Dashboard (Horizon).
- ✦ **Vendor Plugins** – These are pluggable mechanisms responsible for configuring and launching Hadoop on provisioned virtual machines.

Refer to the OpenStack Sahara documentation for more information regarding Sahara architecture: <http://docs.openstack.org/developer/sahara/architecture.html>

3.3 Software Component Overview

This section describes the software components used to develop the reference architecture.

3.3.1 Red Hat Enterprise Linux 7

Red Hat Enterprise Linux (RHEL) 7 is the base operating system used in this reference architecture. All servers run RHEL 7, the latest major release of Red Hat's flagship platform. RHEL 7 provides a stable code base for bare metal servers, virtual machines, Infrastructure-as-a-Service (IaaS), and Platform-as-a-Service (PaaS) across the enterprise data center.

3.3.2 Apache Hadoop

The Apache Hadoop software library is a framework allowing the distributed processing of large data sets across clusters of computers using simple programming models. It is designed to be scalable to thousands of machines. The Hadoop core modules and related projects mentioned in this reference architecture are described below.

3.3.2.1 Hadoop Common

The common utilities that support other Hadoop modules.



3.3.2.2 Hadoop Distributed File System (HDFS)

A fault-tolerant and scalable distributed file system. HDFS provides streaming access for large data sets. Data replication schemes and error detection ensure quick recovery from hardware failure.

3.3.2.3 Hadoop YARN

YARN is a resource manager that acts as a data operating system for Hadoop. YARN allows multiple applications to run natively in Hadoop.

3.3.2.4 Hadoop MapReduce

Hadoop MapReduce is a software framework for writing applications that process large amounts of data in parallel on clusters of commodity hardware. MapReduce jobs split input data into chunks that can be operated on independently by map tasks. The output of these jobs are input to reduce tasks that recombine the data to yield a result. MapReduce is a dominant paradigm for large-scale data analysis.

3.3.2.5 Apache Oozie

Oozie is a scalable and extensible workflow scheduler system for Hadoop jobs. It is integrated with Hadoop to support several types of jobs including Java, MapReduce and Streaming MapReduce, Pig, and Hive.

3.3.2.6 Apache Pig

Pig consists of a high level language for expressing and evaluating large scale data sets. Pig simplifies the tasks of MapReduce parallel programming.

3.3.3 RHEL OpenStack Platform 6

OpenStack is open source software for building private and public clouds. Red Hat is an active contributor to the OpenStack code base¹. RHEL OSP 6 combines the benefits of Red Hat's OpenStack technology, Kernel-based Virtual Machine (KVM), and Red Hat Enterprise Linux. RHEL OSP 6 is based on the tenth OpenStack release code named "Juno". RHEL OSP 6 is certified to run on Red Hat Enterprise Linux 7+.

NOTE: See the product release notes for more information:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/6/html/Release_Notes/

3.3.4 OpenStack Services

Sahara integrates with the following OpenStack services:

3.3.4.1 Identity Service ("Keystone")

This is the central authentication and authorization mechanism for all OpenStack users and

¹ <http://www.redhat.com/infographics/openstack/>



services. Users and services interact with Sahara via a Keystone service endpoint. Furthermore, Keystone tenants manage access to Sahara resources including clusters, nodes, and templates.

3.3.4.2 Orchestration Service (“Heat”)

Heat implements an orchestration engine to launch multiple composite cloud applications based on text template files. In OpenStack Juno, Sahara is configured to use a direct engine for instance creation that makes direct calls to the services required for instance provisioning. Sahara can be configured to use the OpenStack Orchestration service for this task instead of the direct engine. They are at feature parity.

3.3.4.3 Image Service (“Glance”)

This service registers and delivers virtual machine images. Sahara cluster nodes can be installed via pre-built images of vanilla Apache Hadoop or vendor supplied images. Users can interact with the pre-built images via the Sahara API once they are registered with Glance.

3.3.4.4 Compute Service (“Nova”)

The Compute service provisions and manages large networks of virtual machines. Sahara leverages the Nova framework to define and build Hadoop clusters via user defined templates. The templates include node and cluster specific information such as the plugin name, Hadoop version, and Nova hardware flavor. Sahara node group templates are passed to Nova along with a Glance image name to instantiate virtual machine nodes. Cluster templates combine the nodes into clusters.

3.3.4.5 Network Service (“Neutron”)

The Network service is a scalable API-driven service for managing networks and IP addresses. Neutron gives users self-service control over their network configurations. Sahara templates define the network access method for cluster nodes. They are accessed most commonly via Neutron floating IP addresses assigned to the instances during creation.

3.3.4.6 Block Storage Service (“Cinder”)

The Block Storage service provides services and libraries that implement on demand, self-service access to Block Storage resources for OpenStack. The Compute service provides Sahara clusters with ephemeral storage by default. Ephemeral storage disappears when a virtual machine is terminated. Block storage persists regardless of the state of running instances and can be moved between instances. OpenStack Block Storage can be used as an alternative to ephemeral drives for Sahara instances.

3.3.4.7 Object Storage Service (“Swift”)

The Object Storage service provides a fully distributed, API-accessible object storage platform that can be integrated into applications or used for backup, archiving, and data retention. Swift data stores can be used within Hadoop data processing jobs after the application of a patch. Sahara automatically sets information about the Swift filesystem implementation, location awareness, and tenant name for authorization. The only information required when launching a Hadoop job with a Swift backend are the username and password to access Swift. Sahara can use Swift Object Stores to store and retrieve input data, job



binaries, and output data.

3.3.4.8 Dashboard Service (“Horizon”)

The Dashboard service is an extensible web-based application that allows cloud administrators and users to control and provision compute, storage, and networking resources. A Sahara dashboard plugin allows administrators to deploy and manage Sahara clusters via the Dashboard.

3.3.4.9 OpenStack Data Processing (“Sahara”)

The Data Processing service provides a simple means to provision a data-intensive application cluster running Hadoop or Spark on top of OpenStack. The Data Processing service in RHEL OpenStack Platform 6 includes the following features:

- Anti-affinity: This feature leverages Nova server groups to ensure cluster instances are spawned on separate physical hosts. This can increase HDFS replica reliability.
- Cluster scaling: This feature allows users to increase or decrease the number of nodes in a cluster without recreating the cluster.
- Data-locality: Data locality can improve the performance of data intensive applications. This feature can schedule jobs to nodes that are local for the input stream at the rack or volume level. The data locality feature can be enabled for both Block and Object data stores.
- Distributed mode: Sahara services can be run as a single process or the sahara-api and sahara-engine subprocesses can be launched in parallel on one or more nodes. Running in distributed mode can balance requests and engine processes across multiple nodes for redundancy.
- Networking support: Sahara supports both Nova and Neutron networking.

NOTE: Additional Sahara features are described at <http://docs.openstack.org/developer/sahara/userdoc/features.html>.

3.4 Server Roles

This section defines the various server roles used in this reference architecture.

3.4.1 Cloud Controller

Cloud controller is the designation used in this reference architecture for the server that provides the endpoint for REST-based API queries to the majority of the OpenStack services. These include Compute, Image, Identity, Block, Network, and Data processing. Although RHEL OSP allows for multiple, high availability cloud controllers, only one cloud controller is used in this reference architecture.



3.4.2 Compute Node

Compute node refers to an OpenStack server that runs a KVM hypervisor. It is responsible for running virtual machine instances. In this reference architecture, Hadoop clusters are instantiated across multiple compute nodes. By default a new instance is spawned on the compute node with the most free memory in a round robin fashion.

3.4.3 Network Node

The Network Node provides centralized networking control for all compute nodes and tenants. It runs the various Neutron agents that control L3 networking functionality in the cluster. In this reference architecture, the network agents run on a single dedicated server. It is also possible to run these agents on the cloud controller or multiple network nodes in a clustered fashion.

3.4.4 Swift Servers

Understanding the Swift ring is central to understanding the role of a server in the Swift cluster. A Swift ring represents a mapping between the names of entities stored on disk and their physical location. There are separate rings for accounts, containers, and objects. When a component needs to interact with an object, container, or account, it interacts with the appropriate ring to determine the target's location in the cluster.

3.4.4.1 Swift Proxy Server

This server ties together the Swift architecture. For each request, the proxy server looks up the location of the account, object, or container in the appropriate Swift ring and routes the request accordingly. In this reference architecture, the cloud controller also acts as the Swift proxy server.

3.4.4.2 Swift Storage Server

The Swift storage servers are simple blob storage servers that can store, retrieve, and delete objects stored on local devices. Objects are stored using a path derived from the object's name and a time stamp. In this reference architecture there are three dedicated Swift storage servers. They are configured as three-way replication partners to ensure data coherency. It is also possible to run Swift object servers directly on the Compute nodes to improve data locality for Sahara.

3.5 Network Names

A typical OpenStack deployment includes several network roles. In some cases the roles overlap across the same physical interfaces or switches. In others each role has a dedicated network interface and switch. This reference architecture uses the following networks:

1. **External network** – the external network is used to perform system maintenance tasks such as installing software. In a private cloud scenario, users access the cloud infrastructure via the external network.
2. **Service network** – this network exposes the OpenStack APIs. It also handles inter-service communication between the OpenStack services and schedulers.
3. **Tenant network** – virtual machines communicate over this network within the cloud

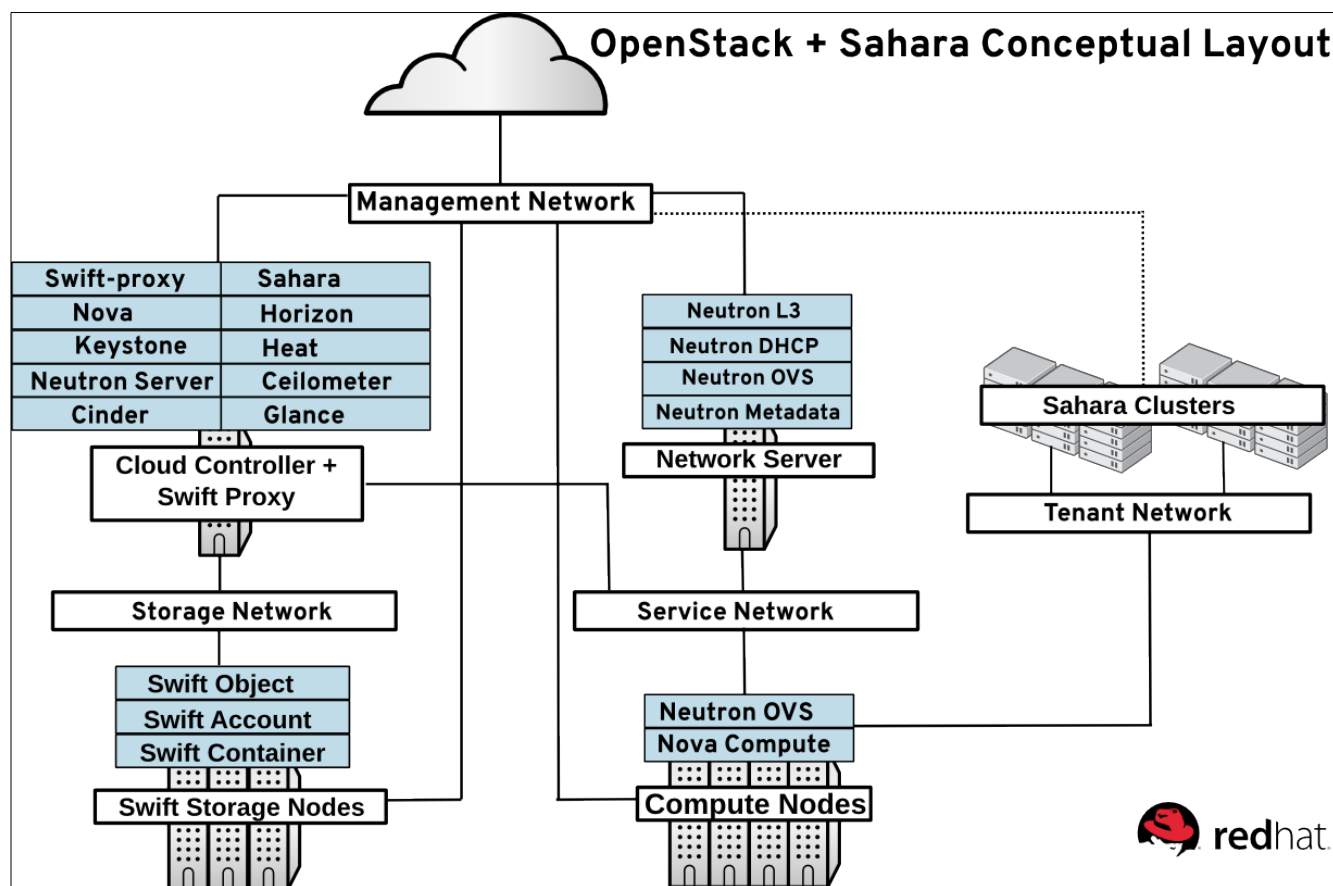


deployment. The addressing requirements of this network depend on the plugin that is used.

4. **Storage Network** – this network is dedicated for storage traffic between the Swift servers and the OpenStack servers.

3.6 Conceptual Diagram of the Solution Stack

3.6.1 Service Placement depicts the solution stack including networks, server roles, and service placement. Section 4 Configuration Details shares complete details.



Graphic 3.6.1: Service Placement

NOTE: Minor variations in OpenStack service placement will not affect the overall solution. For example, collapsing the Network Server and Cloud Controller services to a single system or separating the Swift proxy to a dedicated server are both acceptable changes.

3.7 Sahara Terminology

Sahara uses collections of simple objects to define and execute Hadoop jobs. The objects are stored in the Sahara database and can be reused on subsequent jobs. A job consists of:



1. At least one executable binary
2. Optional supporting libraries
3. Input data
4. Output location
5. Additional configuration values needed to run the job

The job binary object stores a resource locator to a single JAR file or script and the credentials needed to retrieve it. A job refers to the binary when it is launched and all its supporting libraries. The job and job binary are analogous to an instance and an image: multiple jobs can share the same job binary just as multiple instances can be spawned from the same image. The data source objects designate the location of any input or output data required or generated by the job. Sahara supports both Swift and HDFS as data sources.

3.8 Submitting Sahara Jobs

This section describes the typical workflow for running Sahara jobs. It also describes the typical workflow for a single job.

NOTE: More information on these topics can be found in the Sahara documentation: <http://docs.openstack.org/developer/sahara/userdoc/edp.html>.

3.8.1 Sahara Workflow

A lifecycle of a Sahara job follows a general workflow:

- ⤴ Define cluster node templates that reference a Hadoop version, an image, and the number and type of Hadoop services to run
- ⤴ Define a cluster template that references cluster node templates
- ⤴ Launch a cluster from the cluster template
- ⤴ Once the cluster is instantiated, create all the job binaries needed to run the job.
- ⤴ Store the job binaries in the Sahara database or Swift
- ⤴ Create a job that references the job binaries
- ⤴ Create an input data source that references the data to be processed
- ⤴ Create an output data source
- ⤴ Launch the job to the cluster.

Existing objects simplify the workflow. New jobs can reference existing job binaries and already instantiated clusters.

3.8.2 Submitting Sahara Jobs

Hadoop jobs can be submitted to Sahara in several ways:

1. The user can login to the cluster and submit jobs interactively through command line interfaces for Hadoop and Oozie



2. The user can submit a job to the Oozie server via an Oozie client or the Oozie REST API
3. The user can submit a job through the Sahara client or web UI.

The third method highlights several features of using Sahara. First, job binaries and data sources can be combined into reusable templates. Users can re-run the same job multiple times via the template or share it with different users.

Next, jobs submitted through the Sahara EDP workflow can be run on transient clusters. Transient clusters are created specifically for the job and shut down once the job is finished.

Section **6 Validate and Test Sahara** demonstrates the first and third job submission methods.



4. Configuration Details

This section of the paper describes the hardware and software used to configure the reference architecture in the Red Hat Solutions Engineering lab. It includes security details.

4.1 Environment

The reference architecture environment consists of the components required to build a small Red Hat Enterprise Linux OpenStack Platform cloud infrastructure. It includes small form factor servers for the OpenStack servers and Swift storage servers with more internal storage capacity.

4.1.1 Service Placement

Table 4.1.1.1 Service Placement lists the service placement for all OpenStack services. The cloud controller runs the majority of services.

Host	Role	Service
rhos0	Cloud controller	rabbitmq
		neutron-server
		openstack-cinder-api
		openstack-cinder-scheduler
		openstack-cinder-volume
		openstack-glance-api
		openstack-glance-registry
		openstack-heat-api-cfn
		openstack-heat-api
		openstack-heat-engine
		openstack-keystone
		openstack-nova-api
		openstack-nova-cert
		openstack-nova-conductor
		openstack-nova-consoleauth
openstack-nova-novncproxy		
openstack-nova-scheduler		
openstack-sahara-all		
openstack-swift-proxy		
rhos1	Network server	neutron-dhcp-agent
		neutron-l3-agent



		neutron-metadata-agent
		neutron-openvswitch-agent
rhos[2-5]	Compute nodes	neutron-openvswitch-agent
		openstack-nova-compute
rhos[7-9]	Storage nodes	openstack-swift-account
		openstack-swift-container
		openstack-swift-object

Table 4.1.1.1: Service Placement

4.1.2 Network Topology

4.1.2.1 Network Topology shows the network topology of this reference architecture.

- ⤴ All eight servers communicate via the lab network switch on the management network. The management network uses IP addresses in the 10.19.137.0/24 range.

NOTE: There is no server named rhos6. The controller nodes are rhos[0-2], the compute nodes are rhos[3-5], and the Swift object storage nodes are rhos[7-9].

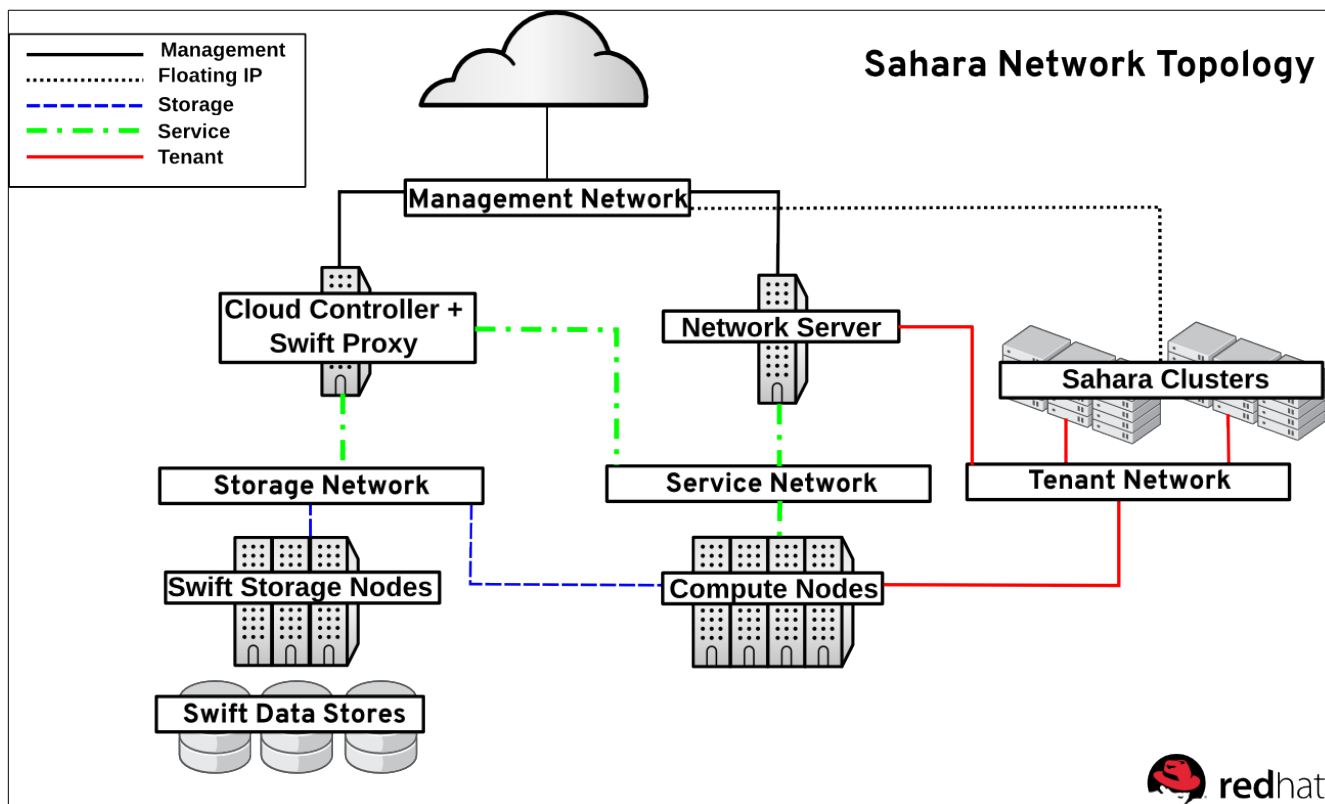
- ⤴ The tenant network carries communication between virtual machines and software-defined networking components. It is the private network over which the instances communicate. In this reference architecture, a network switch connected to 10 GB interfaces on the compute nodes is tagged to VLAN IDs 1000:1010 for tenant communication.

NOTE: The tenant network carries tenant network traffic over tagged VLANs. The interfaces connected to this network are not configured with IPv4 addresses by the OpenStack administrator. Instead, instances and services are allocated addresses within user-created subnets on tenant networks. Network namespaces prevent different users' subnets from conflicting with each other or with the infrastructure's own subnets.

- ⤴ All Swift storage communication occurs via a second 10Gb storage network switch on the 172.31.0.0/16 network. This network delivers the Object storage service communication and delivery.
- ⤴ The Service network carries service requests to the service listeners. These include the various schedulers and agents deployed in the OpenStack environment. The service traffic is segmented from the tenant and management traffic. The service network interfaces are assigned IP addresses in the 172.16.2.0/24 range.



NOTE: This reference architecture uses four physical networks. However it is possible to deploy supported OpenStack solutions with more or fewer networks.



Graphic 4.1.2.1: Network Topology

4.1.3 Network Addresses

Table 4.1.3.1 Host IP Addresses lists the IP addresses used in this reference architecture by server host name and role. All servers have interfaces on four networks: management, service, tenant, and storage. There are four roles: cloud controller, Neutron networker, Swift storage server, or compute node. A minimum of one cloud controller, one Neutron networker, three Swift storage servers, and one compute node are required to implement this reference architecture. In this example three compute nodes were used.



Host	Role	Network	Interface	Network Address
rhos0	Cloud Controller	Management	eno1	10.19.137.100
		Storage	enp20	172.31.139.100
		Service	eno2	172.16.2.100
rhos1	Neutron Networker	Management	eno1	10.19.137.101
		Storage	enp20	172.31.139.101
		Tenant	enp21	VLAN 1000:1010
		Service	eno2	172.16.2.101
rhos2	Compute Node	Management	eno1	10.19.137.102
		Storage	enp20	172.31.139.102
		Tenant	enp21	VLAN 1000:1010
		Service	eno2	172.16.2.102
rhos3	Compute Node	Management	eno1	10.19.137.103
		Storage	enp20	172.31.139.103
		Tenant	enp21	VLAN 1000:1010
		Service	eno2	172.16.2.103
rhos4	Compute Node	Management	eno1	10.19.137.104
		Storage	enp20	172.31.139.104
		Tenant	enp21	VLAN 1000:1010
		Service	eno2	172.16.2.104
rhos5	Compute Node	Management	eno1	10.19.137.105
		Storage	enp20	172.31.139.105
		Tenant	enp21	VLAN 1000:1010
		Service	eno2	172.16.2.105
rhos7	Compute Node	Management	em1	10.19.137.107
		Storage	p1p2	172.31.139.107
rhos8	Compute Node	Management	em1	10.19.137.108
		Storage	p1p2	172.31.139.108
rhos9	Compute Node	Management	em1	10.19.137.109
		Storage	p1p2	172.31.139.109

Table 4.1.3.1: Host IP Addresses



4.2 Software Configuration

This section of the reference architecture lists required software revisions and security details. Customers who use the correct OpenStack and Red Hat Enterprise Linux channels on Red Hat Network (RHN) or Subscription Manager meet the minimum required software versions.

4.2.1 Required Channels

Red Hat Enterprise Linux OpenStack Platform is available via the Red Hat Network channels and RHN Certificate Server repositories listed in **Table 4.2.1.1: Required Channels**.

Channel	Source
rhel-x86_64-server-7	RHN Classic
rhel-x86_64-server-7-ost-6	RHN Classic
rhel-7-server-rpms	RHN Certificate
rhel-7-server-openstack-6.0-rpms	RHN Certificate
rhel-7-server-rh-common-rpms	RHN Certificate

Table 4.2.1.1: Required Channels

NOTE: This reference architecture uses RHN Classic in all examples via a lab satellite server.

4.2.2 Software Versions

Table 4.2.2.1: Software Versions lists the software versions used to develop this reference architecture.

Host	Software	Version
Cloud Controller	openstack-cinder	2014.2.2-2
	openstack-dashboard	2014.2.2-2
	openstack-dashboard-theme	2014.2.2-2
	openstack-glance	2014.2.2-1
	openstack-heat-api	2014.2.2-1
	openstack-heat-common	2014.2.2-1
	openstack-heat-engine	2014.2.2-1
	openstack-keystone	2014.2.2-1
	openstack-neutron	2014.2.2-5
	openstack-nova-api	2014.2.2-5
	openstack-nova-cert	2014.2.2-5
	openstack-nova-common	2014.2.2-19
	openstack-nova-conductor	2014.2.2-19



	openstack-nova-console	2014.2.2-19
	openstack-nova-novncproxy	2014.2.2-19
	openstack-nova-scheduler	2014.2.2-19
	openstack-puppet-modules	2014.2.13-2
	openstack-sahara	2014.2.2-1
	openstack-sahara-doc	2014.2.2-1
	openstack-selinux	0.6.27-1
	openstack-swift	2.2.0-3
	openstack-swift-plugin-swift3	1.7-3
	openstack-swift-proxy	2.2.0-3
	openstack-utils	2014.2-1
	python-django-openstack-auth	1.1.7-4
	python-openstackclient	1.0.1-1
Compute Node	openstack-neutron	2014.2.2-5
	openstack-neutron-openvswitch	2014.2.2-5
	openstack-nova-common	2014.2.2-19
	openstack-nova-compute	2014.2.2-19
	openstack-selinux	0.6.27-1
	openstack-utils	2014.2-1
	python-neutron	2014.2.2-5
	python-neutronclient	2.3.9-1
Neutron Networker	openstack-neutron	2014.2.2-5
	openstack-neutron-ml2	2014.2.2-5
	openstack-neutron-openvswitch	2014.2.2-5
	openstack-selinux	0.6.27-1
	openstack-utils	2014.2-1
	python-neutron	2014.2.2-5
	python-neutronclient	2.3.9-1
Swift Storage Server	openstack-selinux	0.6.27-1
	openstack-swift	2.2.0-3
	openstack-swift-account	2.2.0-3
	openstack-swift-container	2.2.0-3
	openstack-swift-object	2.2.0-3



	openstack-utils	2014.2-1
--	-----------------	----------

Table 4.2.2.1: Software Versions

4.3 Security Details

This section describes the security configuration used in this reference architecture.

4.3.1 Firewall Configuration

Table 4.3.1.1 Allowed Ports by Role lists the allowed ports by host and role. Implement these via either `iptables` or `firewalld`.

Port	Host	Role
22, 80, 443, 873, 3260, 3306, 5000, 5671-5672, 6080, 8000, 8003-8004, 8080, 8386, 8773-8777, 9191, 9292, 9696, 9697, 11211, 35357	rhos0	Cloud Controller
22, 4789	rhos1	Neutron Networker
22, 53, 67, 4789, 5900:5999	rhos[2-5]	Compute node
22, 873, 6000-6002	rhos[7-9]	Storage node

Table 4.3.1.1: Allowed Ports by Role

4.3.2 SELinux Configuration

Red Hat Enterprise Linux OpenStack Platform supports **SELinux** in **enforcing** mode in Red Hat Enterprise Linux 7. **Table 4.3.2.1 Supported SELinux Package Versions** lists the required packages.

Package	Version
libselinux	2.2.2-6
selinux-policy	3.13.1-23
selinux-policy-targeted	3.13.1-23
openstack-selinux	0.6.27-1

Table 4.3.2.1: Supported SELinux Package Versions

4.4 Hardware Details

4.4.1 Server Hardware Configuration

Table 4.4.1.1 Server Hardware lists the hardware specifications for the servers used in this reference architecture.



NOTE: Red Hat Enterprise Linux OpenStack Platform servers do not require identical hardware. The hardware used in this reference architecture meets the minimum requirements outlined in the OpenStack documentation. The Red Hat Enterprise Linux OpenStack Platform 6 installation guide contains further details: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/6/html/Installer_and_Foreman_Guide/index.html

Hardware	Specifications
6 x Dell PowerEdge M520	2 x Intel Xeon CPU E5-2450 0 @ 2.10GHz
	2 x Broadcom 5720 1Gb Dual Port LOMs (4 x 1Gb ports) Broadcom 57810S-k Dual Port 10Gb NIC
	6 x DDR3 8192 MB @1333 MHZ DIMMs
	2 x 146GB SAS internal disk drives
3 x Dell PowerEdge R520	2 x Intel(R) Xeon(R) CPU X5650 @ 2.67 GHz (6 core)
	2 x Broadcom NetXtreme II BCM5709S Gb Ethernet 2 x Emulex Corporation OneConnect 10Gb NIC
	6 x DDR3 8192 MB @1333 MHZ DIMMs
	12 x 146GB SAS internal disk drives

Table 4.4.1.1: Server Hardware

4.4.2 Network Hardware

Each server has up to four interfaces: a management network interface, a private interface for service communication, an interface with tagged VLANs for tenant traffic, and a storage interface.

The management network interfaces and network storage interfaces are connected via a Juniper EX4550 switch. The network storage interfaces are uplinked through a 10GB switch module to 10 GB ports on the Juniper switch.

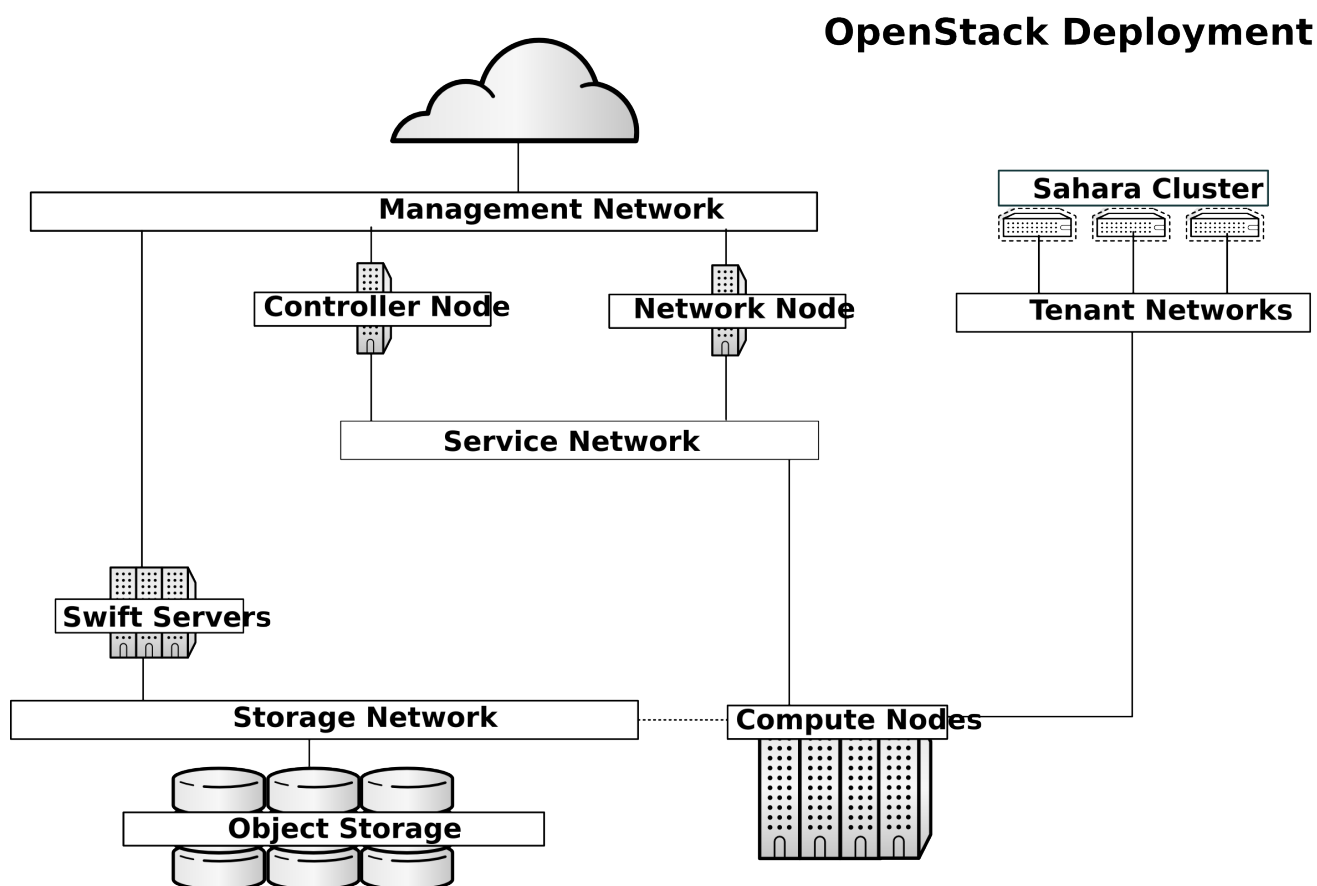


5. Implementing Sahara on OpenStack

This section describes the process that was followed to stand up Sahara on OpenStack in the Systems Engineering lab.

5.1 Prepare the Hosts

Complete the steps described in this section on all servers including the cloud controller, network server, compute nodes, and storage servers.



Graphic 5.1.1: Deployment Topology

5.2 Deploy OpenStack

RHEL OSP 6 supports the Puppet-based OpenStack Installer. This reference architecture demonstrates Sahara on a small but realistic non-HA OpenStack deployment depicted in **Graphic 5.1.1 Deployment Topology**.



The reference architecture OpenStack deployment can be summarized as follows:

- ⤴ All servers run Red Hat Enterprise Linux 7 deployed via a Red Hat Network Satellite server.
- ⤴ All servers run Red Hat Enterprise Linux OpenStack Platform 6, based on the Juno release configured via Puppet manifests.
- ⤴ Core OpenStack services configured in this reference architecture include: Keystone, Nova, Horizon, Heat, Cinder, and Glance.
- ⤴ Neutron networking is used in this reference architecture with ML2 and Open vSwitch plugins.
- ⤴ VLAN 1000:1010 are used for tenant networking.
- ⤴ Swift Object Storage is used as a Sahara data store. The cloud controller acts as the Swift proxy.
- ⤴ Each server has four network interfaces with dedicated networks for management, storage, service, and tenant traffic.

NOTE: Appendix A: References links to installation documentation for Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux OpenStack Platform 6.

5.3 Prepare the Environment

1. On the cloud controller, source the *keystonerc_admin* file.

```
[root@rhos0 ~]# source /root/keystonerc_admin

[root@rhos0 ~(keystone_admin)]# env | grep OS_
OS_REGION_NAME=RegionOne
OS_PASSWORD=54e6a402ae8d4e47
OS_AUTH_URL=http://10.19.137.100:5000/v2.0/
OS_USERNAME=admin
OS_TENANT_NAME=admin
```

2. Create a provider network in the *services* tenant named *ext_net* that maps to the physical lab network.

```
[root@rhos0 ~(openstack_admin)]# neutron net-create ext_net \
--router:external=True --provider:network_type flat \
--provider:physical_network physnet1
```

Created a new network:

Field	Value
admin_state_up	True
id	7792e417-7fc6-4baa-83cb-ae014675877b
name	ext_net



provider:network_type	flat
provider:physical_network	physnet1
provider:segmentation_id	
router:external	True
shared	False
status	ACTIVE
subnets	
tenant_id	e7061eccc4ab4a4c93c871dbbb9fe870

3. Create a subnet named **public** with floating IP addresses and a default gateway.

```
[root@rhos0 ~(openstack_admin)]# neutron subnet-create --name public \  
--gateway 172.16.2.100 --allocation-pool \  
start=172.16.2.112,end=172.16.2.120 ext_net 172.16.2.0/24 --enable_dhcp=False
```

Created a new subnet:

Field	Value
allocation_pools	{"start": "172.16.2.112", "end": "172.16.2.120"}
cidr	172.16.2.0/24
dns_nameservers	
enable_dhcp	False
gateway_ip	172.16.2.100
host_routes	
id	66b7fd45-06d5-46a1-8383-b48795b82e59
ip_version	4
ipv6_address_mode	
ipv6_ra_mode	
name	public
network_id	7792e417-7fc6-4baa-83cb-ae014675877b
tenant_id	e7061eccc4ab4a4c93c871dbbb9fe870

4. Create a tenant and tenant user.

```
[root@rhos0 ~(keystone_admin)]# keystone user-create --name refarch \  
--pass refarch
```

Property	Value
email	
enabled	True
id	20586cafabca4d5b91da08c463d63834
name	refarch
username	refarch

```
[root@rhos0 ~(keystone_admin)]# keystone tenant-create --name refarch-tenant
```

Property	Value
description	
enabled	True
id	24375feeb728444eabb8ce3e51a5170f



name	refarch-tenant

5. Add the ***_member_*** role to the tenant user.

```
[root@rhos0 ~(openstack_admin)]# keystone user-role-add --user-id refarch \
--tenant-id refarch-tenant --role-id _member_
```

```
[root@rhos0 ~(keystone_admin)]# keystone user-role-list --user-id refarch \
--tenant-id refarch-tenant
```

id	name	user_id
tenant_id		
9fe2ff9ee4384b1894a90878d3e92bab	_member_	20586cafabca4d5b91da08c463d63834 24375feeb728444eabb8ce3e51a5170f

6. Create a ***keystonerc*** file for the tenant user.

```
[root@rhos0 ~(openstack_admin)]# cat /root/keystonerc_refarch << EOF
export OS_USERNAME=refarch
export OS_TENANT_NAME=refarch-tenant
export OS_PASSWORD=refarch
export OS_AUTH_URL=http://10.19.137.100:35357/v2.0/
export PS1='\u@\h \W(refarch_member)]\$ '
EOF
```

7. Switch to ***refarch*** tenant user.

```
[root@rhos0 ~(openstack_admin)]# source /root/keystonerc_refarch
```

```
[root@rhos0 ~(refarch_member)]# env | grep OS_
OS_PASSWORD=refarch
OS_AUTH_URL=http://10.19.137.100:35357/v2.0/
OS_USERNAME=refarch
OS_TENANT_NAME=refarch-tenant
```

8. Create a tenant network named ***net1*** and a subnet named ***private***.

```
[root@rhos0 ~(refarch_member)]# neutron net-create net1
Created a new network:
```

Field	Value
admin_state_up	True
id	a718e96f-df8c-474b-ad34-1c5a01a8ed9d
name	net1
router:external	False
shared	False
status	ACTIVE
subnets	
tenant_id	24375feeb728444eabb8ce3e51a5170f



```
[root@rhos0 ~(refarch_member)]# neutron subnet-create --name private net1 \
172.16.3.0/24
```

Created a new subnet:

Field	Value
allocation_pools	{"start": "172.16.3.2", "end": "172.16.3.254"}
cidr	172.16.3.0/24
dns_nameservers	
enable_dhcp	True
gateway_ip	172.16.3.1
host_routes	
id	0265b876-5f39-4924-b889-b53cd99b8a7c
ip_version	4
ipv6_address_mode	
ipv6_ra_mode	
name	private
network_id	a718e96f-df8c-474b-ad34-1c5a01a8ed9d
tenant_id	24375feeb728444eabb8ce3e51a5170f

9. View the network and subnet.

```
[root@rhos0 ~(refarch_member)]# neutron net-list
```

id	name	subnets
a718e96f-df8c-474b-ad34-1c5a01a8ed9d	net1	0265b876-5f39-4924-b889-b53cd99b8a7c 172.16.3.0/24
7792e417-7fc6-4baa-83cb-ae014675877b	ext_net	66b7fd45-06d5-46a1-8383-b48795b82e59

```
[root@rhos0 ~(refarch_member)]# neutron subnet-list
```

id	name	cidr	allocation_pools
0265b876-5f39-4924-b889-b53cd99b8a7c	private	172.16.3.0/24	{"start": "172.16.3.2", "end": "172.16.3.254"}

10. Create a router named *route1*.

```
[root@rhos0 ~(refarch_member)]# neutron router-create route1
```

Created a new router:

Field	Value
admin_state_up	True
external_gateway_info	
id	24b8aba0-8701-47ad-902b-44d134dfabd7
name	route1
routes	
status	ACTIVE
tenant_id	24375feeb728444eabb8ce3e51a5170f



```

+-----+
[root@rhos0 ~(refarch_member)]# neutron router-show route1
+-----+
| Field                | Value
+-----+
| admin_state_up      | True
| external_gateway_info | {"network_id": "7792e417-7fc6-4baa-83cb-ae014675877b",
"enable_snat": true, "external_fixed_ips": [{"subnet_id": "66b7fd45-06d5-46a1-
8383-b48795b82e59", "ip_address": "172.16.2.112"}]} |
| id                   | 24b8aba0-8701-47ad-902b-44d134dfabd7
| name                 | route1
| routes              |
| status              | ACTIVE
| tenant_id           | 24375feeb728444eabb8ce3e51a5170f
+-----+

```

11. Attach the *private* subnet to the router.

```

[root@rhos0 ~(refarch_member)]# subnet_id=$(neutron subnet-list | \
awk ' /172.16.3.0/ { print $2 } ')

[root@rhos0 ~(refarch_member)]# echo $subnet_id
0265b876-5f39-4924-b889-b53cd99b8a7c

[root@rhos0 ~(refarch_member)]# neutron router-interface-add route1 $subnet_id
Added interface a534c38a-a506-4b97-a1e7-88ba61873420 to router route1.

```

12. Create security group rules to allow traffic to instances in this security group. **Table 4.3.1.1 Allowed Ports by Role** lists the ports required by OpenStack and Sahara.

```

[root@rhos0 ~(refarch_member)]# default_id=$(neutron security-group-list | \
awk ' /default/ { print $2 } ')

[root@rhos0 ~(refarch_member)]# echo -e "default id: $default_id"
default id: 68183225-5706-4765-99e7-1dc00cf7c220

[root@rhos0 ~(refarch_member)]# neutron security-group-rule-create -direction \
ingress --protocol icmp $default_id
Created a new security_group_rule:
+-----+
| Field                | Value
+-----+
| direction            | ingress
| ethertype            | IPv4
| id                   | 4dcc5bc6-ae36-4a63-bb2b-762c3744d360
| port_range_max       |
| port_range_min       |
| protocol             | icmp
| remote_group_id      |
| remote_ip_prefix     |
| security_group_id    | 68183225-5706-4765-99e7-1dc00cf7c220
| tenant_id           | 24375feeb728444eabb8ce3e51a5170f
+-----+

[root@rhos0 ~(refarch_member)]# neutron security-group-rule-create -direction \

```



```
ingress --protocol tcp --port_range_min 22 --port_range_max 22 $default_id  
Created a new security_group_rule:
```

Field	Value
direction	ingress
ethertype	IPv4
id	891a3f72-9a51-40b5-a3d6-4e3abe78dd70
port_range_max	22
port_range_min	22
protocol	tcp
remote_group_id	
remote_ip_prefix	
security_group_id	68183225-5706-4765-99e7-1dc00cf7c220
tenant_id	24375feeb728444eabb8ce3e51a5170f

```
[root@rhos0 ~(refarch_member)]# neutron security-group-rule-create -direction \\  
ingress --protocol tcp --port_range_min 80 --port_range_max 80 $default_id  
Created a new security_group_rule:
```

Field	Value
direction	ingress
ethertype	IPv4
id	a25c2da2-b2dd-461d-b9a8-95feffe3e918
port_range_max	80
port_range_min	80
protocol	tcp
remote_group_id	
remote_ip_prefix	
security_group_id	68183225-5706-4765-99e7-1dc00cf7c220
tenant_id	24375feeb728444eabb8ce3e51a5170f

```
[root@rhos0 ~(refarch_member)]# neutron security-group-rule-create -direction \\  
ingress --protocol tcp --port_range_min 873 --port_range_max 873 $default_id  
Created a new security_group_rule:
```

Field	Value
direction	ingress
ethertype	IPv4
id	55142d61-43e9-43c4-a7a0-0e30a3dd1fa6
port_range_max	873
port_range_min	873
protocol	tcp
remote_group_id	
remote_ip_prefix	
security_group_id	68183225-5706-4765-99e7-1dc00cf7c220
tenant_id	24375feeb728444eabb8ce3e51a5170f



NOTE: Command output is shown only for the first few ports and protocols. The remainder are truncated for brevity. Repeat the **neutron-securitygroup-rule-create** command for all remaining ports in **Table 4.3.1.1 Allowed Ports by Role**.

13. Display the new security group rules.

```
[root@rhos0 ~(refarch_member)]# neutron security-group-show --max-width 50 \
default
```

Field	Value
description	default
id	68183225-5706-4765-99e7-1dc00cf7c220
name	default
security_group_rules	{ "remote_group_id": null, "direction": "ingress", "remote_ip_prefix": null, "protocol": "tcp", "tenant_id": "24375feeb728444eabb8ce3e51a5170f", "port_range_max": 50075, "security_group_id": "68183225-5706-4765-99e7-1dc00cf7c220", "port_range_min": 50075, "ethertype": "IPv4", "id": "14ef8675-8011-4012-b251-ad3d8e0c8b72"} { "remote_group_id": null, "direction": "ingress", "remote_ip_prefix": null, "protocol": "tcp", "tenant_id": "24375feeb728444eabb8ce3e51a5170f", "port_range_max": 10020, "security_group_id": "68183225-5706-4765-99e7-1dc00cf7c220", "port_range_min": 10020, "ethertype": "IPv4", "id": "2506f9a8-e0b6-46f8-876e-67d335937222"} { "remote_group_id": null, "direction": "ingress", "remote_ip_prefix": null, "protocol": "tcp", "tenant_id": "24375feeb728444eabb8ce3e51a5170f", "port_range_max": 50060, "security_group_id": "68183225-5706-4765-99e7-1dc00cf7c220", "port_range_min": 50060, "ethertype": "IPv4", "id": "34eb83ed-7b11-4d50-8e0a-e9475ca1db65"} [Output truncated for brevity...] { "remote_group_id": null, "direction": "ingress", "remote_ip_prefix": null, "protocol": "tcp", "tenant_id": "24375feeb728444eabb8ce3e51a5170f", "port_range_max": 8088, "security_group_id": "68183225-5706-4765-99e7-1dc00cf7c220", "port_range_min": 8088, "ethertype": "IPv4", "id": "efb0c40b-0f71-44a6-9660-241a3cf55a93"} { "remote_group_id": null, "direction": "ingress", "remote_ip_prefix": null, "protocol": "tcp", "tenant_id": "24375feeb728444eabb8ce3e51a5170f", "port_range_max": 8033, "security_group_id": "68183225-5706-4765-99e7-1dc00cf7c220", "port_range_min": 8030, "ethertype": "IPv4", "id": "f768bc0a-ff2f-45d8-a089-815c882c1642"}
tenant_id	24375feeb728444eabb8ce3e51a5170f

14. Set the external network gateway on the router to the external network.



```
[root@rhos0 ~(refarch_member)]# route_id=$(neutron router-list | \
awk ' /route1/ { print $2 }')

[root@rhos0 ~(refarch_member)]# ext_net_id=$(neutron net-list | \
awk ' /ext_net/ { print $2 } ')

[root@rhos0 ~(refarch_member)]# echo -e "route_id: $route_id"
route_id: 24b8aba0-8701-47ad-902b-44d134dfabd7

[root@rhos0 ~(refarch_member)]# echo -e "ext_net_id: $ext_net_id"
ext_net_id: 7792e417-7fc6-4baa-83cb-ae014675877b

[root@rhos0 ~(refarch_member)]# neutron router-gateway-set $route_id $ext_net_id
Set gateway for router 24b8aba0-8701-47ad-902b-44d134dfabd7
```

15. Create a key pair.

```
[root@rhos0 ~(refarch_member)]# nova keypair-add refarchkp > /root/refarchkp.pem

[root@rhos0 ~(refarch_member)]# chmod 600 /root/refarchkp.pem
```

5.4 Install Sahara

1. Open port 8386 for Sahara API communication on the cloud controller.

NOTE: This example uses **iptables** due to the underlying Puppet installer. The equivalent **firewalld** command may also be used with Red Hat Enterprise Linux 7.

```
[root@rhos0 ~(keystone_admin)]# iptables -I INPUT 1 -p tcp --dport 8386 -j ACCEPT

[root@rhos0 ~(keystone_admin)]# iptables-save > /etc/sysconfig/iptables.save
```

2. Install the Sahara packages.

```
[root@rhos0 ~(keystone_admin)]# yum install -yq openstack-sahara \
python-saharaclient openstack-sahara-doc
```

3. Create the Sahara database. In this example the OpenStack state database is used for Sahara. It is possible to use an alternate database.

```
[root@rhos0 ~(keystone_admin)]# mysql -u root -e "CREATE DATABASE sahara;"

[root@rhos0 ~(keystone_admin)]# mysql -u root -e "GRANT ALL ON sahara.* TO \
'sahara'@'%' IDENTIFIED BY 'saharatest';"

[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
database connection mysql://sahara:saharatest@$IPADDR/sahara

[root@rhos0 ~(keystone_admin)]# sahara-db-manage -config-file \
/etc/sahara/sahara.conf upgrade head
INFO [alembic.migration] Context impl MySQLImpl.
```



```
INFO [alembic.migration] Will assume non-transactional DDL.
INFO [alembic.migration] Running upgrade None -> 001, Icehouse release
INFO [alembic.migration] Running upgrade 001 -> 002, placeholder
INFO [alembic.migration] Running upgrade 002 -> 003, placeholder
INFO [alembic.migration] Running upgrade 003 -> 004, placeholder
INFO [alembic.migration] Running upgrade 004 -> 005, placeholder
INFO [alembic.migration] Running upgrade 005 -> 006, placeholder
INFO [alembic.migration] Running upgrade 006 -> 007, convert
clusters.status_description to LongText
INFO [alembic.migration] Running upgrade 007 -> 008, add security_groups field
to node groups
INFO [alembic.migration] Running upgrade 008 -> 009, add rollback info to
cluster
INFO [alembic.migration] Running upgrade 009 -> 010, add auto_security_groups
flag to node group
INFO [alembic.migration] Running upgrade 010 -> 011, add Sahara settings info
to cluster
```

```
[root@rhos0 ~(keystone_admin)]# mysql -u sahara --password=saharatest \
sahara -e "show tables;"
```

```
+-----+
| Tables_in_sahara |
+-----+
| alembic_version |
| cluster_templates |
| clusters |
| data_sources |
| instances |
| job_binaries |
| job_binary_internal |
| job_executions |
| jobs |
| libs_association |
| mains_association |
| node_group_templates |
| node_groups |
| templates_relations |
+-----+
```

4. Source *keystonerc_admin* and create the Sahara user in the **services** tenant.

```
[root@rhos0 ~(keystone_admin)]# source /root/keystonerc_admin
```

```
[root@rhos0 ~(keystone_admin)]# keystone user-create --name sahara \
--pass saharatest
```

```
+-----+
| Property | Value |
+-----+
| email | |
| enabled | True |
| id | 6dfa4ac5f2194b7898e9802545e00f94 |
| name | sahara |
| username | sahara |
+-----+
```

```
[root@rhos0 ~(keystone_admin)]# keystone user-role-add --user sahara \
--role admin --tenant services
```



5. Create a Keystone service for Sahara of type *data_processing*.

```
[root@rhos0 ~(keystone_admin)]# keystone service-create --name sahara --type data_processing --description "Sahara data processing"
```

Property	Value
description	Sahara data processing
enabled	True
id	7fb2653febee4eb89c6f3c73a8ed9210
name	sahara
type	data_processing

6. Create a Keystone service endpoint for Sahara on port 8386.

```
[root@rhos0 ~(keystone_admin)]# keystone endpoint-create \
--service-id $(keystone service-list | awk '/ sahara / {print $2}') \
--publicurl http://10.19.137.100:8386/v1.1/%(tenant_id)s \
--internalurl http://10.19.137.100:8386/v1.1/%(tenant_id)s \
--adminurl http://10.19.137.100:8386/v1.1/%(tenant_id)s \
--region RegionOne
```

Property	Value
adminurl	http://10.19.137.100:8386/v1.1/%(tenant_id)s
id	b91177e5f22944c383963f0a35ef2ef2
internalurl	http://10.19.137.100:8386/v1.1/%(tenant_id)s
publicurl	http://10.19.137.100:8386/v1.1/%(tenant_id)s
region	RegionOne
service_id	7fb2653febee4eb89c6f3c73a8ed9210

7. Configure the Sahara service to use Keystone authentication and Neutron networking. Add the vanilla Hadoop plugin to the installed plugins list.

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
keystone_auth token admin_user sahara
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
keystone_auth token admin_tenant_name services
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
keystone_auth token admin_password saharatest
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
keystone_auth token auth_uri http://10.19.137.100:5000/v2.0/
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
keystone_auth token identity_uri http://10.19.137.100:35357/
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
DEFAULT plugins hdp,vanilla
```

```
[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
```



```

DEFAULT use_neutron true

[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
DEFAULT log_dir /var/log/sahara

[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
DEFAULT verbose true

[root@rhos0 ~(keystone_admin)]# openstack-config --set /etc/sahara/sahara.conf \
DEFAULT debug false

```

8. Start and enable Sahara on the cloud controller.

```

[root@rhos0 ~(keystone_admin)]# systemctl restart openstack-sahara-all

[root@rhos0 ~(keystone_admin)]# systemctl enable openstack-sahara-all
ln -s '/usr/lib/systemd/system/openstack-sahara-all.service'
'/etc/systemd/system/multi-user.target.wants/openstack-sahara-all.service'

[root@rhos0 ~(keystone_admin)]# systemctl status openstack-sahara-all
openstack-sahara-all.service - Sahara API Server
  Loaded: loaded (/usr/lib/systemd/system/openstack-sahara-all.service;
  enabled)
  Active: active (running) since Wed 2015-04-22 17:37:37 CDT; 249ms ago
  Main PID: 42511 (sahara-all)
  CGroup: /system.slice/openstack-sahara-all.service
          └─42511 /usr/bin/python2 /usr/bin/sahara-all --config-file
  /etc/sahara/sahara.conf
Apr 22 17:37:37 rhos0.cloud.lab.eng.bos.redhat.com systemd[1]: Starting Sahara
API Server...
Apr 22 17:37:37 rhos0.cloud.lab.eng.bos.redhat.com systemd[1]: Started Sahara
API Server.

```

9. List the installed and enabled Sahara plugins.

```

[root@rhos0 ~(keystone_admin)]# sahara plugin-list
+-----+-----+-----+
| name   | versions          | title                               |
+-----+-----+-----+
| hdp    | 2.0.6             | Hortonworks Data Platform         |
| vanilla | 1.2.1, 2.3.0, 2.4.1 | Vanilla Apache Hadoop             |
+-----+-----+-----+

```

5.5 Configure Sahara

1. Configure Sahara on the cloud controller as the tenant user.

```

[root@rhos0 ~(refarch_member)]# source /root/keystonerc_admin

[root@rhos0 ~(keystone_admin)]# env | grep OS_
OS_REGION_NAME=RegionOne
OS_PASSWORD=54e6a402ae8d4e47
OS_AUTH_URL=http://10.19.137.100:5000/v2.0/
OS_USERNAME=admin
OS_TENANT_NAME=admin

```



2. Export Hadoop environment variables.

```
[root@rhos0 ~(keystone_admin)]# sahara plugin-list
+-----+-----+-----+
| name      | versions                | title                               |
+-----+-----+-----+
| hdp       | 2.0.6                   | Hortonworks Data Platform         |
| vanilla   | 1.2.1, 2.3.0, 2.4.1    | Vanilla Apache Hadoop             |
+-----+-----+-----+

[root@rhos0 ~(keystone_admin)]# export PLUGIN_NAME=vanilla

[root@rhos0 ~(keystone_admin)]# export PLUGIN_VERSION=2.4.1

[root@rhos0 ~(keystone_admin)]# export MASTER_FLAVOR=m1.large

[root@rhos0 ~(keystone_admin)]# export WORKER_FLAVOR=m1.large

[root@rhos0 ~(keystone_admin)]# export WORKER_COUNT=3
```

3. Export OpenStack environment variables.

```
[root@rhos0 ~(keystone_admin)]# export MANAGEMENT_NETWORK=net1

[root@rhos0 ~(keystone_admin)]# export EXT_NET_ID=$(nova net-list | \
awk ' /ext_net/ { print $2 } ')

[root@rhos0 ~(keystone_admin)]# echo $EXT_NET_ID
7792e417-7fc6-4baa-83cb-ae014675877b

[root@rhos0 ~(keystone_admin)]# export KEYPAIR=refarchkp
```

5.6 Import Glance Images for Hadoop

1. Create or download a Glance image for Hadoop. The image used in this example was created with Sahara Image Elements².

NOTE: Links to community-maintained images can be found at https://sahara.readthedocs.org/en/stable-juno/userdoc/vanilla_plugin.html

2. Create the Glance images for Hadoop.

```
[root@rhos0 ~(keystone_admin)]# glance image-create --progress -name \
centos64_hadoop2 --is-public true --disk-format qcow2 \
--container-format bare \
--file /pub/rhos/images/sahara-juno-vanilla-2.4.1-centos-6.5.qcow2

[=====>] 100%
+-----+-----+-----+
| Property          | Value                               |
+-----+-----+-----+
| checksum          | 872c5dd330fb18a070afdbba8811ae25  |
+-----+-----+-----+
```

² <http://docs.openstack.org/developer/sahara/userdoc/diskimagebuilder.html>



```
| container_format | bare |
| created_at      | 2015-04-22T22:47:42 |
| deleted         | False |
| deleted_at      | None |
| disk_format     | qcow2 |
| id              | ecadd42f-61bc-4056-b8fd-64da8c55d336 |
| is_public       | True |
| min_disk        | 0 |
| min_ram         | 0 |
| name            | centos64_hadoop2 |
| owner           | e7061eccc4ab4a4c93c871dbbb9fe870 |
| protected       | False |
| size            | 1105592320 |
| status          | active |
| updated_at      | 2015-04-22T22:47:55 |
| virtual_size    | None |
+-----+-----+
```

3. Ensure the Glance image is **Active**.

```
[root@rhos0 ~(keystone_admin)]# glance image-list
+-----+-----+-----+-----+
| ID | Name | Disk Format |
+-----+-----+-----+-----+
| ecadd42f-61bc-4056-b8fd-64da8c55d336 | centos64_hadoop2 | qcow2 | bare |
| 1105592320 | active |
+-----+-----+-----+-----+
```

4. Register the image with Sahara.

```
[root@rhos0 ~(keystone_admin)]# export IMAGE_ID=$(glance image-list | \
awk ' /centos64_hadoop2/ { print $ 2 } ')

[root@rhos0 ~(keystone_admin)]# echo -e "Image ID: $IMAGE_ID"
Image ID: ecadd42f-61bc-4056-b8fd-64da8c55d336

[root@rhos0 ~(keystone_admin)]# sahara image-register --id $IMAGE_ID \
--username cloud-user

[root@rhos0 ~(keystone_admin)]# sahara image-add-tag --id $IMAGE_ID \
--tag $PLUGIN_NAME

[root@rhos0 ~(keystone_admin)]# sahara image-add-tag --id $IMAGE_ID \
--tag $PLUGIN_VERSION

[root@rhos0 ~(keystone_admin)]# sahara image-list
+-----+-----+-----+-----+
| name | id | username | tags |
| description |
+-----+-----+-----+-----+
| centos64_hadoop2 | ecadd42f-61bc-4056-b8fd-64da8c55d336 | cloud-user |
vanilla, 2.4.1 | None |
+-----+-----+-----+-----+
```



5.7 Define Templates

1. Node and cluster templates are reusable objects that define clusters. Define environment variables common to the templates.

```
[root@rhos0 ~(keystone_admin)]# source /root/keystonerc_refarch

[root@rhos0 ~(refarch_member)]# env | grep OS_
OS_REGION_NAME=RegionOne
OS_PASSWORD=refarch
OS_AUTH_URL=http://10.19.137.100:35357/v2.0/
OS_USERNAME=refarch
OS_TENANT_NAME=refarch-tenant

[root@rhos0 ~(refarch_member)]# export MASTER_FLAVOR_ID=$(nova \
flavor-show $MASTER_FLAVOR | awk ' / id / {print $4}')

[root@rhos0 ~(refarch_member)]# export WORKER_FLAVOR_ID=$(nova \
flavor-show $WORKER_FLAVOR | awk ' / id / {print $4}')

[root@rhos0 ~(refarch_member)]# export MANAGEMENT_NETWORK_ID=$(neutron \
net-show $MANAGEMENT_NETWORK | awk ' / id / {print $4}')
```

5.7.1 Define Master Group Template

1. Create a JSON file named `/root/master_group_template.json` that defines the master node of the cluster.

```
[root@rhos0 ~(refarch_member)]# tee /root/master_group_template.json << EOF
{
  "plugin_name": "$PLUGIN_NAME",
  "node_processes": ["namenode", "resourcemanager", "oozie", "historyserver"],
  "flavor_id": "$MASTER_FLAVOR_ID",
  "floating_ip_pool": "$EXT_NET_ID",
  "hadoop_version": "$PLUGIN_VERSION",
  "name": "master"
}
EOF
{
  "plugin_name": "vanilla",
  "node_processes": ["namenode", "resourcemanager", "oozie", "historyserver"],
  "flavor_id": "4",
  "floating_ip_pool": "7792e417-7fc6-4baa-83cb-ae014675877b",
  "hadoop_version": "2.4.1",
  "name": "master"
}

[root@rhos0 ~(refarch_member)]# sahara node-group-template-create \
--json /root/master_group_template.json | awk '/id/ { print $4 } '
```

2. Create the *master* group template.

```
[root@rhos0 ~(refarch_member)]# sahara node-group-template-create \
--json /root/master_group_template.json | awk '/id/ { print $4 } '
0026bb51-b8e9-4e53-84a0-cc62b5018d5e

[root@rhos0 ~(refarch_member)]# sahara node-group-template-show \
```




```
--id 0026bb51-b8e9-4e53-84a0-cc62b5018d5e
+-----+
| Property          | Value                                     |
+-----+-----+
| description       | None                                     |
| volume_mount_prefix | /volumes/disk                           |
| updated_at       | None                                     |
| plugin_name      | vanilla                                  |
| floating_ip_pool  | 7792e417-7fc6-4baa-83cb-ae014675877b   |
| image_id         | None                                     |
| volumes_size     | 0                                         |
| volumes_per_node | 0                                         |
| hadoop_version   | 2.4.1                                    |
| id               | 0026bb51-b8e9-4e53-84a0-cc62b5018d5e   |
| security_groups  | None                                     |
| name             | master                                   |
| tenant_id       | 24375feeb728444eabb8ce3e51a5170f      |
| created_at      | 2015-04-22 22:50:53                     |
| node_processes   | namenode, resourcemanager, oozie, historyserver |
| flavor_id       | 4                                         |
| node_configs    | {}                                       |
| auto_security_group | False                                    |
+-----+-----+
```

3. Set the **MASTER_TEMPLATE_ID** environment variable.

```
[root@rhos0 ~(refarch_member)]# export MASTER_TEMPLATE_ID=$(sahara node-group-
template-show --name master | awk '/\| id/ { print $4 } ')
```

5.7.2 Define the Worker Group Template

1. Create a JSON file named `/root/worker_group_template.json` that defines the worker nodes of the cluster.

```
[root@rhos0 ~(refarch_member)]# tee /root/worker_group_template.json << EOF
{
"plugin_name": "$PLUGIN_NAME",
"node_processes": ["datanode", "nodemanager"],
"flavor_id": "$WORKER_FLAVOR_ID",
"floating_ip_pool": "$EXT_NET_ID",
"hadoop_version": "$PLUGIN_VERSION",
"name": "worker"
}
EOF
{
"plugin_name": "vanilla",
"node_processes": ["datanode", "nodemanager"],
"flavor_id": "4",
"floating_ip_pool": "7792e417-7fc6-4baa-83cb-ae014675877b",
"hadoop_version": "2.4.1",
"name": "worker"
}
```

2. Create the **worker** group template.

```
[root@rhos0 ~(refarch_member)]# sahara node-group-template-create \
--json /root/worker_group_template.json | awk '/id/ { print $4 } '
```



```
c28c5fad-394d-421c-9e06-121589d4475c
```

```
[root@rhos0 ~(refarch_member)]# sahara node-group-template-show \
--id c28c5fad-394d-421c-9e06-121589d4475c
```

Property	Value
description	None
volume_mount_prefix	/volumes/disk
updated_at	None
plugin_name	vanilla
floating_ip_pool	7792e417-7fc6-4baa-83cb-ae014675877b
image_id	None
volumes_size	0
volumes_per_node	0
hadoop_version	2.4.1
id	c28c5fad-394d-421c-9e06-121589d4475c
security_groups	None
name	worker
tenant_id	24375feeb728444eabb8ce3e51a5170f
created_at	2015-04-22 22:51:35
node_processes	datanode, nodemanager
flavor_id	4
node_configs	{}
auto_security_group	False

3. Set the **WORKER_TEMPLATE_ID** environment variable.

```
[root@rhos0 ~(refarch_member)]# export WORKER_TEMPLATE_ID=$(sahara node-group-
template-show --name worker | awk '/\| id/ { print $4 }')
```

4. List the node templates.

```
[root@rhos0 ~(refarch_member)]# sahara node-group-template-list
```

name	id	plugin_name	node_processes
master	0026bb51-b8e9-4e53-84a0-cc62b5018d5e	vanilla	namenode,
			resourcemanager, oozie, historyserver
worker	c28c5fad-394d-421c-9e06-121589d4475c	vanilla	datanode,
			nodemanager

5.7.3 Define Cluster Templates

1. Create a JSON file named `/root/cluster_template.json` that defines a cluster of one master node and three worker nodes using the node group templates.

```
[root@rhos0 ~(refarch_member)]# tee /root/cluster_template.json << EOF
{
"plugin_name": "$PLUGIN_NAME",
"node_groups": [
  { "count": 1,
    "name": "master",
    "node_group_template_id": "$MASTER_TEMPLATE_ID" },
```



```

        { "count": $WORKER_COUNT,
          "name": "worker",
          "node_group_template_id": "$WORKER_TEMPLATE_ID" }],
"hadoop_version": "$PLUGIN_VERSION",
"name": "cluster"
}
EOF
{
"plugin_name": "vanilla",
"node_groups": [
  { "count": 1,
    "name": "master",
    "node_group_template_id": "0026bb51-b8e9-4e53-84a0-cc62b5018d5e" },
  { "count": 3,
    "name": "worker",
    "node_group_template_id": "c28c5fad-394d-421c-9e06-121589d4475c" }],
"hadoop_version": "2.4.1",
"name": "cluster"
}

```

2. Create a cluster template from the JSON definition.

```
[root@rhos0 ~(refarch_member)]# sahara cluster-template-create --json /root/cluster_template.json
```

Property	Value
neutron_management_network	None
description	None
cluster_configs	{}
created_at	2015-04-22 22:51:45
default_image_id	None
updated_at	None
plugin_name	vanilla
anti_affinity	[]
tenant_id	24375feeb728444eabb8ce3e51a5170f
node_groups	master: 1, worker: 3
hadoop_version	2.4.1
id	d83b8f39-e93f-41cf-9cea-cbe8866cc355
name	cluster

3. List the available cluster templates.

```
[root@rhos0 ~(refarch_member)]# sahara cluster-template-list
```

name	id	plugin_name	node_groups
cluster	d83b8f39-e93f-41cf-9cea-cbe8866cc355	vanilla	master: 1, worker: 3
	None		

4. Set the **CLUSTER_TEMPLATE_ID** environment variable.



```
[root@rhos0 ~(refarch_member)]# export CLUSTER_TEMPLATE_ID=$(sahara cluster-
template-show --name cluster | awk '/\| id/ { print $4 } ')
```

5.8 Launch a Cluster via the Command Line

1. Create a cluster definition file named `/root/cluster.json`.

```
[root@rhos0 ~(refarch_member)]# tee /root/cluster.json << EOF
{
"cluster_template_id": "$CLUSTER_TEMPLATE_ID",
"default_image_id": "$IMAGE_ID",
"hadoop_version": "$PLUGIN_VERSION",
"name": "cluster-instance-$(date +%s)",
"plugin_name": "$PLUGIN_NAME",
"user_keypair_id": "$KEYPAIR",
"neutron_management_network": "$MANAGEMENT_NETWORK_ID"
}
EOF
{
"cluster_template_id": "d83b8f39-e93f-41cf-9cea-cbe8866cc355",
"default_image_id": "ecadd42f-61bc-4056-b8fd-64da8c55d336",
"hadoop_version": "2.4.1",
"name": "cluster-instance-1429743107",
"plugin_name": "vanilla",
"user_keypair_id": "refarchkp",
"neutron_management_network": "a718e96f-df8c-474b-ad34-1c5a01a8ed9d"
}
```

2. Launch a cluster from the template via the Sahara command line.

```
[root@rhos0 ~(refarch_member)]# sahara cluster-create --json /root/cluster.json
```

Property	Value
status	Validating
neutron_management_network	a718e96f-df8c-474b-ad34-1c5a01a8ed9d
is_transient	False
description	None
user_keypair_id	refarchkp
updated_at	2015-04-22 22:51:56.804344
plugin_name	vanilla
anti_affinity	[]
node_groups	[{u'count': 3, u'name': u'worker', u'instances': [], u'volume_mount_prefix': u'/volumes/disk', u'created_at': u'2015-04-22 22:51:56', u'updated_at': None, u'floating_ip_pool': u'7792e417-7fc6-4baa-83cb-ae014675877b', u'image_id': None, u'volumes_size': 0, u'node_configs': {}, u'node_group_template_id': u'c28c5fad-394d-421c-9e06-121589d4475c', u'volumes_per_node': 0, u'node_processes': [u'datanode', u'nodemanager'], u'auto_security_group': False, u'security_groups': None, u'flavor_id': u'4'}, {u'count': 1, u'name': u'master', u'instances': [], u'volume_mount_prefix': u'/volumes/disk', u'created_at': u'2015-04-22 22:51:56',



```

management_public_key | u'updated_at': None, u'floating_ip_pool':
                        | u'7792e417-7fc6-4baa-83cb-ae014675877b',
                        | u'image_id': None, u'volumes_size': 0,
                        | u'node_configs': {}, u'node_group_template_id':
                        | u'0026bb51-b8e9-4e53-84a0-cc62b5018d5e',
                        | u'volumes_per_node': 0, u'node_processes':
                        | [u'namenode', u'resourcemanager', u'oozie',
                        | u'historyserver'], u'auto_security_group':
                        | False, u'security_groups': None, u'flavor_id':
                        | u'4']}]
                        | ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDDFW+tTtw
                        | TYbXBunae6mnebkbdthJQ2XlXQkwwAKmfvz6Bp3gtdNfh02
                        | dUxosgZyCNej4BthW+jJa0++qivj/adAWgaQQ0Q8CSJ6Ymr
                        | eMiHINCmTwiYFgZiIm7xiyCZfbb5UvdHvknqrBb5VTQJi22
                        | 9qZfOR5mJB9W6qzDD1zRXvoaEs0cRM3NxqeJaOdQWdjiihC
                        | Q40ghXfatQE4M/FgTbebJm432XmgcXzZ48n4n05U3i+GKX
                        | xC2RcWMacFbSqDG5SSUE+2rS/6YntUt0j1aLHmI150aR/FW
                        | 2JXelW0hfSw+nBipw+1vUEfdN0Ujq6Mnc9treYT0cRNgmDI
                        | oXbz Generated by Sahara

status_description
hadoop_version        | 2.4.1
id                    | d38d9dad-867d-4e88-8b51-ac193fe42cf1
trust_id              | None
info                  | {}
cluster_template_id  | d83b8f39-e93f-41cf-9cea-cbe8866cc355
name                  | cluster-instance-1429743107
cluster_configs       | {}
created_at            | 2015-04-22 22:51:56
default_image_id      | ecadd42f-61bc-4056-b8fd-64da8c55d336
tenant_id             | 24375feeb728444eabb8ce3e51a5170f

```

3. On a successful create, the cluster status will go from **Spawning** to **Waiting** to **Active**. Verify the cluster is in an **Active** state.

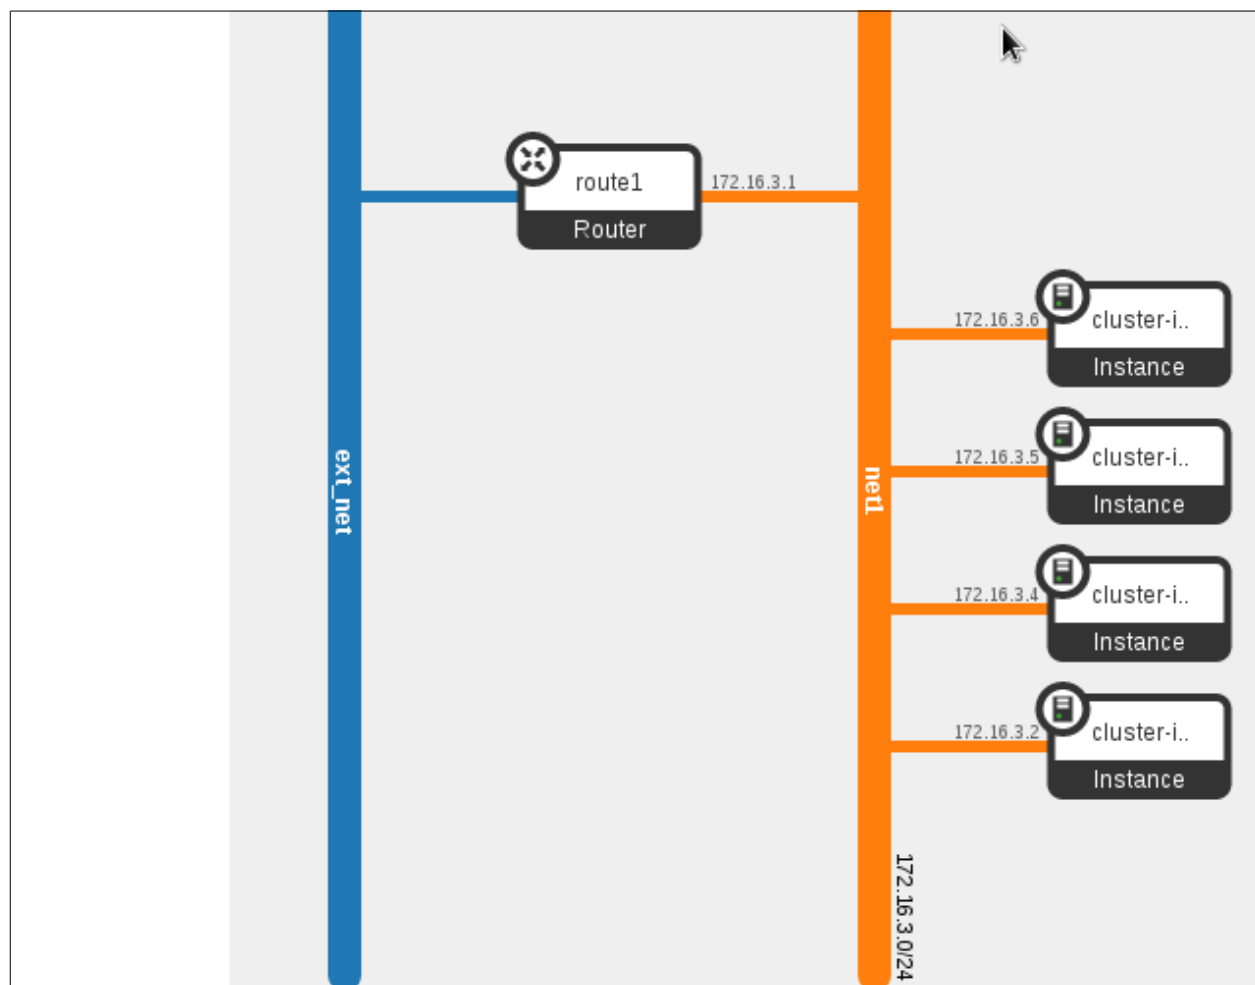
```

[root@rhos0 ~(refarch_member)]# sahara cluster-list
+-----+-----+-----+
| name          | id          | status |
+-----+-----+-----+
| cluster-instance-1429743107 | d38d9dad-867d-4e88-8b51-ac193fe42cf1 | Active |
+-----+-----+-----+

```



4. After the cluster is deployed successfully, the network and cluster layout should match the topology outlined in **Graphic 5.8.1: Sahara Cluster Topology**.



Graphic 5.8.1: Sahara Cluster Topology



6. Validate and Test Sahara

This section describes how to explore and validate the cluster installation by connecting to the master node and running a test MapReduce job.

6.1 Explore the Cluster

1. Find the floating IP address of the master node. In this example it is 10.19.137.115.

```
[root@rhos0 update(refarch_member)]# nova list
+-----+-----+-----+-----+-----+
| ID                | Name                                     |
| Status | Task State | Power State | Networks |
+-----+-----+-----+-----+
| 2350a4d3-57db-4290-8275-ec137d284e3e | cluster-instance-1429743107-master-001 | | |
| ACTIVE | -          | Running    | net1=172.16.3.6, 10.19.137.115 |
| bd5a4b30-3d61-484e-a029-54612272293b | cluster-instance-1429743107-worker-001 |
| ACTIVE | -          | Running    | net1=172.16.3.2, 10.19.137.116 |
| f42543c6-0549-4af8-97fd-a54aa81543b1 | cluster-instance-1429743107-worker-002 |
| ACTIVE | -          | Running    | net1=172.16.3.4, 10.19.1137.117 |
| 89da0462-6183-4f77-8070-4562b365b01f | cluster-instance-1429743107-worker-003 |
| ACTIVE | -          | Running    | net1=172.16.3.5, 10.19.137.118 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

2. SSH to the master node by floating IP address. By default, the image allows SSH login from the cloud-user account with authentication via the key pair specified in section 5.8
Launch a Cluster via the Command Line.

```
[root@rhos0 ~(refarch_member)]# ssh -l cloud-user -i /root/refarchkp.pem
10.19.137.115
```

3. After login, switch to the root user.

```
[cloud-user@cluster-instance-1429743107-master-001 ~]$ sudo su
```

4. View the running Java processes on the master node. In this example the master is running **NameNode**, **ResourceManager**, and the **JobHistoryServer**.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# jps
1836 ResourceManager
3200 Bootstrap
10875 Jps
2196 JobHistoryServer
1723 NameNode
```

5. View the environment variables related to Hadoop.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# env | grep -i hadoop
PATH=/sbin:/bin:/usr/sbin:/usr/bin:/usr/local/bin:/usr/java/jdk1.7.0_51/bin:/opt/hadoop/bin:/opt/hadoop/sbin:/usr/java/jdk1.7.0_51/bin
HADOOP_HDFS_HOME=/opt/hadoop
HADOOP_COMMON_HOME=/opt/hadoop
HADOOP_YARN_HOME=/opt/hadoop
```



```
HADOOP_MAPRED_HOME=/opt/hadoop
```

6. View the unique HDFS version ID, cluster ID, and block pool ID.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# grep ID \  
/mnt/hdfs/namenode/current/VERSION  
namespaceID=1075783263  
clusterID=CID-28df7086-a588-462b-89e1-b497631a52c5  
blockpoolID=BP-386020788-172.16.3.6-1429743928531
```

7. View the **NameNode** and **secondaryNameNode** configuration. In this example both run on the master node.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# su hadoop -c "hdfs  
getconf -namenodes"  
cluster-instance-1429743107-master-001
```

```
[root@cluster-instance-1429743107-master-001 cloud-user]# su hadoop -c "hdfs  
getconf -secondaryNameNodes"  
0.0.0.0
```

8. Run a report on the HDFS with **dfsadmin**.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# su - hadoop \  
-c "hdfs dfsadmin -report"  
Configured Capacity: 242043899904 (225.42 GB)  
Present Capacity: 222234882048 (206.97 GB)  
DFS Remaining: 222006140928 (206.76 GB)  
DFS Used: 228741120 (218.14 MB)  
DFS Used%: 0.10%  
Under replicated blocks: 0  
Blocks with corrupt replicas: 0  
Missing blocks: 0  
-----  
Datanodes available: 3 (3 total, 0 dead)  
  
Live datanodes:  
Name: 172.16.3.4:50010 (cluster-instance-1429743107-worker-002.novalocal)  
Hostname: cluster-instance-1429743107-worker-002.novalocal  
Decommission Status : Normal  
Configured Capacity: 80681299968 (75.14 GB)  
DFS Used: 76247040 (72.71 MB)  
Non DFS Used: 6603010048 (6.15 GB)  
DFS Remaining: 74002042880 (68.92 GB)  
DFS Used%: 0.09%  
DFS Remaining%: 91.72%  
Configured Cache Capacity: 0 (0 B)  
Cache Used: 0 (0 B)  
Cache Remaining: 0 (0 B)  
Cache Used%: 100.00%  
Cache Remaining%: 0.00%  
Last contact: Thu Apr 23 06:13:33 MSK 2015  
  
Name: 172.16.3.2:50010 (cluster-instance-1429743107-worker-001.novalocal)  
Hostname: cluster-instance-1429743107-worker-001.novalocal  
Decommission Status : Normal
```




```

Configured Capacity: 80681299968 (75.14 GB)
DFS Used: 76247040 (72.71 MB)
Non DFS Used: 6603001856 (6.15 GB)
DFS Remaining: 74002051072 (68.92 GB)
DFS Used%: 0.09%
DFS Remaining%: 91.72%
Configured Cache Capacity: 0 (0 B)
Cache Used: 0 (0 B)
Cache Remaining: 0 (0 B)
Cache Used%: 100.00%
Cache Remaining%: 0.00%
Last contact: Thu Apr 23 06:13:33 MSK 2015

```

```

Name: 172.16.3.5:50010 (cluster-instance-1429743107-worker-003.novalocal)
Hostname: cluster-instance-1429743107-worker-003.novalocal
Decommission Status : Normal
Configured Capacity: 80681299968 (75.14 GB)
DFS Used: 76247040 (72.71 MB)
Non DFS Used: 6603005952 (6.15 GB)
DFS Remaining: 74002046976 (68.92 GB)
DFS Used%: 0.09%
DFS Remaining%: 91.72%
Configured Cache Capacity: 0 (0 B)
Cache Used: 0 (0 B)
Cache Remaining: 0 (0 B)
Cache Used%: 100.00%
Cache Remaining%: 0.00%
Last contact: Thu Apr 23 06:13:33 MSK 2015

```

9. List the nodes seen by Yarn.

```

[root@cluster-instance-1429743107-master-001 cloud-user]# su - hadoop \
-c "yarn node --list"
Total Nodes:3
      Node-Id          Node-State   Node-Http-Address  Number-of-Running-
Containers
cluster-instance-1429743107-worker-002.novalocal:43790          RUNNING
      cluster-instance-1429743107-worker-002.novalocal:8042
0
cluster-instance-1429743107-worker-001.novalocal:48393          RUNNING
      cluster-instance-1429743107-worker-001.novalocal:8042
0
cluster-instance-1429743107-worker-003.novalocal:41236          RUNNING
      cluster-instance-1429743107-worker-003.novalocal:8042
0

```

10. Run a file system check.

```

[root@cluster-instance-1429743107-master-001 cloud-user]# su - hadoop \
-c "hdfs fsck /"
Connecting to namenode via http://cluster-instance-1429743107-master-001:50070
FSCK started by hadoop (auth:SIMPLE) from /172.16.3.6 for path / at Thu Apr 23
06:14:05 MSK 2015
.....Status: HEALTHY
Total size:      74686911 B
Total dirs:      18
Total files:     139

```



```
Total symlinks:          0
Total blocks (validated): 139 (avg. block size 537315 B)
Minimally replicated blocks: 139 (100.0 %)
Over-replicated blocks:  0 (0.0 %)
Under-replicated blocks:  0 (0.0 %)
Mis-replicated blocks:    0 (0.0 %)
Default replication factor: 3
Average block replication: 3.0
Corrupt blocks:          0
Missing replicas:        0 (0.0 %)
Number of data-nodes:    3
Number of racks:         1
FSCK ended at Thu Apr 23 06:14:05 MSK 2015 in 40 milliseconds

The filesystem under path '/' is HEALTHY
```

6.2 Run a Test Hadoop Job from the Command Line

1. Switch to the Hadoop user.

```
[root@cluster-instance-1429743107-master-001 cloud-user]# su - hadoop
```

2. Execute the *Pi* MapReduce example to verify Hadoop functionality.

```
[hadoop@cluster-instance-1429743107-master-001 ~]$ cd \
/opt/hadoop-2.4.1/share/hadoop/mapreduce/

[hadoop@cluster-instance-1429743107-master-001 mapreduce]$ hadoop jar \
hadoop-mapreduce-examples-2.4.1.jar pi 10 100
Number of Maps = 10
Samples per Map = 100
Wrote input for Map #0
Wrote input for Map #1
Wrote input for Map #2
Wrote input for Map #3
Wrote input for Map #4
Wrote input for Map #5
Wrote input for Map #6
Wrote input for Map #7
Wrote input for Map #8
Wrote input for Map #9
Starting Job
15/04/23 06:15:23 INFO client.RMProxy: Connecting to ResourceManager at cluster-
instance-1429743107-master-001/172.16.3.6:8032
15/04/23 06:15:24 INFO input.FileInputFormat: Total input paths to process : 10
15/04/23 06:15:24 INFO mapreduce.JobSubmitter: number of splits:10
15/04/23 06:15:25 INFO mapreduce.JobSubmitter: Submitting tokens for job:
job_1429743940375_0001
15/04/23 06:15:26 INFO impl.YarnClientImpl: Submitted application
application_1429743940375_0001
15/04/23 06:15:26 INFO mapreduce.Job: The url to track the job: http://cluster-
instance-1429743107-master-001:8088/proxy/application_1429743940375_0001/
15/04/23 06:15:26 INFO mapreduce.Job: Running job: job_1429743940375_0001
15/04/23 06:15:34 INFO mapreduce.Job: Job job_1429743940375_0001 running in uber
mode : false
15/04/23 06:15:34 INFO mapreduce.Job: map 0% reduce 0%
15/04/23 06:15:42 INFO mapreduce.Job: map 20% reduce 0%
```



```
15/04/23 06:15:50 INFO mapreduce.Job: map 100% reduce 0%
15/04/23 06:15:52 INFO mapreduce.Job: map 100% reduce 100%
15/04/23 06:15:52 INFO mapreduce.Job: Job job_1429743940375_0001 completed
successfully
15/04/23 06:15:52 INFO mapreduce.Job: Counters: 49
  File System Counters
    FILE: Number of bytes read=226
    FILE: Number of bytes written=1050159
    FILE: Number of read operations=0
    FILE: Number of large read operations=0
    FILE: Number of write operations=0
    HDFS: Number of bytes read=2940
    HDFS: Number of bytes written=215
    HDFS: Number of read operations=43
    HDFS: Number of large read operations=0
    HDFS: Number of write operations=3
  Job Counters
    Launched map tasks=10
    Launched reduce tasks=1
    Data-local map tasks=10
    Total time spent by all maps in occupied slots (ms)=111237
    Total time spent by all reduces in occupied slots (ms)=6098
    Total time spent by all map tasks (ms)=111237
    Total time spent by all reduce tasks (ms)=6098
    Total vcore-seconds taken by all map tasks=111237
    Total vcore-seconds taken by all reduce tasks=6098
    Total megabyte-seconds taken by all map tasks=113906688
    Total megabyte-seconds taken by all reduce tasks=6244352
  Map-Reduce Framework
    Map input records=10
    Map output records=20
    Map output bytes=180
    Map output materialized bytes=280
    Input split bytes=1760
    Combine input records=0
    Combine output records=0
    Reduce input groups=2
    Reduce shuffle bytes=280
    Reduce input records=20
    Reduce output records=0
    Spilled Records=40
    Shuffled Maps =10
    Failed Shuffles=0
    Merged Map outputs=10
    GC time elapsed (ms)=934
    CPU time spent (ms)=6080
    Physical memory (bytes) snapshot=2622341120
    Virtual memory (bytes) snapshot=9670066176
    Total committed heap usage (bytes)=2118123520
  Shuffle Errors
    BAD_ID=0
    CONNECTION=0
    IO_ERROR=0
    WRONG_LENGTH=0
    WRONG_MAP=0
    WRONG_REDUCE=0
  File Input Format Counters
```



```

Bytes Read=1180
File Output Format Counters
Bytes Written=97
Job Finished in 28.965 seconds
Estimated value of Pi is 3.14800000000000000000000000000000

```

6.3 Run a Test EDP Job Through Sahara

In the previous section a user submitted a Hadoop job via the command line. Users may also define and submit reusable jobs directly via Sahara. This section demonstrates how to define a job template and submit a Pig job that uses Swift as a data store.

1. As OpenStack admin, add the *SwiftOperator* role to the tenant user.

```
[root@rhos0 ~(openstack_admin)]# keystone user-role-add --user refarch \ --role SwiftOperator --tenant refarch-tenant
```

2. Switch to the tenant user environment.

```
[root@rhos0 ~(refarch_member)]# source /root/keystonerc_admin

[root@rhos0 ~(keystone_admin)]# env | grep OS_
OS_REGION_NAME=RegionOne
OS_PASSWORD=54e6a402ae8d4e47
OS_AUTH_URL=http://10.19.137.100:5000/v2.0/
OS_USERNAME=admin
OS_TENANT_NAME=admin
```

3. Copy the job source files to a temporary directory.

NOTE: Get example code from the Juno-stable openstack/sahara repository on Github: <https://github.com/openstack/sahara/archive/stable/juno.zip>

```
[root@rhos0 ~(refarch_member)]# wget -quiet \
https://github.com/openstack/sahara/archive/stable/juno.zip

[root@rhos0 ~(refarch_member)]# unzip -q juno.zip

[root@rhos0 ~(refarch_member)]# cd sahara-stable-juno/etc/edp-examples/pig-job/

[root@rhos0 pig-job(refarch_member)]# ls
data example.pig README.rst udf.jar
```

4. Create a Swift container and Sahara data sources for input and output.

```
[root@rhos0 pig-job(refarch_member)]# swift post container4

[root@rhos0 pig-job(refarch_member)]# sahara data-source-create \
--name rainput --type swift --url swift://container4.sahara/input \
--user refarch --password refarch
+-----+-----+

```



Property	Value
description	
url	swift://container4.sahara/input
tenant_id	24375feeb728444eabb8ce3e51a5170f
created_at	2015-04-23 02:35:40.122561
type	swift
id	a65cce11-a60a-4c4b-bc8a-9d7b8c291f01
name	rainput

```
[root@rhos0 pig-job(refarch_member)]# sahara data-source-create \  
--name raoutput --type swift --url swift://container4.sahara/output \  
--user refarch --password refarch
```

Property	Value
description	
url	swift://container4.sahara/output
tenant_id	24375feeb728444eabb8ce3e51a5170f
created_at	2015-04-23 02:35:52.115554
type	swift
id	75a03a46-2335-4cfd-98ce-01fc952c83c4
name	raoutput

5. Upload the input files to the Sahara input data store. Copy the *input* file from *data/input* to the local directory.

```
[root@rhos0 pig-job(refarch_member)]# sahara data-source-list
```

name	id	type	description
rainput	a65cce11-a60a-4c4b-bc8a-9d7b8c291f01	swift	
raoutput	75a03a46-2335-4cfd-98ce-01fc952c83c4	swift	

```
[root@rhos0 pig-job(refarch_member)]# cp data/input .
```

```
[root@rhos0 pig-job(refarch_member)]# for UPFILES in example.pig udf.jar input;  
do swift upload container4 ${UPFILES}; done
```

```
example.pig  
udf.jar  
input
```

```
[root@rhos0 pig-job(refarch_member)]# swift list container4
```

```
input  
example.pig  
udf.jar
```

6. Create the Sahara job binaries.

```
[root@rhos0 pig-job(refarch_member)]# for BINFILE in example.pig udf.jar;  
do sahara job-binary-create --name ${BINFILE} \  
--url swift://container4.sahara/${BINFILE} --user refarch \  
--password refarch; done
```



```

+-----+
| Property | Value |
+-----+
| description |
| url | swift://container4.sahara/example.pig |
| tenant_id | 24375feeb728444eabb8ce3e51a5170f |
| created_at | 2015-04-23 02:36:33.906692 |
| id | e405322b-2902-4dca-874e-82db4453303e |
| name | example.pig |
+-----+

```

```

+-----+
| Property | Value |
+-----+
| description |
| url | swift://container4.sahara/udf.jar |
| tenant_id | 24375feeb728444eabb8ce3e51a5170f |
| created_at | 2015-04-23 02:36:35.099204 |
| id | 9be8babd-a788-4586-9bef-48706f8f0a4a |
| name | udf.jar |
+-----+

```

```

[root@rhos0 pig-job(refarch_member)]# sahara job-binary-list
+-----+
| id | name | description |
+-----+
| 9be8babd-a788-4586-9bef-48706f8f0a4a | udf.jar |
| e405322b-2902-4dca-874e-82db4453303e | example.pig |
+-----+

```

7. Get the Sahara job binary IDs.

```

[root@rhos0 pig-job(refarch_member)]# JOB_MAIN_ID=$(sahara job-binary-list | awk
'/ example.pig / {print $2}')

[root@rhos0 pig-job(refarch_member)]# JOB_LIB_ID=$(sahara job-binary-list | awk
'/ udf.jar / {print $2}')

[root@rhos0 pig-job(refarch_member)]# echo $JOB_MAIN_ID
e405322b-2902-4dca-874e-82db4453303e

[root@rhos0 pig-job(refarch_member)]# echo $JOB_LIB_ID
9be8babd-a788-4586-9bef-48706f8f0a4a

```

8. Create a job template of type Pig that includes the job binaries.

```

[root@rhos0 pig-job(refarch_member)]# sahara job-template-create \
--name pig-test --type Pig --main ${JOB_MAIN_ID} --lib ${JOB_LIB_ID}
+-----+
| Property | Value |
+-----+
| description |
| tenant_id | 24375feeb728444eabb8ce3e51a5170f |
| created_at | 2015-04-23 02:38:28.548361 |
| mains | [{u'description': u'', u'url':
u'swift://container4.sahara/example.pig', u'tenant_id':
u'24375feeb728444eabb8ce3e51a5170f', u'created_at': u'2015-04-23 02:36:33',
u'updated_at': None, u'id': u'e405322b-2902-4dca-874e-82db4453303e', u'name':

```



```

u'example.pig'}}] |
| libs      | [{u'description': u'', u'url':
u'swift://container4.sahara/udf.jar', u'tenant_id':
u'24375feeb728444eabb8ce3e51a5170f', u'created_at': u'2015-04-23 02:36:35',
u'updated_at': None, u'id': u'9be8babd-a788-4586-9bef-48706f8f0a4a', u'name':
u'udf.jar'}}] |
| type      | Pig
| id        | 516dd0e3-14ff-4788-b19e-259c914d734b
| name      | pig-test
+-----+

[root@rhos0 pig-test(refarch_member)]# sahara job-template-create --name pig-
test --type Pig --main ${JOB_MAIN_ID} --lib ${JOB_LIB_ID}
+-----+
| Property  | Value
| description |
| tenant_id | 644a0e9be3ab4a42b1a19ca4f1fcf7f1
| created_at | 2015-04-21 06:28:20.190124
| mains     | [{u'description': u'', u'url':
u'swift://container4.sahara/example.pig', u'tenant_id':
u'644a0e9be3ab4a42b1a19ca4f1fcf7f1', u'created_at': u'2015-04-21 06:24:17',
u'updated_at': None, u'id': u'05a6be44-37b2-4391-8184-44a9f06c2dd7', u'name':
u'example.pig'}}] |
| libs      | [{u'description': u'', u'url':
u'swift://container4.sahara/udf.jar', u'tenant_id':
u'644a0e9be3ab4a42b1a19ca4f1fcf7f1', u'created_at': u'2015-04-21 06:24:18',
u'updated_at': None, u'id': u'ee9f6304-0eb6-4376-ada5-e7712ec55084', u'name':
u'udf.jar'}}] |
| type      | Pig
| id        | 30080ba1-2e08-4cf3-9982-58cd57d1dacb
| name      | pig-test
+-----+

```

9. Set the ID for the job template to an environment variable.

```

[root@rhos0 pig-job(refarch_member)]# JOB_TEMPLATE_ID=$(sahara job-template-
list | awk '/ pig-test / {print $2}')

[root@rhos0 pig-job(refarch_member)]# echo $JOB_TEMPLATE_ID
516dd0e3-14ff-4788-b19e-259c914d734b

```

10. Get the cluster ID of the previously launched cluster.

```

[root@rhos0 pig-job(refarch_member)]# CLUSTER_ID=$(sahara cluster-list | tail
-n +4 | awk '{print $4}' | head -n 1)

[root@rhos0 pig-job(refarch_member)]# echo $CLUSTER_ID
d38d9dad-867d-4e88-8b51-ac193fe42cf1

```

11. Set the job input data store ID to an environment variable.

```

[root@rhos0 pig-job(refarch_member)]# sahara data-source-list
+-----+-----+-----+-----+
| name      | id                | type  | description |
+-----+-----+-----+-----+
| rainput   | a65cce11-a60a-4c4b-bc8a-9d7b8c291f01 | swift |             |

```



```

| raoutput | 75a03a46-2335-4cfd-98ce-01fc952c83c4 | swift |
+-----+-----+-----+
[root@rhos0 pig-job(refarch_member)]# DATA_INPUT_ID=$(sahara data-source-list |
awk '/^\| rainput / {print $4}')

[root@rhos0 pig-job(refarch_member)]# echo $DATA_INPUT_ID
a65cce11-a60a-4c4b-bc8a-9d7b8c291f01

```

12. Set the job output ID to an environment variable.

```

[root@rhos0 pig-job(refarch_member)]# DATA_OUTPUT_ID=$(sahara data-source-list
| awk '/^\| raoutput / {print $4}')

[root@rhos0 pig-job(refarch_member)]# echo $DATA_OUTPUT_ID
75a03a46-2335-4cfd-98ce-01fc952c83c4

```

13. Launch a job from the job template to the cluster, specifying the input and output data stores.

```

[root@rhos0 pig-job(refarch_member)]# sahara job-create --job-template \
${JOB_TEMPLATE_ID} --cluster ${CLUSTER_ID} --input-data ${DATA_INPUT_ID} \
--output-data ${DATA_OUTPUT_ID}
+-----+-----+-----+
| Property | Value |
+-----+-----+-----+
| output_id | 75a03a46-2335-4cfd-98ce-01fc952c83c4 |
| job_id | 516dd0e3-14ff-4788-b19e-259c914d734b |
| tenant_id | 24375feeb728444eabb8ce3e51a5170f |
| created_at | 2015-04-23 02:52:23.837552 |
| status | PENDING |
| input_id | a65cce11-a60a-4c4b-bc8a-9d7b8c291f01 |
| cluster_id | d38d9dad-867d-4e88-8b51-ac193fe42cf1 |
| job_configs | {u'configs': {}, u'args': [], u'params': {}} |
| id | 6ceb7d6c-6eb1-4d31-a25d-f4a6f05f4666 |
+-----+-----+-----+

```

14. Verify that the job ran successfully.

```

[root@rhos0 pig-job(refarch_member)]# sahara job-list
+-----+-----+-----+
| id | cluster_id |
| status |
+-----+-----+-----+
| 6ceb7d6c-6eb1-4d31-a25d-f4a6f05f4666 | d38d9dad-867d-4e88-8b51-ac193fe42cf1 |
| SUCCEEDED |
+-----+-----+-----+

[root@rhos0 pig-job(refarch_member)]# sahara job-show \
--id 6ceb7d6c-6eb1-4d31-a25d-f4a6f05f4666
+-----+-----+-----+
| Property | Value |
+-----+-----+-----+
| output_id | 75a03a46-2335-4cfd-98ce-01fc952c83c4 |
| job_id | 516dd0e3-14ff-4788-b19e-259c914d734b |
| tenant_id | 24375feeb728444eabb8ce3e51a5170f |

```




```
job -oozie http://localhost:11000/oozie -definition 0000001-150423030653942-oozie-hado-w
```

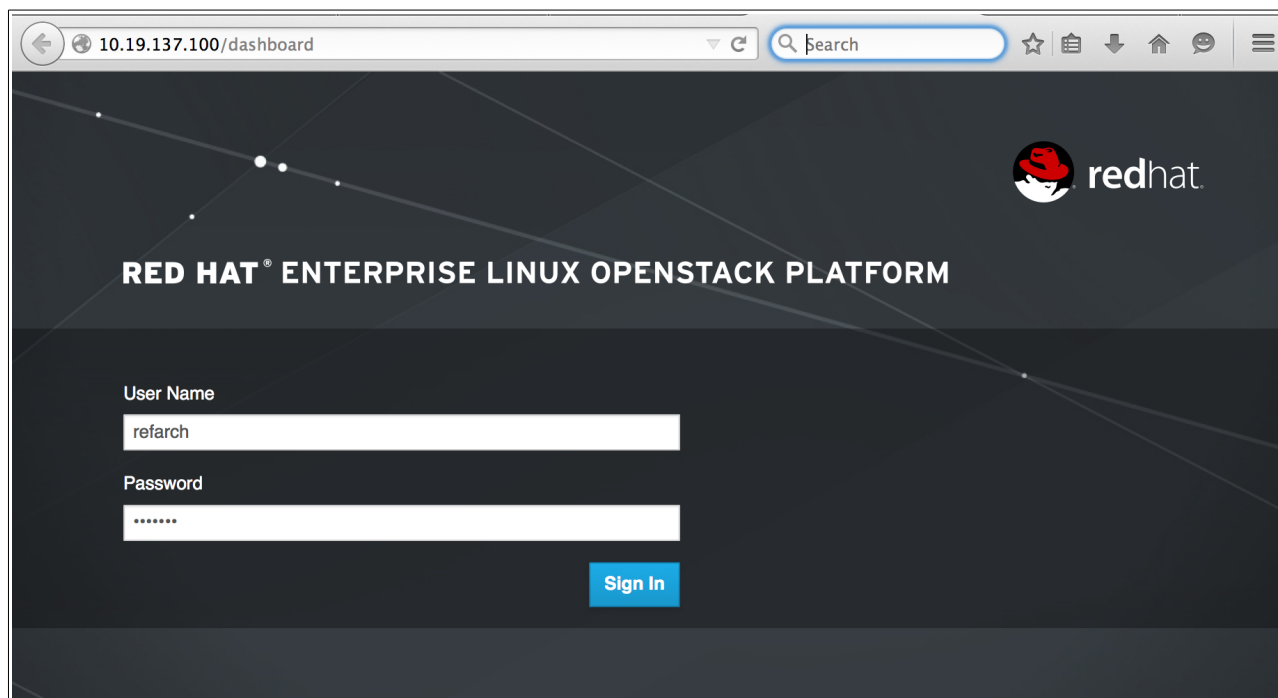
```
<?xml version="1.0" ?>
<workflow-app name="job-wf" xmlns="uri:oozie:workflow:0.2">
  <start to="job-node"/>
  <action name="job-node">
    <pig>
      <job-tracker>${jobTracker}</job-tracker>
      <name-node>${nameNode}</name-node>
      <configuration>
        <property>
          <name>fs.swift.service.sahara.password</name>
          <value>refarch</value>
        </property>
        <property>
          <name>fs.swift.service.sahara.username</name>
          <value>refarch</value>
        </property>
      </configuration>
      <script>example.pig</script>
      <param>INPUT=swift://container4.sahara/input</param>
      <param>OUTPUT=swift://container4.sahara/output</param>
    </pig>
    <ok to="end"/>
    <error to="fail"/>
  </action>
  <kill name="fail">
    <message>Workflow failed, error message[
{wf:errorMessage(wf:lastErrorNode())}]</message>
  </kill>
  <end name="end"/>
</workflow-app>
```

6.4 Verify Dashboard

This section describes brief steps for verifying Data Processing management via the OpenStack Dashboard.

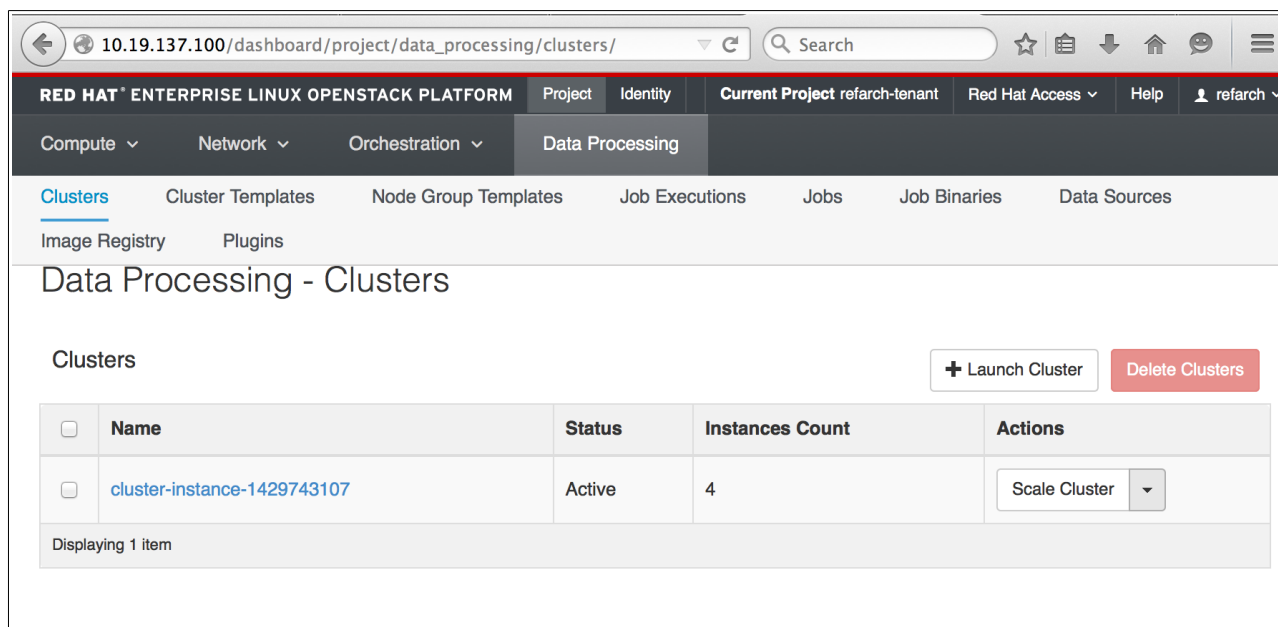


1. Connect to the OpenStack dashboard by IP address or hostname. Log in as the tenant user.



Graphic 6.4.1: Dashboard login screen

2. Select the Data Processing tab.
3. View the cluster instance by selecting the Clusters tab.



Graphic 6.4.2: Data processing clusters



NOTE: If the Dashboard does not include the Data processing tab, verify the following:

- Region is set correctly
- Service name is defined correctly. IE – data_processing



7. Conclusion

Sahara enables OpenStack users to deploy virtual Hadoop clusters and run data processing jobs on them. It combines the power of large scale data analytics with the convenience and flexibility of OpenStack. Jobs are defined by a collection of simple objects. Users create modular job, node, and cluster objects from JSON templates and store them in the Sahara database for persistence and reuse. Sahara empowers its users to focus on data analytics instead of administering Hadoop infrastructure.

This reference architecture demonstrates how to implement Sahara in a RHEL OSP 6 environment and submit a test job. It also describes Sahara's high level architecture and provides an overview of how it leverages the OpenStack ecosystem for image management, authentication, and data storage.

This reference architecture focuses on the single user to small group use case of deploying a Hadoop cluster for Dev or QA. Other use cases such as utilizing unused compute power from a general purpose OpenStack cloud for Hadoop jobs or "Analytics-as-a-Service" for bursty or ad-hoc workloads could be topics for future reference architectures.



Appendix A: References

1. Red Hat Enterprise Linux 7 installation instructions: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Installation_Guide/index.html
2. Red Hat Enterprise Linux OpenStack Platform 6 deployment guide: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/6/html/Installer_and_Foreman_Guide/index.html
3. Building images for Sahara Plugins_ <http://docs.openstack.org/developer/sahara/userdoc/diskimagebuilder.html>



Appendix B: Revision History

Revision 2.0	Monday April 27, 2015	Jacob Liberman
<ul style="list-style-type: none">• Updated to OSP 6 and RHEL 7• Incorporated engineering feedback		
Revision 1.0	Thursday October 9, 2014	Jacob Liberman
<ul style="list-style-type: none">⤴ Incorporated Services review feedback⤴ Final layout and spacing⤴ Verified internal document links		
Revision 0.4	Wednesday October 1, 2014	Jacob Liberman
<ul style="list-style-type: none">⤴ Finalized content⤴ Incorporated Systems Engineering review feedback		
Revision 0.3	Monday September 29, 2014	Jacob Liberman
<ul style="list-style-type: none">⤴ Incorporated Product Engineering review feedback		
Revision 0.2	Friday September 5, 2014	Jacob Liberman
<ul style="list-style-type: none">⤴ Table of contents⤴ Executive summary⤴ Ported to template version 11		