


BRNO UNIVERSITY OF TECHNOLOGY
VYSOKÉ UČENÍ TECHNICKÉ V BRNĚ



FACULTY OF INFORMATION TECHNOLOGY
DEPARTMENT OF INTELLIGENT SYSTEMS

FAKULTA INFORMAČNÍCH TECHNOLOGIÍ
ÚSTAV INTELIGENTNÍCH SYSTÉMŮ

EXTENSION OF OPENSTACK MODULES FOR ANSIBLE PLATFORM

ROZŠÍŘENÍ MODULŮ OPENSTACK PRO PLATFORMU ANSIBLE

BACHELOR'S THESIS

BAKALÁŘSKÁ PRÁCE

AUTHOR

AUTOR PRÁCE

ADAM ŠAMALÍK

SUPERVISOR

VEDOUČÍ PRÁCE

Ing. MARTIN HRUŠKA

BRNO 2016

Abstract

OpenStack is a cloud platform with distributed architecture that is very complex to deploy. In this thesis, I will design an Ansible playbok (automatic deployment script) to deploy a custom OpenStack architecture.

Abstrakt

OpenStack je cloudová platforma s distribuovanou architekturou, jejíž instalace je velice komplexní. V této práci navrhnu Ansible Playbook (automatický instalační skript), který OpenStack nainstaluje.

Keywords

OpenStack, Ansible, deployment, cloud, automation

Klíčová slova

OpenStack, Ansible, instalace, cloud, automatizace

Reference

ŠAMALÍK, Adam. *Extension of OpenStack Modules for Ansible Platform*. Brno, 2016. Bachelor's thesis. Brno University of Technology, Faculty of Information Technology. Supervisor Hruška Martin.

Extension of OpenStack Modules for Ansible Platform

Declaration

I declare that I have wrote this Bachelor's Thesis individually, with help from my supervisor, Ing. Martin Hruška, and technical consultant, Miroslav Suchý from Red Hat. I have included all resources and literature I have used for this Thesis.

.....
Adam Šamalík
May 16, 2016

Acknowledgements

I would like to thank to my supervisor, Ing. Martin Hruška, who gave me extensive feedback about technical writing and who helped me with the overall structure of this thesis. His feedback and constructive criticism was much appreciated. I would also like to thank to my technical consultant, Miroslav Suchý from Red Hat, who gave me strong support about the technologies and shared his technical experience. Thank you.

© Adam Šamalík, 2016.

This thesis was created as a school work at the Brno University of Technology, Faculty of Information Technology. The thesis is protected by copyright law and its use without author's explicit consent is illegal, except for cases defined by law.

Contents

1	Introduction	3
1.1	Relation with Red Hat	3
2	Technology Overview	4
2.1	Cloud Computing Introduction	4
2.1.1	Difference Between Traditional Virtualisation and Cloud Computing	4
2.1.2	Cloud Applications and Scaling	5
2.1.3	Models of Cloud Computing	5
2.1.4	Models of Cloud Deployment	6
2.2	OpenStack	6
2.2.1	Core OpenStack Components	6
2.3	Ansible	8
2.3.1	What is Ansible	8
2.3.2	Advantages of the Ansible Technology	9
2.3.3	Ansible Playbooks	10
3	OpenStack Architecture	13
3.1	Conceptual Architecture	13
3.2	Logical Architecture	13
3.3	Physical Architecture	14
3.4	Installation of Core Components	14
3.4.1	OpenStack Identity Service	15
3.4.2	OpenStack Image Service	17
3.4.3	OpenStack Compute Service	17
3.4.4	OpenStack Networking Service	20
3.4.5	OpenStack Block Storage Service	22
3.4.6	OpenStack Dashboard Service	24
4	Existing Methods of Automated OpenStack Deployment	25
4.1	Packstack	25
4.2	OpenStack-Ansible	25
5	Implementation Design	26
5.1	Design of the Physical Architecture	26
5.1.1	Selecting OpenStack Services	26
5.1.2	Selecting Specific Implementations	27
5.1.3	Physical architecture	27
5.2	Design of the Ansible Playbook	28

5.2.1	Ansible Roles	30
5.2.2	Applying Roles to the Hosts	35
6	Implementation and Testing	37
6.1	Third-party Modules in the Playbook	37
6.2	Testing of the Deployment	37
6.2.1	Description of the Host Environment	37
6.2.2	Using Ansible with Vagrant	38
6.2.3	Running the Deployment	38
6.2.4	Testing the Deployment	39
7	Conclusion	42
7.1	Future Development	42
	Bibliography	43
	Appendices	45
	List of Appendices	46
A	CD Contents	47
B	Manual	48
B.1	Using Roles with Variables	48
B.2	Variables of Individual Roles	49
B.2.1	Role <code>keystone</code>	49
B.2.2	Role <code>glance</code>	49
B.2.3	Role <code>nova-controller</code>	50
B.2.4	Role <code>nova-compute</code>	50
B.2.5	Role <code>neutron-controller</code>	51
B.2.6	Role <code>neutron-compute</code>	52
B.2.7	Role <code>cinder-controller</code>	52
B.2.8	Role <code>cinder-storage</code>	53
B.2.9	Role <code>dashboard</code>	53
B.2.10	Role <code>rabbit</code>	53
B.2.11	Role <code>sql-database</code>	54

Chapter 1

Introduction

Originally, server applications were running on dedicated physical servers. If the application needed more resources, the server would be upgraded with additional CPUs, memory, or storage which required physical access and only allowed adding whole new disk, or whole new memory board. Virtualisation (running virtual machines on a single physical host) solved both limitations, as the virtual machines could be managed remotely and only small chunks of storage or memory could be added. Cloud computing is another step in the evolution, as it creates an abstract layer of resources over multiple servers and provides advanced techniques such as scheduling and load balancing. All of this is described in more detail in section 2.1.

OpenStack is a cloud platform with distributed architecture, and is briefly described in section ???. Deploying OpenStack can be a complex task requiring deep knowledge in distributed system architecture, networking, storage, and virtualisation. Because of its complexity, it is very hard to create an installation script. There are several projects that try to simplify the process. They are described in chapter 4.

However, before I start describing the methods of installing OpenStack, the reader should understand automatic application deployments which is described in section 2.3 using technology called Ansible.

As said before, OpenStack is a complex system with distributed architecture. It is described in more detail in chapter 3 which describes the architecture of OpenStack and its components.

The goal of this thesis is to design and implement an installation script which would enable to install basic OpenStack deployment and would also offer flexibility in configuring more complex scenarios by reusing the individual components of the script. The design of a reference OpenStack architecture and the script is described in chapter 5.

The final chapter 6 shows some details about the implementation process, describes the testing environment and will also present results of the testing.

1.1 Relation with Red Hat

This thesis has been written in cooperation with Red Hat¹, a leading open-source technology company. Some parts of the installation script designed and created as part of this thesis will be reused by the Fedora Project² in its own infrastructure.

¹<https://www.redhat.com>

²<https://getfedora.org>

Chapter 2

Technology Overview

2.1 Cloud Computing Introduction

Originally, we used dedicated servers to run enterprise applications. One application would run on a single physical machine. These machines could be small servers or large mainframes. When an organization grown and their application required more resources, the server would scale up, which means adding more processors, memory, or storage to the host. Is is also referred to as vertical scaling. This method was not flexible enough, as every change required physical access to the datacentre and buying new hardware. Also, it only allowed upgrading by rather large steps, for example by adding whole new hard drive, memory module, etc. This problem has been solved by virtualization.

Virtualization is a technology that enables running several virtual servers on a single host. Every virtual server would have its own operating system and would act as a physical host. These servers could be dynamically scaled up and down without the need of a physical access to the datacentre. New virtual servers could be provisioned, unused servers could be deleted to make space for scaling, which would reflect the business needs in more flexible way then dedicated servers. This method is also referred to as a traditional virtualization. It is because there is a new technology called cloud computing that will be described right now in the following section.

2.1.1 Difference Between Traditional Virtualisation and Cloud Computing

In a traditional virtualization, virtual servers are created and managed as part of a single physical host. Virtual machines on every physical host must be managed separately, directly on the host. This approach, again, might not seem flexible enough for large corporations running thousands of virtual servers on many physical hosts in several datacentres. This is why a technology called cloud computing has emerged.

Cloud computing, in a similar way as traditional virtualization, is a technology that manages running of virtual machines. The main difference is that in cloud computing, a virtual layer of resources, such as computing power, memory, and storage, is created over one or more datacentres. These resources are then used to provision virtual servers. It also offers scheduling and load-balancing. Scheduling means, that the cloud computing platform automatically chooses on which host the virtual server would run. Load balancing means that requests coming to an application running on multiple virtual machines are evenly distributed across all of them.

This approach is very flexible in terms of creating, deleting, and scaling virtual servers. It also brings new technology, such as software-defined networking (SDN), or an object store, both of which will be described later in the thesis.

2.1.2 Cloud Applications and Scaling

The cloud computing flexibility opens new possibilities in the way enterprise applications are designed and created. Instead of building a large monolithic application that would run on a single host, the application has been broken down to smaller parts communicating with each other. These parts are called microservices and each part runs on a separate virtual machine. This approach of creating an application consisting of several microservices is called distributed architecture.




Figure 2.1: Vertical and horizontal scalability [Source: [1]]

This architecture style enabled another innovation in terms of scaling. As I described in the beginning of this section, with physical servers and also in traditional virtualization, applications are scaled up by adding more processors, memory, or storage. In cloud computing, the application can scale by adding more virtual machines to handle the workload. This approach is called scaling out. This approach is also referred to as horizontal scaling. See the picture 2.1 for more details. Such scaling can be also done automatically, on demand. This means that the application does not need to have all the resources reserved for itself all the time. For example, an application can run in business' private datacentre, which provides the computing power for the majority of time. In a need of extra computing power, the application can automatically scale out and use resources from another datacentre, which might be run by different company as a paid service. This leads us to different models of cloud computing and deployment models.

2.1.3 Models of Cloud Computing

There are several models [17] describing the the terms in which are cloud computing resources used. It varies from using the whole infrastructure to using only an environment, in which an application would run.

- **Infrastructure as a Service (IaaS)** provides computing resources, which can be used to provision virtual machines and software-defined networks. In this model, the cloud consumer manages the application, operating system, storage, networking, and computing resources. OpenStack is an IaaS solution.
- **Platform as a Service (PaaS)** provides an operating system including libraries and programming languages. In this model, the cloud consumer manages the application, but not the underlying infrastructure.
- **Software as a Service (SaaS)** provides an operating system and the application. In this model, the cloud consumer does not manage the application nor the underlying infrastructure.

2.1.4 Models of Cloud Deployment

There are several ways [17] where to deploy the cloud computing infrastructure. A company can a cloud in their own datacentre, or they can use a public provider. A combination of both is also common.

- **Public cloud** is a cloud run by a cloud provider, provided as a service.
- **Private cloud** is a cloud run and used by a single organization.
- **Hybrid cloud** is a combination of public and private cloud. The cloud consumer has the ability to run an application in their datacentre and expand to the private cloud when more resources are needed. It is commonly achieved by two independent cloud deployment that use single interface to manage the resources. An example of such interface can be Red Hat CloudForms.

2.2 OpenStack

OpenStack is an open-source cloud platform that controls large pools of compute power, storage, and networking resources of one ore more datacentres. These resources are then used to provision virtual machines.

The service is managed by a single web interface called dashboard or a command-line interface, which gives the administrators control over the whole cloud, and users the ability to provision and manage resources via single web interface. OpenStack consists of several components. The base six components provide identity, networking, object storage, block storage, a service for managing images, and the compute service. Other services may include an orchestration service, telemetry service to measure usage of the cloud, or a database as a service solution (DBaaS). These services will be described in more detail later in the thesis.

All services communicate with each other via public HTTPS endpoints and use Advanced Message Queuing Protocol (AMQP), which is a protocol supporting sending and receiving messages between distributed systems. [17]

2.2.1 Core OpenStack Components

OpenStack is composed of several components [17] called services. Some of them can be seen on picture 2.2, and are also described in this section.

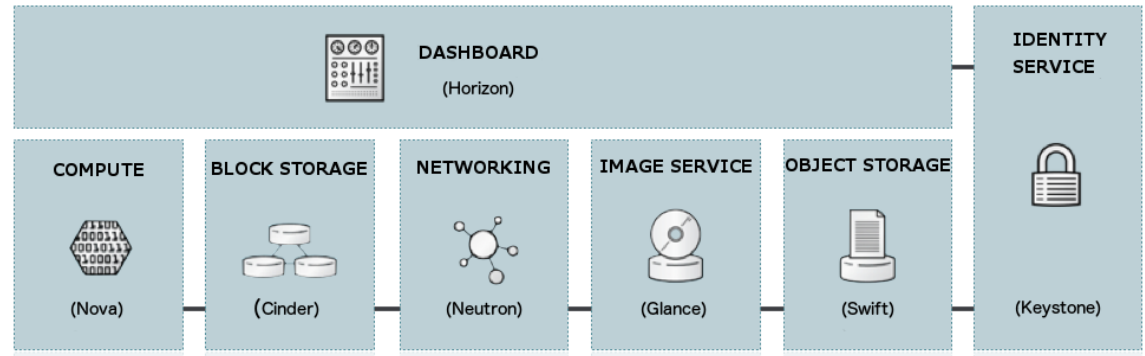


Figure 2.2: OpenStack core services [Source: [2]]

- **Horizon (dashboard)** is a web interface for managing OpenStack services providing graphical user interface. Horizon supports operations like launching instances, managing networking, and setting access controls.
- **Keystone (identity)** is a centralized identity service providing authentication and authorization for other services. It also provides a catalogue of services running in the OpenStack cloud.
- **Neutron (networking)** provides software-defined networking (SDN) in the OpenStack cloud, including networks, subnets and routers. Other advanced services such as firewalls, virtual private networks (VPN) or quality of service (QoS) are also supported. This service handles the creating and management of the networking infrastructure for the cloud administrator and users.
- **Cinder (block storage)** manages persistent block storage volumes used by virtual machines. The service supports creating of snapshots which can be used for backing up data. Then backup can be then used to restore data or to create some new block storage volumes. This service is often used by the virtual machines as a storage.
- **Swift (object storage)** provides an object storage that allows users to store and retrieve files. Swift has a distributed architecture that enable horizontal scaling and redundancy. Data replication is managed by software, which allows larger scalability and redundancy than dedicated hardware.
- **Glance (image)** is a registry of virtual machine images. Users can copy server images and use them as templates when setting up new virtual servers.
- **Nova (compute)** is a service that manages virtual machines running on compute nodes. Nova is designed to scale horizontally on standard hardware and to download images to launch new instances. Nova is a distributed component that interact with Keystone for authentication, Glance for images and Horizon for a web interface. It uses libvirt, qemu, and kvm for the hypervisor.
- **Ceilometer (metering)** is a service that provides a centralized source for metering and monitoring data, which can be used to meter and bill users.
- **Heat (orchestration)** an advanced service to orchestrate multiple cloud application using the Amazon Web Services (AWS) template format. The software integrates

other core components of OpenStack into a one-file template system. Templates can be used to create most of OpenStack resources, such as instances, floating IPs, volumes, security groups or users. Heat also offers advanced functionality such as high availability, auto scaling and nested stacks.

2.3 Ansible

2.3.1 What is Ansible

Ansible is an automation engine which can be used for application deployment. It is also suitable as a configuration management tool, for orchestration of deployment and provisioning new servers, all of which will be described later in the section. [4]

Deployment

Application deployment is a process of taking a source code of a software that has been written internally, building it to get the binaries, copying the required files to the servers that are supposed to run the application, and then starting the necessary services.

Tools other than Ansible that can be used for deployment are for example Capistrano, or Fabric, which are both open-source. [16]

As Ansible can manage both deployment and configuration management, using a same tool for these tasks can make it easier and well-arranged for the people using it.

Configuration management

Configuration management typically means writing some kind of description of our servers - a state in which we want the servers to be. This definition of a state can include information about packages that should be installed (or, more precisely, present) on the server, configuration files containing specific values and have specific permission and ownership, that the correct services are running, and so on.

Except Ansible, common configuration management tools are for example Chef, Salt, Puppet, or cfEngine. [16]

Orchestration of deployment

We need to use orchestration when there are more servers involved and we need to run tasks in a specific order. This can mean, for example, bringing up the database server first and then starting the web servers. Or in case of updating web servers, we need to take them out of the load balancer one at a time to preserve stable service without outages.

Ansible is designed to support orchestration by using a simple model of dependencies build directly into the core architecture - so the tasks are automatically executed in the correct order.

Provisioning

The term provisioning a new server is often used in a cloud environment. It means creating a new virtual machine instance with operating system installed and having the networking and access configured in a way that the machine can be used.

However, provisioning is also used with baremetal machines, because, in the end, all virtual and cloud machines run on a physical server. Ansible can provision all the machines

in the datacenter and can be also integrated with several datacentre management tools, like (not limited to) Red Hat Satellite, Hanlon, or Cobler.

Ansible supports several cloud providers including Amazon EC2, Microsoft Azure, Digital Ocean, Google Compute Engine, Linode, Rackspace and clouds supporting the Open-Stack APIs. [16]

2.3.2 Advantages of the Ansible Technology

There are several features of the Ansible technology that are important to mention:

Syntax is Easy to Read

The ansible configuration management scripts are build on top of YAML syntax, which is a data format language that was designed to be easy to read for people.

When used properly, the ansible playbooks might be seen as an executable documentation. And it will never be outdated, because it is also the code, that gets executed.

Push-based mechanism

There are two main concepts of deploying a change to servers: push-based and pull-based mechanisms.

Some configuration management, that use agents, are pull-based. An example of these tools can be Puppet or Chef. The pull-based mechanism work in the following way:

1. Administrator makes change to the configuration scripts
2. Administrator pushes the changes to a central service
3. Agent running on the server checks for changes periodically
4. Agent downloads the change from the central service
5. Agent executes the configuration script locally and changes the state of the server

Ansible uses the push-based mechanism, in which the administrator controls when a change is applied to the server. [16] The administrator does not need to wait for a timer to expire before the change is applied. This is an advantage of the Ansible solution, as it offers more control over the overall system. The push-based mechanism works in the following way:

1. Administrator makes change to the configuration scripts
2. Administrator runs the playbook
3. Ansible connects to the servers and executes modules, making changes to the server

Ansible, however, also supports a pull-based mechanism, using a tool called ansible-pull.

No Agents Required

To manage a server by ansible, it only needs a Python version 2.5 or later and an SSH to be installed. Ansible does not need any agents or additional software to be installed on the servers. It also does not need any special management interfaces, as it runs on the existing network infrastructure using SSH.

Thin Layer of Abstraction

Some configuration management tools provide a thick layer of abstraction, which enables the administrator to manage multiple operation systems using a single script. For example, an abstraction called package could be used to install a package on a system, regardless its type. It could be a yum-based, or a apt-based system. However, these abstractions can be even bigger which means that administrator often need to learn how these abstractions work. Learning these abstractions would be an extra effort, as the administrators should know how their system work anyway.

Ansible does not use these abstractions. This allows the administrator to write scripts for a specific system without the need of learning the abstractions, which makes the learning process shorter, and reading the scripts easier. [16]

2.3.3 Ansible Playbooks

As already mentioned previously, ansible configuration scripts are called playbooks. Playbooks can define the configuration of the target servers, and they can orchestrate the steps in which the configuration is applied. Playbooks consist of several parts: plays, tasks, and modules. [16] [14]




Figure 2.3: Ansible playbook anatomy [Source: [16]]

Plays

Playbook is a list of plays. Plays can be thought of as a connection of tasks to hosts. A play must contain:

- A list of hosts
- A list of tasks that would be executed on the hosts

Plays also support several options that affect the way it is executed. An example of these settings can be:

- **name** - A description of the respective task. The names should be always used, as they will be printed at the time of run, showing what changes has been made to a particular server.
- **sudo** - A boolean value determining if the tasks will be executed with root privileges.

- `vars` - A list of variables that would be used for the particular task. They act as a parameters.

An example of a play, taken from [16], is shown below:

```
- name: Configure webserver with nginx
  hosts: webservers
  sudo: True
  tasks:
    - name: install nginx
      apt: name=nginx update_cache=yes

    - name: copy nginx config file
      copy: >
        src=files/nginx.conf
        dest=/etc/nginx/sites-available/default

    - name: enable configuration
      file: >
        dest=/etc/nginx/sites-enabled/default
        src=/etc/nginx/sites-available/default
        state=link

    - name: copy index.html
      template: >
        src=templates/index.html.j2
        dest=/usr/share/nginx/html/index.html mode=0644

    - name: restart nginx
      service: name=nginx state=restarted
```

Tasks

Play is a list of tasks. Every task represents a single module that would be executed. Every task must contain:

- A name of the particular module
- A list of arguments for the module

Task can also include a definition of its name. Using of names is a good practise, as they are printed at the time of run and can be used to track changes made by the scripts. They also act as comments, improving readability of the playbook.

Modules

Modules are scripts that perform the desired action on the target server.

They are designed to be *idempotent*, which means that they can be run multiple times, but will only make changes when there is a difference between the present and desired state.

In other words, running a playbook for the second time without making changes will not affect the target server in any way.

An example of modules can be:

- `copy` - Copies a file from the local machine to the remote host.
- `service` - Manages the services running on the remote host. It can start, stop, or restart a services.

Handlers

When ansible execute modules on a particular server, they may or may not make a change. They only make change when the present state is different from the desired state. At the time of run, ansible prints out the names of tasks that are being executed. If the state has changed, the module will return changed state. And if no action was needed, an ok state will be returned instead. Ansible will then print the result, as seen on the following example:

```
TASK: [Install packages] *****
changed: [webserver1]

TASK: [Set configuration] *****
ok: [webserver1]
```

This detection of a change can be used by a mechanism called handlers. [16] Handlers are actions triggered at the end of run only when a specific change has occurred. However, the action will be executed only once.

For example, a change of a web server configuration might require restart of the particular service. Also, an update of the web server package might require a restart of the service. If any of these actions occur, a handler responsible for restarting the service will be notified and executed at the end of run. It will be executed only once, even when both of the changes above will occur.

Roles

Roles are the main mechanism to break large playbooks into multiple files, which simplifies writing complex playbooks. It also makes them easier to reuse.

When managing a large number of servers, a single playbook describing changes of all of the servers would become too long and hard to maintain. Also, applying a common configuration to multiple servers would result in a duplication of code. To solve this problem, a mechanism called Roles has been introduced to Ansible. [14]

A role can be thought of as something that should be assigned to one or more hosts. For example, a database role will be assigned to servers acting as a database servers.

Roles can also be dependent, which means that a role can require other role to be applied before. Using this feature, Ansible will automatically ensure that roles are applied in the correct order.

Chapter 3

OpenStack Architecture

This chapter describes the OpenStack architecture, how the components communicate and what they are used for. Selected components, which are needed for the deployment option used later in this thesis, will be described in more detail.

3.1 Conceptual Architecture

Conceptual architecture is a high level architecture that shows the core components of a system and their relations including description of responsibilities for each component.

The relationships among the OpenStack services can be seen on picture [3.1](#).

3.2 Logical Architecture

Logical architecture shows processes and functions that make the system functioning.

OpenStack consists of several independent parts called OpenStack services.[\[17\]](#) Authentication for all services is provided by a common identity service. The services interact with each other via public APIs. [\[5\]](#)

OpenStack services are internally composed of several processes. Each service has at least one API process. This API process listens for API requests, does necessary preprocessing and passes them to the particular process or processes within the service. The actual work of a service is done by distinct processes. An exception to this rule is the identity service.

The processes within a service communicate with each other via message broker using AMQP protocol. The most used AMQP broker is RabbitMQ, which is also used in this thesis.

The state of a service is stored in a database. OpenStack support several databases, such as MySQL, MariaDB, and SQLite. The database solution used in this thesis is MariaDB. [\[5\]](#)

Users can interact with OpenStack in several ways. There is a graphical web-based interface, which is implemented by OpenStack Dashboard. Other ways include command-line clients provided by all basic services, and by API requests using web browser plugins or curl.

The most common, but not the only possible logical architecture can be seen on picture [3.2](#).

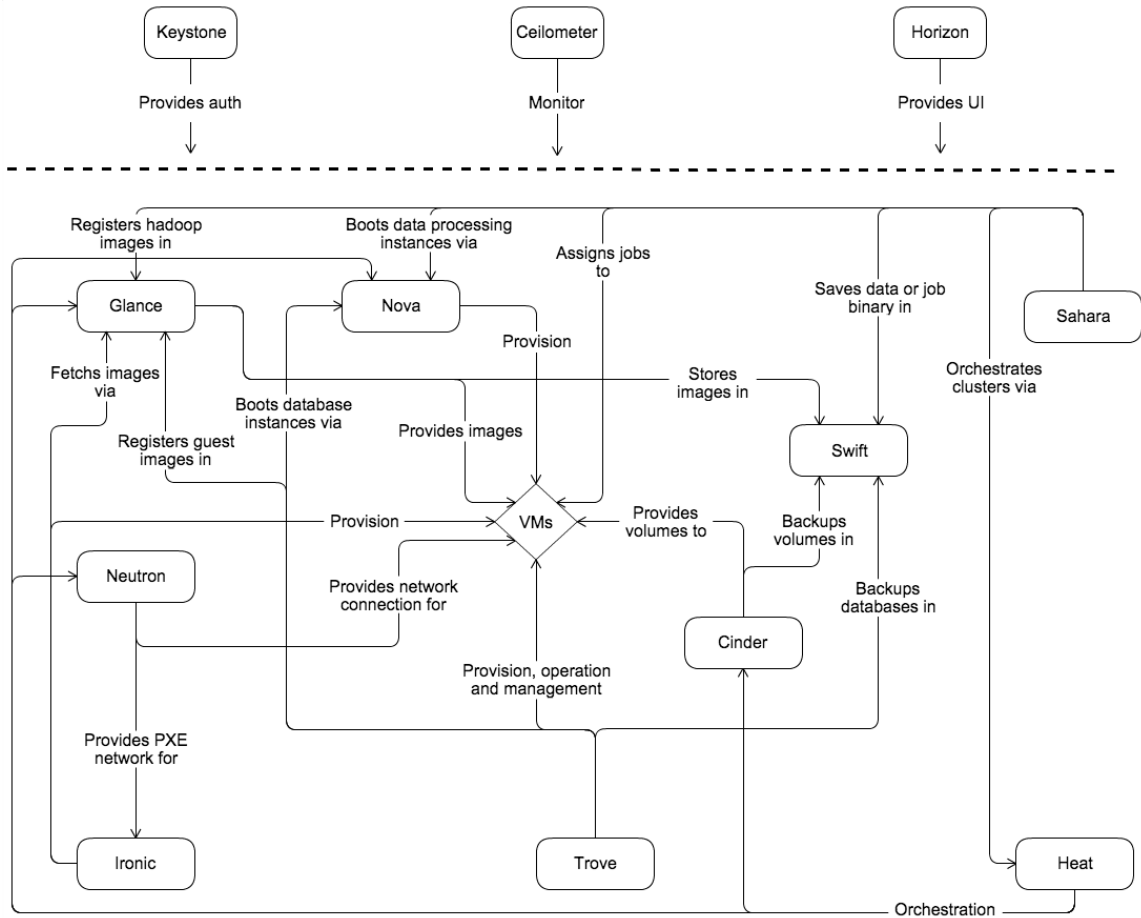


Figure 3.1: OpenStack conceptual architecture [Source: [3]]

3.3 Physical Architecture

OpenStack is a distributed system, which enables cloud architects to design varieties of physical architecture reflecting needs of the cloud consumer.

“Designing an OpenStack cloud is a great achievement. It requires a robust understanding of the requirements and needs of the cloud’s users to determine the best possible configuration to meet them. OpenStack provides a great deal of flexibility to achieve your needs...” [15]

3.4 Installation of Core Components

This section will explain each service of the OpenStack cloud in more detail including general description, some basic principles of the function, and a list of components of the service. Some services, such as Compute or Storage, can use several backends for the actual work. List and short description of these backends is also included in this section.

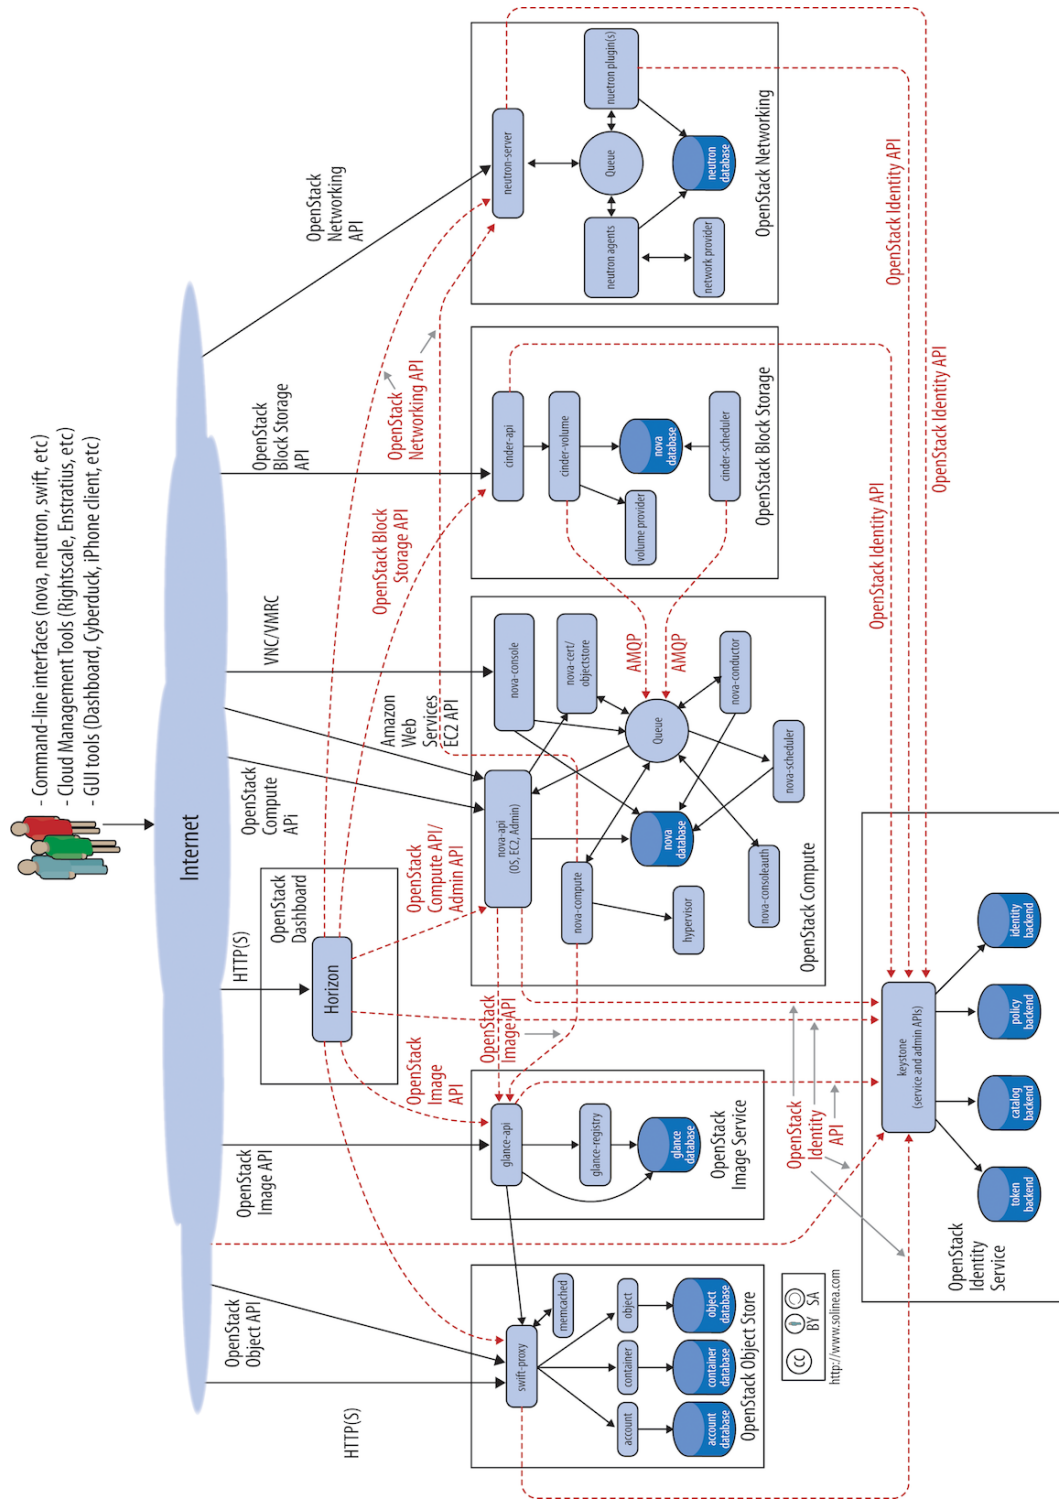


Figure 3.2: OpenStack logical architecture [Source: [15]]

3.4.1 OpenStack Identity Service

The OpenStack identity service provides a single point of authentication and authorisation. When OpenStack services receive a request from a user, they check with the identity service

to verify the authorisation of the user.

It also provides a catalog of OpenStack services. When installing OpenStack cloud, each service needs to be registered in the identity service. The identity service then tracks which services are installed and where they are located on the network.

OpenStack services also use the identity service as a common unified API.

Identity Abstractions

The identity service uses the following abstractions:

- **Service** is an OpenStack service, such as the identity (Keystone) service, compute (Nova) service, etc.
 - **Endpoint** is an address accessible via network and acts as an access-point of a service.
 - **Region** represents a general division of the OpenStack deployment. The identity service also supports sub-regions, which allows the creation of a tree-structured hierarchy.
 - **Authentication** is the process of confirming the identity of a user. To confirm an incoming authentication request, the identity service validates a set of credentials supplied by the user. These can be a name and password, or a name and API key. After validation of the user credentials, the identity services then issues an authentication token. This process is done only once on in the beginning of a session and user then provide this token in all subsequent requests.
 - **Credentials** are data to identify the user, such as name and password, or name and API key.
 - Authentication **token** is an alpha-numeric string used to access to OpenStack APIs and resources. Token is valid for a finite duration and can be revoked at any time.
 - **Domain** is a collection of projects and users. This collection defines administrative boundaries for managing identity entities. Domains can be used to represent an individual person, a company, or an operator-owned space. User can be granted the administrator role within a domain. This domain administrator can then create users, projects, and groups in the domains as well as create and assign roles to users or groups.
 - **Project** is used to isolate resources or identity objects. Projects can be used to map to a customer, account, organisation, or tenant.
 - **Role** is a personality containing a set of user rights and privileges. A user contains a list of roles.
 - **Group** is a collection of users and is owned by a domain. A group can be assigned a role that applies to all users in the group.
 - **User** is a representation of a person, system or service using the OpenStack services.
- [17]

OpenStack Identity Components

The identity service consists of these components:

- **Server** - A centralised server, providing authentication and authorisation services. It uses a RESTful interface.
- **Drivers** - Drivers are used to access identity information provided by external repositories like LDAP, or existing SQL databases.
- **Modules** - Middleware modules that check service requests, extract user credentials and send them to the Server for authorisation. The modules are integrated with the OpenStack components by using the Python Web Server Gateway Interface (Python WSGI).
[6]

3.4.2 OpenStack Image Service

The OpenStack Image service controls image storage and management and allows users to discover, register, and retrieve images.

Images provide templates for virtual machine filesystems. Each virtual machine runs from a copy of a base image and several virtual machines can be run from a single base image. Any changes made to the virtual machines will not affect the image. Users can also create a snapshot - a state of a virtual machines running disk - and build a new image based on these snapshots.

The implementation of the image service is called Glance. It supports various backends as a storage, which can be normal filesystems, OpenStack object storage, HTTP, RADOS block devices, and Amazon S3. Some of the storage backends can be read-only. The architecture in this thesis will use normal filesystem mounted to the controller node.

Basic Components

The image service consists of the following components:

- **glance-api** - Accepts API calls.
- **glance-registry** - Stores, processes, and retrieves metadata about images like size and type. This is an internal service only and should not be exposed outside of the image service.
- **Database** - An SQL database to store image metadata.
- **Storage** repository for image files - Backend that stores the images itself.

3.4.3 OpenStack Compute Service

The OpenStack Compute service is the major part of an Infrastructure as a Service (IaaS) system. It runs and manages the virtual machines running in the OpenStack cloud. It can also provide networking for the virtual machines. However, Nova networking is not often used. The OpenStack Networking (Neutron) , described in this section, is used instead. [6]

The OpenStack Compute service scales horizontally and is commonly deployed on multiple hosts, [15] often called compute nodes.

The service manages virtualisation, but it does not include any virtualisation software. Instead, it uses drivers to interact with the underlying virtualisation backend, also called virtualisation provider. It is also possible to use multiple providers in different availability zones.

The OpenStack Compute service supports these following virtualisation providers: [5]

- Baremetal
- Docker
- Hyper-V
- Kernel-based Virtual Machine (KVM)
- Linux Containers (LXC)
- Quick Emulator (QEMU)
- User Mode Linux (UML)
- VMware vSphere
- Xen

Service Architecture

The architecture of the Compute service could be divided into the following four parts:

- **API Server** is the compute controller that commands and controls the hypervisor, storage, and networking available to the end users. It manages the API endpoints, which are basic HTTP web services that handle the authentication, authorisation, and basic command and control functions. It support various API interfaces including Amazon, Rackspace, and others, which enables compatibility with multiple existing tools already created for other cloud platforms. This open approach also prevents vendor lock-in.
- **Message Queue** manages the interaction between compute nodes, the networking controllers, API server, and the scheduler. It is provided by the Advanced Message Queuing Protocol (AMQP). Several implementations are available, the most common are RabbitMQ, Qpid, ZeroMQ, and others.
- **Compute Worker** manages the virtual machines that run on the compute nodes. It is managed via API that provides commands to run, delete, and reboot instances, attaching and detaching volumes, and to get local console output.
- **Network Controller** manages the networking resources available on the compute nodes. It is managed via API that provides commands for allocating fixed IP addresses, configuring VLANs for projects, or configuring networks for compute nodes. [5]

OpenStack Compute Components

The OpenStack Compute service consists of the following components:

- **nova-api service** - Accepts API calls and responds to them. It supports the OpenStack Compute API, the Amazon EC2 API, and a special Admin API.
- **nova-api-metadata service** - Accepts metadata requests from instances.
- **nova-compute service** - A worker daemon that manages virtual machines via virtualisation provider API. It supports XenAPI, libvirt, VMwareAPI, and others. The state of the instances is always saved in the database.
- **nova-scheduler service** - This service determines on which compute node a new virtual machine should be started.
- **nova-conductor module** - It is a point of access to the database for other nova services. It eliminates direct access to the database by other services. This service scales horizontally and can be deployed on multiple hosts.
- **nova-cert module** - This is only needed to use with the Amazon EC2 API, used to generate certificates.
- **nova-network worker daemon** - Similar to the nova-compute service, but for networking. Accepts networking tasks from the queue and manipulates the network.
- **nova-consoleauth daemon** - Provides authorisation for users provided by the console proxies - see nova-novncproxy and nova-xvpngproxy for more information. It needs to run for the proxies to work.
- **nova-novncproxy daemon** - Proxy for accessing running virtual machines via a VNC connection. Supports browser-based HTML5 clients.
- **nova-spicehtml5proxy daemon** - Proxy for accessing running virtual machines via a SPICE connection. Supports browser-based HTML5 clients.
- **nova-xvpngproxy daemon** - Proxy for accessing running virtual machines via a VNC connection. Supports an OpenStack-specific Java client.
- **nova-cert daemon** - Manages x509 certificates.
- **euca2ools client** - A set of command-line interpreter commands to manage the cloud resources via Amazon EC2 interface.
- **nova client** - A command-line client for the end user.
- **Message Queue** - Passes messages between nova daemons. Provided by the Advanced Message Queuing Protocol (AMQP). It is usually implemented by RabbitMQ.
- **SQL Database** - Stores most of the build-time and run-time states of the infrastructure, including information about available virtual machine types, virtual machines in use, available networks, and information about projects. OpenStack Compute supports SQLite3 for test and development work, MySQL and MariaDB, and PostgreSQL.

[6]

3.4.4 OpenStack Networking Service

The OpenStack Networking service enables users to create virtual networks and attach the devices managed by other OpenStack services to them. It provides an API that lets users to configure and manage network connectivity and addressing. The network services include L3 forwarding, NAT, load balancing, firewalls, and VPN.

Networking provides the end users the ability to create virtual networks, subnets, routers, and firewalls. All of these will be explained later in this section.

It also supports variety of plug-ins which can enable interoperability with several commercial and open source technologies. This plug-in architecture provides flexibility when designing custom OpenStack architecture and deploying it.

Networking

Before using or deploying the OpenStack Networking service, the following general facts about networking [7] should be understood:

- **Ethernet** is a networking protocol that is being used by most wired network interface cards. In the OSI model of networking, Ethernet operates on the second layer, also referred to as layer 2, L2, link layer, or data link layer. Every host has unique identification called Media Access Control (MAC) address.

Ethernet can be conceptually think of as a single bus, to which each of the network host is connected. However, modern networks use devices called switches, and every device is connected directly to them.

The OpenStack Dashboard uses this simple model to visualise the network topology to the end user. This ethernet network is sometimes referred to as a layer 2 segment.

- **VLAN** is a networking technology that creates separate virtual network on a single switch in a way, that devices connected to the same switch can not each other's traffic, if they are on different VLANs.

OpenStack uses VLANs to isolate the traffic of different tenants, even when their instances run on a single host.

Each VLAN has a numerical ID between 1 and 4095. For example, a VLAN with an id of 15 will be referred to as VLAN 15.

- **Address Resolution Protocol (ARP)** - As pointed out above, network devices use MAC addresses to be identified. However, TCP/IP applications use IP addresses as an identifiers. The Address Resolution Protocol (ARP) bridges the gap by translating IP addresses into MAC addresses.
- **Dynamic Host Configuration Protocol (DHCP)** dynamically assigns IP addresses to network hosts. These hosts are called DHCP clients.
- **TCP, UDP, and ICMP** - Software applications communicating over an IP network use another protocols above IP. In the OSI model of networking, they use the fourth layer, which is also referred to as layer 4, or transport layer. There are three main protocols:

- *Transmission Control Protocol (TCP)* is the most commonly used layer 4 protocol. It is a connection-oriented protocol. Delivery of packets via this protocol is guaranteed.
- *User Datagram Protocol (UDP)* is mostly used to transfer real-time information like voice or video. It is a connectionless and unreliable protocol which means that the delivery of packets via this protocol is not guaranteed.
- *Control Message Protocol (ICMP)* is used for sending control messages.
- **Switch** is device that allow packets to travel from one node to another. They connect hosts belonging to the same layer-2 network. Switches forward the traffic based on the destination Ethernet address in the packet header.
- **Router** is device that allows communication between nodes on different layer-3 networks. They route the traffic based on the destination IP address in the packet header.
- **Firewall** is device that restricts traffic to and from hosts on a network by special rules defined on the device. They are supposed to protect hosts from unauthorised access and attacks.
- **Load Balancer** is device that allows even distribution of traffic across several hosts. They are supposed to avoid overload of a single host. They also prevent a single point of failure as they enable the traffic to be processed by multiple hosts.

Networking Concepts in OpenStack

OpenStack uses the following concepts to enable users to create their own virtual network infrastructure:

- **Network** is an isolated L2 segment. There are two types of network:
 - *Tenant Networks* are managed by the end user and are used within their projects. These networks are fully isolated from other projects.
 - *Provider Networks* are managed by the OpenStack administrator and map to the existing physical network in the datacenter. These networks are mainly used to enable external connectivity to the virtual machines running in the cloud. Each project should have at least one public provider network.
- **Subnet** is a block of IP addresses and associated configuration state. They are used to allocate IP addresses to ports that are created on a network.
- **Port** - Not to be confused with TCP or UDP port. In the OpenStack terminology, port is a connection point for attaching a single device (such as NIC of a virtual machine) to a virtual network. It also describes the configuration like MAC address and IP address.
- **Security Groups** enable users to define their own firewall rules in groups. They can control traffic in both direction (called ingress and egress). These rules are then applied to a port. A port can be assigned with multiple security groups.

[7]

OpenStack Networking Components

The OpenStack Networking service is composed of the following components [6]:

- **neutron-server** - accepts API requests and routes them to the appropriate Networking plug-in
- **OpenStack Networking plug-ins and agents** - the main logic is implemented by several plug-ins. They plug and unplug ports, create networks, subnets, and provide IP addressing. OpenStack Networking ships with plug-ins and agents for Cisco virtual and physical switches, NEC OpenFlow products, Open vSwitch, Linux bridging, and the VMware NSX product. The architecture in this thesis will use Linux bridging.
- **Messaging queue** - routes the information between neutron-server and other agents. It also act as a database for particular plug-ins. The architecture in this thesis will use centralised RabbitMQ service running on a controller node.

The architecture in this thesis will use the following plug-ins and agents for networking:

- **Modular Layer 2 plug-in** - builds layer-2 (bridging and switching) virtual networking infrastructure for virtual machines. IT uses the Linux bridge mechanism.
- **Linux bridge agent** - builds the networking infrastructure for virtual machines and also manages VXLAN tunnels for private networks and security groups.
- **Layer 3 agent** - is the most commonly used agent that provides layer-3 (routing) services for virtual networks.
- **DHCP agent** - provides DHCP services for virtual networks.

3.4.5 OpenStack Block Storage Service

The OpenStack Block Storage service adds persistent block storage to a virtual machine instances and provides an infrastructure for managing volumes. This service is similar to the Amazon EC2 Elastic Block Storage (EBS) offering.

The service supports various backends which can be consumed using a Block Storage driver. An example of supported drivers that are available can be NAS/SAN, NFS, iSCSI, Ceph, and more. A usage of multi-backend configuration is also supported. [6]

To install the OpenStack Block storage service, it is important to understand a number of concepts, because there are several choices of deployment. [15] Apart from choosing the right storage backend, it mostly depends on the final architecture - which can be single node or multi-node. For example, in this thesis I use a multi-node architecture with separate block storage node, and will use a local LVM storage as a backend.

Logical Volume Management (LVM)

Because this thesis will use an LVM as a block storage provider, it is important to understand the basic concepts of LVM.

LVM creates a layer of abstraction over physical storage, which allows you to create multiple logical storage volumes. From an application view, these logical volumes are the same as traditional disk partitions. The hardware storage configuration is also hidden from the

software, so it can be resized and moved without stopping services or unmounting filesystems. And in an opposite way, the logical volumes can be also resized without changing the underlying physical storage. This solution provides much more flexibility than using traditional partitions directly. [8]




Figure 3.3: Logical volume architecture [Source: [8]]

As shown on the picture 3.3, LVM consists of three basic layers:

- **Physical Volume (PV)** is the underlying physical block storage device, which can be a partition or the whole disk. To use the device for LVM, it needs to be initialised as a Physical Volume (PV).
- **Volume Group (VG)** is a combination of several physical volumes. However, volume group can also consist of a single physical volume. This layer provides a pool of disk space used to create logical volumes (LV) in the same way disks are divided into partitions.
- **Logical Volume (LV)** is the volume used by filesystems and applications. In LVM, a volume group is divided into several logical volumes.

Block Storage Components

The OpenStack Block Storage service consists of the following components [6]:

- **cinder-api** - Accepts API requests and routes them to the cinder-volume process for action.
- **cinder-volume** - This process manages the read and write requests sent to the Block Storage service. It interacts with the cinder-scheduler and cinder-backup processes and the storage providers, using respective drivers.

- **cinder-scheduler daemon** - In multi-node deployments, it selects the optimal storage node on which the volume will be created. This process is similar to the nova-scheduler.
- **cinder-backup daemon** - Provides the ability to back up volumes. It can interact with multiple backup storage providers using drivers, in a similar way as the cinder-volume process.
- **Messaging queue** - Routes information between the processes within the OpenStack Block Storage service. In this thesis, I use a centralised RabbitMQ service, running on the controller node.

3.4.6 OpenStack Dashboard Service

The OpenStack Dashboard service provides you a way to manage the OpenStack resources and services. It is used by administrators and the end user. It provides a web-based interface that communicates through the OpenStack APIs. It also allows customising the brand of the dashboard to match the cloud provider's needs. [6]

Handling of User Session Data

The Dashboard service supports several session backends [9] to handle the user session data:

- **Local Memory Cache** - The quickest and easiest session backend, which does not require external dependencies and is the easiest to set up. However, it does not support shared storage across processes or workers, and does not offer data persistency after a process terminates. This is the reason why it is not recommended for production use.
- **Memcached** - An external caching service. Supports shared storage and offers data persistency after process or worker terminates. It is extremely fast and efficient cache backend. It requires the memcached service to be running and accessible and a python memcached module installed.
- **Database** - A database can be also used as a caching backend. It is scalable, it can be highly-available and offers data persistency. However, it is slow in comparison to other caching methods.
- **Cached Database** - This is a hybrid setting using database and caching infrastructure together.
- **Cookies** - Stores data in a cookie in the user's browser. It supports a cryptographic signing to ensure that the data has not been changed during transport. It should be noted, that signing is not the same as encryption, so that the data are still readable by potential attacker.

In this thesis, I will use the Memcached option.

Chapter 4

Existing Methods of Automated OpenStack Deployment

4.1 Packstack

Packstack is a command line installation utility that support deployment of OpenStack on existing server using SSH connection. The installation can be configured interactively or by a configuration file called answer file. [17]

It supports these two basic types of deployment [13]:

- **An all-on-one installation** - All services are installed on a single physical host that would run all controller services and the virtual machines.
- **Multiple nodes** - Using several hosts to run the installations, where there is a single controller node running the controller services, and one or more compute nodes that would run the virtual machines.

However, packstack is not suitable for production deployments. This is mainly because it makes many assumptions about the configuration in order to simplify the installation process. It can not deploy the services in a high availability (HA) mode or using load balancers. It also does not support advanced networking configuration, which might be required by more complex setups. [12]

4.2 OpenStack-Ansible

OpenStack-Ansible is an official OpenStack project, still under development. The goal of this project is to be able to deploy OpenStack cloud in a production environment directly from source code. It is focused on Ubuntu Linux and the OpenStack components are installed into Linux Containers (LXC). [10] [11]

Chapter 5

Implementation Design

5.1 Design of the Physical Architecture

Before I can write the Ansible playbook to deploy the OpenStack cloud, I need to know exactly what I need to deploy. I have designed a basic physical architecture that I will use, and this chapter will describe it.

5.1.1 Selecting OpenStack Services

The first decision when designing an OpenStack cloud deployment is about the services we need to have running.

For this thesis, I have chosen these five core services:

- OpenStack Identity (Keystone)
- OpenStack Image (Glance)
- OpenStack Compute (Nova)
- OpenStack Networking (Neutron)
- OpenStack Block Storage (Cinder)
- OpenStack Dashboard (Horizon)

OpenStack consists of several components. I have chosen these six as they are the minimum to have a cloud deployment where I can create virtual machines, attach them to networks, use routers and other networking devices, attach a persistent block storage, and also manage the cloud by a web interface. Implementing more services would be out of the scope of the Thesis.

The services mentioned above will require a database to store their states, and a message broker for some processes to communicate. Specifically, these two services will be needed:

- SQL Database
- AMQP message broker

5.1.2 Selecting Specific Implementations

The second decision is about the specific implementation of each technology. The OpenStack cloud comes in several releases. The newest release at the time of writing this Thesis is OpenStack Liberty. I will use this release as I want to develop for the newest stable version available.

I also need to choose an operating system platform that will run on the physical host machines. Choosing the right Linux distribution is important as there might be big differences in terms of available packages, stability, the lifecycle, and an option of commercial support. For this Thesis, I have chosen a CentOS 7.2. CentOS is free and open source operating system based on the Red Hat Enterprise Linux platform. It uses the RPM packaging model and the OpenStack Liberty packages are included in the repositories. The system is stable, has a long lifecycle and potential migration to to a commercially supported environment would be easy, as it is based on the Red Hat Enterprise Linux platform. I have also chose this particular distribution as I know the environment well enough to perform needed configurations and debugging.

OpenStack also needs an SQL database to store the state of the services, and it also might be used as a storage for the Identity service. In this deployment, it will be used in this way. I have chosen MariaDB because it is open source and widely used SQL database. It also comes with the CentOS distribution by default.

The next component is an AMQP broker. It is used by processes within the OpenStack services for communication, and it also might be used as a short-term storage by some networking plugins. For this particular deployment, I will use RabbitMQ because it is an open source solution and it is widely used in production environments. It is also well-supported by the OpenStack itself.

The last decision will be about storage backend for the OpenStack Block Storage service. I will use a local disk and an LVM technology. LVM is a robust and stable open source solution for storage that runs on commodity hardware, which is important especially for this testing environment.

5.1.3 Physical architecture

With the knowledge of the specific components and their implementations, I can design a physical architecture of the deployment I will be using in this Thesis. The physical architecture can be seen on picture 5.1. It consists of three hosts, referred to as nodes. The description of the nodes is as follows:

Controller node

The controller node is the central point of this OpenStack architecture. It will control and command the whole cloud, and it will also store the state of the cloud. It will host both of the external services: MariaDB and RabbitMQ. There will be also the controller parts of the Compute and Block Storage services. To make this setup easier, it will also host the whole Networking service. However, in large production deployments, the Networking service might be installed on a separate node.

The communication between all OpenStack services will go through this node, as it also hosts the Identity service that will act as an API. Also, the Dashboard service will be deployed to this nodes. The administrators and the end users will be accessing the cloud via the controller node.

Compute node

The compute node will host the Nova Compute service and will act as a hypervisor. All instances in the cloud will run on this node. It is designed in a way it can scale horizontally in need of extra resources. It means that there can be several compute nodes deployed at the same time. If this happens, the Nova scheduler service deployed on the controller node will automatically choose the most suitable compute node for all new virtual machines.

Storage node

The storage node will provide persistent storage to the virtual machines. It has been separated so the storage can scale independently of the compute resources and vice versa. The storage node is designed in a way that it can scale horizontally, in the same way as the compute node which means that there can be more storage nodes deployed at the same time.

Networks

I have used three different networks in my architecture and each one have different purpose in the whole system. Their descriptions are as follows:

- **Utility Network** - This network is used the datacenter administrators to maintain the physical hosts. In this example it will also be used by Ansible to deploy the OpenStack and all other services.
- **Management Network** - This network is used by the OpenStack services to communicate with each other. It will also provide connectivity for the services to the MySQL database and the RabbitMQ Message Queue. All Keystone endpoints based on hostnames will also use this network. It will be also used to connect storage from Storage1 host to virtual machines running on the Compute1 host.
- **VM Network** - This network is used to provide outside connectivity to the virtual machines running on the compute node.

There is a separate Utility and Management network. This is because the IP addresses on the Management network are setted up during the deployment. These two networks can be joined together if the IP addresses are properly set before and will not change during deployment. This information might be important when testing on physical nodes with limited number of network connectors.

It is also important to note that the eth2 interfaces on Controller and Compute1 hosts will use special configuration without an IP address and will be attached to virtual network bridges. Their specific configuration will be mentioned later in the text.

5.2 Design of the Ansible Playbook

The Ansible playbook will be designed in a way that it will deploy the physical architecture designed above. To achieve this, it will require minimal configuration, such as IP addresses, hostnames, and passwords - so the level of abstraction will be high enough for the administrator to deploy OpenStack without deep knowledge about it as they will only need to know the IP addresses, hostnames, and passwords.

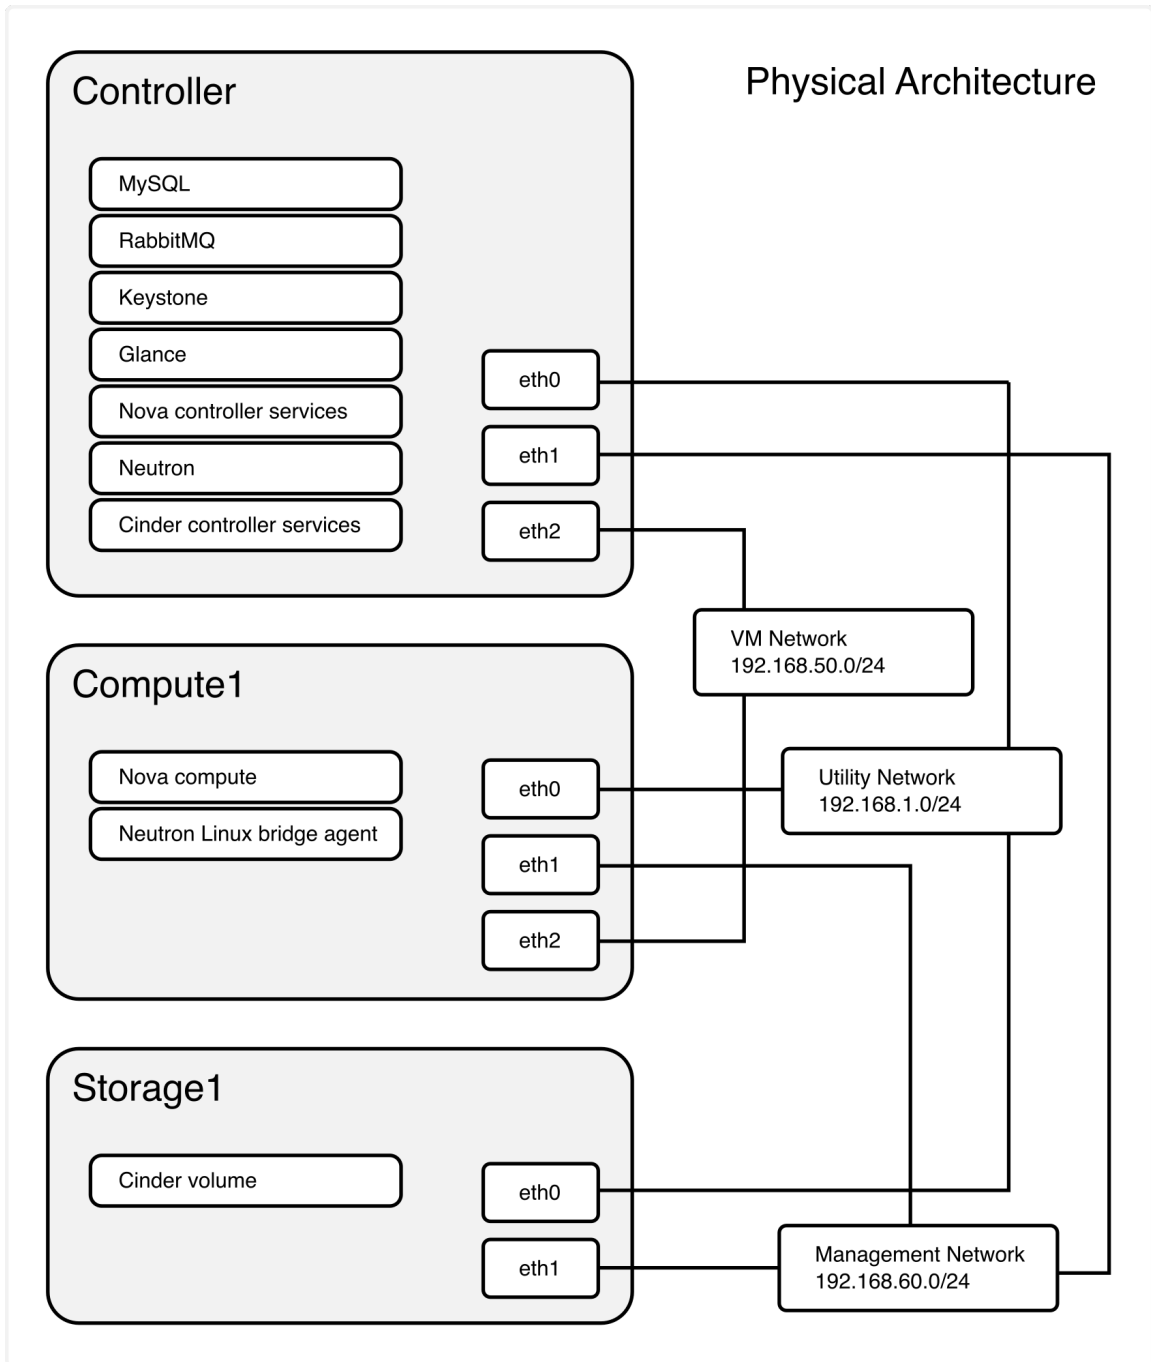


Figure 5.1: OpenStack physical architecture design

However, the real benefit of this playbook will be in the reusability for different architectures. In addition of configuring the basic options mentioned above, it will be possible to change the script itself to match other architectures. To make this as easy as possible, I will use the minimal level of abstraction when writing the playbook. I will only use the most basic operations such as installing specific packages using yum, changing the OpenStack configuration files, and restarting systemd services. This minimal abstraction will allow to make much more customisations of the playbook than it is possible to do with Packstack.

However, it will require much deeper understanding of the OpenStack architecture which might make it harder to use for some people.

5.2.1 Ansible Roles

The services will be installed on the target hosts using Ansible roles. Roles offer flexibility in a way, that each role can be installed separately on a given host, or, if needed, on multiple hosts. These roles can be then reused to deploy a different physical architectures.

All roles use handlers to restart services in case of a configuration change. All roles are also designed to be idempotent.

There is a role for each OpenStack service. Some services, such as OpenStack Compute and OpenStack Block Storage, will be divided into two roles: the controller part, installed on the controller node, and the functional part, installed on respective nodes.

There will be two other roles to install the database and the AMQP message bus.

Following is the description of each individual role.

Role `sql-database`

The role `sql-database` will deploy a MariaDB database on the target host. It will set the root password accordingly and will also remove the anonymous user as well as the test database.

The following packages will be installed:

- `mariadb`
- `mariadb-server`
- `MySQL-python`

And the following MariaDB service will be enabled and started:

- `mariadb`

Role `rabbit`

The `rabbit` role will deploy a RabbitMQ message broker on the target host. It will also create a new user needed for the OpenStack deployment.

The following package will be installed:

- `rabbitmq-server`

And the following RabbitMQ service will be enabled and started:

- `rabbitmq-server`

Roles `controller-basic` and `compute-basic`

The roles `controller-basic` and `compute-basic` enable the OpenStack Liberty repository for CentOS and install the SELinux package. These roles are the first that should be run before any other openstack roles.

These roles will install these two packages:

- `python-openstackclient`
- `openstack-selinux`

Role keystone

The keystone role will install the OpenStack Identity service, codename Keystone, on the target host. It requires an SQL database to be running on the network. The design used in this thesis will run the database on the controller node.

The role will create a database for the Keystone service called keystone.

This role will install the following packages:

- `openstack-keystone`
- `httpd`
- `mod_wsgi`
- `memcached`
- `python-memcached`
- `python-keystoneclient`

And it will enable and start the httpd service on the target host.

Role glance

The glance role will install the OpenStack Image service, codename Glance, on the target host. It requires an SQL database, and the OpenStack Identity service to be running on the network. The example in this thesis will run both on the controller node.

The role will create a database for the Glance service called glance, and registers the Glance service and creates endpoints in the Keystone service.

Glance supports several backends for storing images. This role uses local filesystem to do so. Metadata about the images will be stored in the SQL database running on the controller node.

This role will install the following packages:

- `openstack-glance`
- `python-glance`
- `python-glanceclient`

And the two following services will be enabled and started:

- `openstack-glance-api`
- `openstack-glance-registry`

Role nova-controller

The nova-controller role will install some parts of the OpenStack Compute service, code-name Nova, on the target host. It requires an SQL database, RabbitMQ message bus, and the OpenStack Identity service to be running on the network. The example in this thesis will run all these services on the controller node.

The role will create a database for the Nova service called nova, and registers the Nova service and creates endpoints in the Keystone service.

Besides the standard configurations such as setting hostname, configuration of database, and message bus access, the Nova service will be configured to use the OpenStack Networking (as opposed to legacy Nova networking) with the Linux bridge driver.

This role will install the following packages:

- `openstack-nova-api`
- `openstack-nova-cert`
- `openstack-nova-conductor`
- `openstack-nova-console`
- `openstack-nova-novncproxy`
- `openstack-nova-scheduler`
- `python-novaclient`

And the following services will be enabled and started:

- `openstack-nova-api`
- `openstack-nova-cert`
- `openstack-nova-consoleauth`
- `openstack-nova-scheduler`
- `openstack-nova-conductor`
- `openstack-nova-novncproxy`

Role `nova-compute`

The `nova-compute` role will install the `nova-compute` process of the OpenStack Compute service, codename Nova, on the target host. It requires the RabbitMQ message bus to be running on the network. The example in this thesis will run it on the controller node.

The role will configure an access to the necessary Nova processes installed by the `nova-controller` role, and will also configure Nova to use the OpenStack Networking with the Linux bridge driver.

It will also configure the hypervisor. This role uses the `libvirt` provider with `QEMU`. This is because `VirtualBox`, which is used for testing the deployment, does not support nested virtualisation, so `KVM` can not be used.

This role will install these two packages:

- `openstack-nova-compute`
- `sysfsutils`

And will start and enable the two following services:

- `libvirtd`
- `openstack-nova-compute`

Role neutron-controller

The neutron-controller role will install the OpenStack Networking service, codename Neutron, on the target host. It requires an SQL database, RabbitMQ message bus, and the OpenStack Identity service to be running on the network. The example in this thesis will run all these services on the controller node. This role also requires the nova-compute role to be run on the same host before.

The role will create a database for the Neutron service called neutron, and registers the Neutron service and creates endpoints in the Keystone service.

The OpenStack Networking service uses plug-ins and agents for the actual networking functionality. This role will use:

- Modular Layer 2 (ML2) plugin
- Linux bridge agent
- Layer-3 agent
- DHCP agent

This setup will allow to create tenant networks as well as public provider networks.

The ML2 plugin will be configured to use the Linux bridge technology and VXLAN to create the tenant networks. The Linux bridge agent will be configured to use VXLAN for the tenant networks and iptables firewall driver to manage security groups. The layer-3 agent will be configured to use the Linux bridge driver and to support multiple external networks.

The DHCP agent will be also configured to use the Linux bridge driver and the MTU is set to 1450 bytes. This is because the VXLAN includes additional packet header and virtual machines running in the cloud use the default MTU of 1500 bytes. Using this settings, the virtual machines will use the smaller MTU, which would allow space for the additional header.

Finally, the Nova service will be configured to use the OpenStack Networking service, installed by this role.

This role will install the following packages:

- openstack-neutron
- openstack-neutron-ml2
- openstack-neutron-linuxbridge
- python-neutronclient
- ebtables
- ipset

It will restart this service:

- openstack-nova-api

And the following services will be started and enabled:

- neutron-server

- `neutron-linuxbridge-agent`
- `neutron-dhcp-agent`
- `neutron-metadata-agent`
- `neutron-l3-agent`

Role `neutron-compute`

The `neutron-compute` role will install the Linux bridge agent of the OpenStack Networking service, codename Neutron, on the target host. It requires the RabbitMQ message bus to be running on the network. The example in this thesis will run it on the controller node. It also requires the role `nova-compute` to be run on the same host before.

This role will configure the Linux bridge agent to use the correct network interface as a public interface, enables VXLAN for the tenant networks, and configures iptables as a firewall driver to manage security groups.

It will also configure the Nova service to use the OpenStack Networking.

This role will install the following packages:

- `openstack-neutron`
- `openstack-neutron-linuxbridge`
- `ebtables`
- `ipset`

It will restart this Nova service:

- `openstack-nova-compute`

And also the Linux bridge agent service will be started and enabled:

- `neutron-linuxbridge-agent`

Role `cinder-controller`

The `cinder-controller` role will install the OpenStack Block Storage service, codename Cinder, on the target host. It requires an SQL database, RabbitMQ message bus, and the OpenStack Identity service to be running on the network. The example in this thesis will run all these services on the controller node.

The role will create a database for the Cinder service called `cinder`, and registers two services called `cinder` and `cinderv2`, and creates endpoints for each service. Two services are created, because the OpenStack Block Storage service currently supports two versions of API. This role will use both for better compatibility.

The following packages will be installed:

- `openstack-cinder`
- `python-cinderclient`

The `openstack-nova-api` service needs to be restarted first, and then the following services will be enabled and started:

- `openstack-cinder-api`
- `openstack-cinder-scheduler`

Role `cinder-storage`

This `cinder-storage` role will install the storage part of the OpenStack Block Storage service, codename Cinder, on the target host. It requires the RabbitMQ message bus to be running on the network. The example in this thesis will run it on the controller node.

At first, the role makes sure that there is a partition available for the storage, and creates an LVM volume group. Then it configures the Block Storage service to use this volume group to create individual volumes of persistent storage for the virtual machines.

The following packages will be installed:

- `lvm2`
- `openstack-cinder`
- `targetcli`
- `python-oslo-policy`

And the following services will be started and enabled:

- `openstack-cinder-volume`
- `target`

Role `dashboard`

The `dashboard` role will install the OpenStack Dashboard service, codename Horizon, on the target host. The Dashboard service installed by this role uses the Apache `httpd` web server.

It will install the following packages:

- `openstack-dashboard`
- `httpd`
- `memcached`

And the following services will be started and enabled:

- `httpd`
- `memcached`

5.2.2 Applying Roles to the Hosts

The main Ansible playbook will apply the roles designed above to particular hosts. See the table 5.1 for reference about what roles will be applied to which hosts.

Each role will require several parameters such as IP addresses, hostnames, or passwords to be set. This can be also done in the main playbook.

Host Name	Applied Roles
Controller node	controller-basic rabbit sql-database dashboard glance keystone neutron-controller nova-controller cinder-controller
Compute node	compute-basic neutron-compute nova-compute
Storage node	storage-basic cinder-storage

Table 5.1: Roles applied to the particular hosts

Chapter 6

Implementation and Testing

6.1 Third-party Modules in the Playbook

The implementation requires two third-party modules for managing Keystone endpoints and services. Both modules have been created by Davide Guerri ¹, licensed under Apache License, Version 2.0, and are included as part of the playbook.

6.2 Testing of the Deployment

I have prepared the following setup to for testing the environ

6.2.1 Description of the Host Environment

The reference testing environment consists of three virtual machines with the following specification:

Controller

- CPUs: 2
- RAM: 2048 MB
- Disk: 30 GB

Compute

- CPUs: 4
- RAM: 6144 MB
- Disk: 30 GB

Storage

- CPUs: 1
- RAM: 1536 MB

¹<davide.guerri@gmail.com>

- Disk1: 20 GB
- Disk2: 60 GB

These virtual machines were running on a laptop with the following configuration:

- CPU: 2.8 GHz Intel i7 4980HQ
- RAM: 16 GB
- SSD: PCIe 3.0 x4 8.0 GT/s (25.6 Gbit/s)

6.2.2 Using Ansible with Vagrant

Vagrant is a tool that manages virtual machines for development environment. It uses private key authentication and generates its own keys for the virtual machines.

First, Vagrant needs to be configured to use a single private key to authenticate to all three virtual machines. To achieve this, the following line needs to be put in the `Vagrantfile`:

```
config.ssh.insert_key = false
```

Vagrant routes SSH ports of the guest virtual machines to the localhost and uses different port for each. A command `vagrant ssh-config` will show the ports of each virtual machine. These ports need to be used in the `hosts` file. An example might look like this:

```
controller ansible_ssh_host=127.0.0.1 ansible_ssh_port=2222
compute1 ansible_ssh_host=127.0.0.1 ansible_ssh_port=2200
compute1 ansible_ssh_host=127.0.0.1 ansible_ssh_port=2201
```

The last step is to configure Ansible to use the correct user name and private key. This can be done by creating file called `ansible.cfg` with the following content:

```
[defaults]
hostfile = hosts
remote_user = vagrant
private_key_file = ~/.vagrant.d/insecure_private_key
host_key_checking = False
```

6.2.3 Running the Deployment

To deploy the OpenStack cloud in the testing environment using the Ansible playbook, you need to provision the virtual machines and run the playbook. The virtual machines can be provisioned using Vagrant by issuing the following command:

```
$ vagrant up
```

This will create three virtual machines matching the physical architecture. When the virtual machines are ready, you can run the Ansible playbook by issuing the following command:

```
$ ansible-playbook deploy-openstack.yml
```

This will deploy the OpenStack cloud on the virtual machines created before. The playbook assumes that you have an internet connection available to all three virtual machines. When running the deployment, the installation took approximately 7 minutes.

6.2.4 Testing the Deployment

When the deployment is finished, the local instance of the OpenStack cloud can be accessed using the Dashboard via web browser. With the default configuration provided with this thesis, the Dashboard will be accessible on `http://10.0.0.11/dashboard`. To log in, enter the following credentials:

- Domain: `default`
- User: `admin`
- Password: `redhat`

Checking Running Services

The Dashboard is able to show all OpenStack services running in the current deployment. A screenshot of the running services is shown as picture 6.1). There is also a more detailed information about distributed services running on two nodes. These are the Compute Services (shown on picture 6.2), the block storage services (shown on picture 6.3), and the network agents (shown on picture 6.4).

All necessary services are running which means that the deployment was successful.

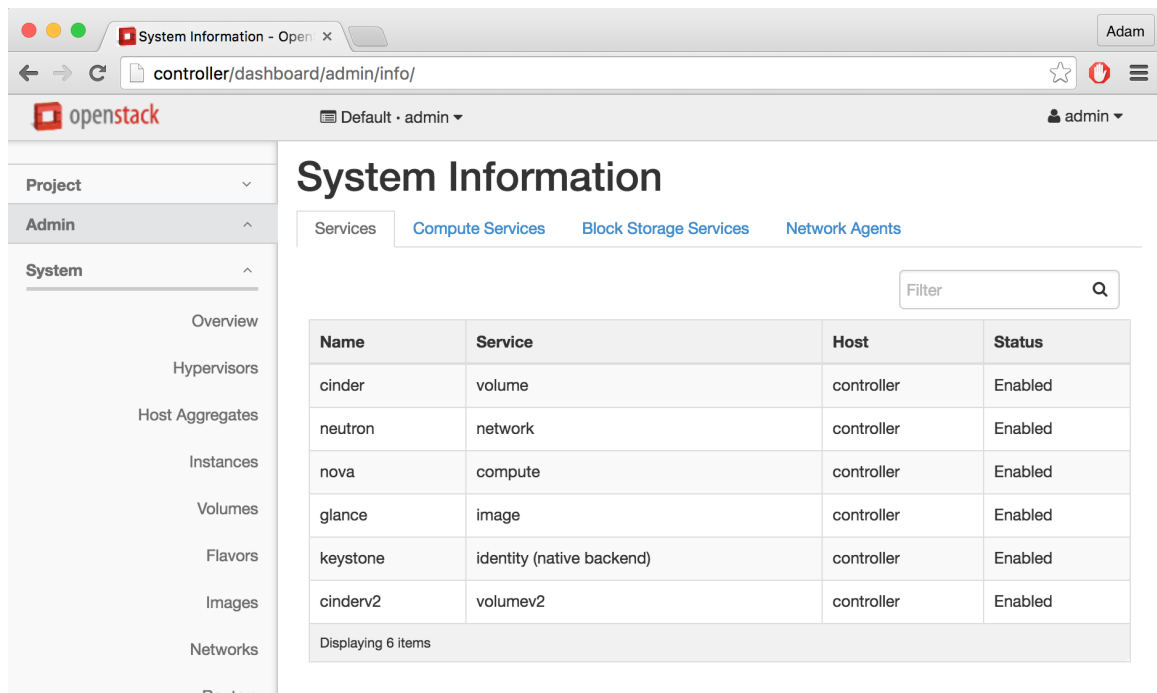


Figure 6.1: The OpenStack Dashboard shows the OpenStack services running

The screenshot shows the OpenStack Dashboard interface. The top navigation bar includes the OpenStack logo, the current project 'Default', and the user 'admin'. The left sidebar contains a navigation menu with 'Project', 'Admin', and 'System' sections. The main content area is titled 'System Information' and has tabs for 'Services', 'Compute Services', 'Block Storage Services', and 'Network Agents'. The 'Compute Services' tab is active, displaying a table of running services. A search filter is present above the table.

Name	Host	Zone	Status	State	Last Updated
nova-cert	controller	internal	Enabled	Up	0 minutes
nova-consoleauth	controller	internal	Enabled	Up	0 minutes
nova-scheduler	controller	internal	Enabled	Up	0 minutes
nova-conductor	controller	internal	Enabled	Up	0 minutes
nova-compute	compute1	nova	Enabled	Up	0 minutes

Displaying 5 items

Figure 6.2: The OpenStack Dashboard shows the compute services running

The screenshot shows the OpenStack Dashboard interface. The top navigation bar includes the OpenStack logo, the current project 'Default', and the user 'admin'. The left sidebar contains a navigation menu with 'Project', 'Admin', and 'System' sections. The main content area is titled 'System Information' and has tabs for 'Services', 'Compute Services', 'Block Storage Services', and 'Network Agents'. The 'Block Storage Services' tab is active, displaying a table of running services. A search filter is present above the table.

Name	Host	Zone	Status	State	Last Updated
cinder-scheduler	controller	nova	Enabled	Up	0 minutes
cinder-volume	storage1@lvm	nova	Enabled	Up	0 minutes

Displaying 2 items

Figure 6.3: The OpenStack Dashboard shows the block storage services running

The screenshot shows the OpenStack Dashboard interface. The top navigation bar includes the OpenStack logo, the current project 'Default', and the user 'admin'. The left sidebar contains a navigation menu with 'Project', 'Admin', and 'System' sections. The main content area is titled 'System Information' and has tabs for 'Services', 'Compute Services', 'Block Storage Services', and 'Network Agents'. The 'Network Agents' tab is active, displaying a table of running agents. A search filter is present above the table.

Type	Name	Host	Status	State	Last Updated
L3 agent	neutron-l3-agent	controller	Enabled	Up	0 minutes
Linux bridge agent	neutron-linuxbridge-agent	controller	Enabled	Up	0 minutes
Linux bridge agent	neutron-linuxbridge-agent	compute1	Enabled	Up	0 minutes
Metadata agent	neutron-metadata-agent	controller	Enabled	Up	0 minutes
DHCP agent	neutron-dhcp-agent	controller	Enabled	Up	0 minutes

Displaying 5 items

Figure 6.4: The OpenStack Dashboard shows the network agents running

Testing the Functionality

To test the deployment functionality, I have created six virtual machines, two tenant networks, one router, and one provider network. The picture 6.5 is a screenshot showing the whole environment running in the OpenStack Dashboard.

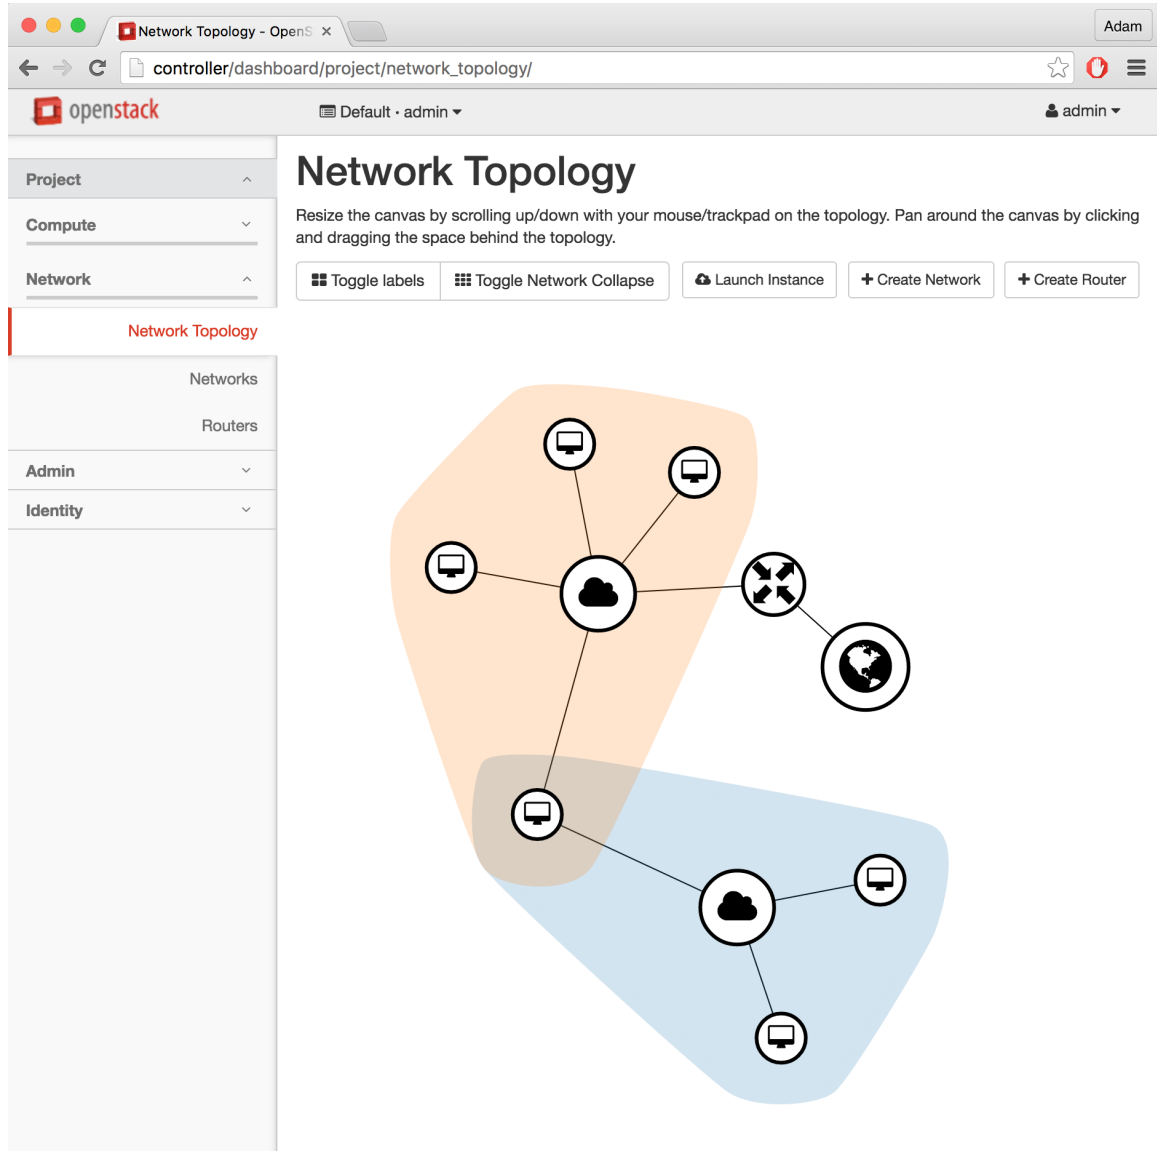


Figure 6.5: The OpenStack Dashboard shows the testing environment

Chapter 7

Conclusion

In this thesis, I have designed simple OpenStack architecture able to run virtual machines, create tenant networks using software-defined networking (SDN), and use persistent storage. This runs in a multi-tenant environment and is managed by a web interface.

I have also designed and implemented an Ansible playbook which can automatically deploy the architecture on multiple physical hosts. This playbook has been designed for an easy deployment of the reference architecture, and, at the same time, to be reusable with modifications for deploying different production environments. The reusability is also simplified by using Ansible roles. There is a separate role for each functionality and they can be deployed on separate hosts.

In comparison to another automatic deployment system called Packstack, it offers more flexibility in terms of architecture changes. This is because I have not used the high level of abstraction Packstack uses. On the other hand, it means that customising the playbook will require much deeper knowledge of the OpenStack cloud compared to using Packstack. While this is an advantage for experienced OpenStack administrators, it can be also disadvantage for beginners.

Compared to the official OpenStack-Ansible project, the playbook developed in this thesis installs the OpenStack cloud using RPM packages that can be signed and certified before use.

7.1 Future Development

Some roles will be offered to the Fedora Infrastructure team¹ for their own OpenStack infrastructure which is currently installed by Packstack and customized by Ansible. It might require further development and changes to match their architecture, and it will also require extensive testing before using in production. The project has been also published on GitHub².

¹<https://fedoraproject.org/wiki/Infrastructure>

²<https://github.com/asamalik/openstack-ansible>

Bibliography

- [1] How is scalability achieved?
<https://hadoop4usa.wordpress.com/2012/04/13/scale-out-up/>, 2012 [quoted 2016-04-5].
- [2] *Introduction to OpenStack [online]*. <http://docs.openstack.org/security-guide/introduction/introduction-to-openstack.html>, 2015-10-22 [quoted 2015-10-22].
- [3] *Conceptual architecture [online]*. http://docs.openstack.org/liberty/install-guide-obs/common/get_started_conceptual_architecture.html, 2015-10-25 [quoted 2015-10-25].
- [4] *How Ansible Works [online]*. <https://www.ansible.com/how-ansible-works>, 2015 [quoted 2015-11-14].
- [5] *OpenStack Administrator Guide [online]*.
<http://docs.openstack.org/admin-guide/>, 2015 [quoted 2015-11-27].
- [6] *OpenStack Installation Guide for Red Hat Enterprise Linux and CentOS [online]*.
<http://docs.openstack.org/liberty/install-guide-rdo/>, 2015 [quoted 2015-11-27].
- [7] *OpenStack Networking Guide [online]*.
<http://docs.openstack.org/liberty/networking-guide/>, 2015 [quoted 2015-11-27].
- [8] *Red Hat Enterprise Linux 6 Logical Volume Manager Administration [online]*.
https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/pdf/Logical_Volume_Manager_Administration/Red_Hat_Enterprise_Linux-6-Logical_Volume_Manager_Administration-en-US.pdf, 2015 [quoted 2015-11-27].
- [9] *Deploying Horizon [online]*.
<http://docs.openstack.org/developer/horizon/topics/deployment.html>, 2015 [quoted 2015-11-28].
- [10] *About OpenStack-Ansible [online]*. <http://docs.openstack.org/developer/openstack-ansible/install-guide/overview-osa.html>, 2015 [quoted 2015-12-09].
- [11] *Ansible playbooks for deploying OpenStack [online]*.
<https://github.com/openstack/openstack-ansible>, 2015 [quoted 2015-12-09].

- [12] *Deploying OpenStack: Proof-of-Concept Environments (PackStack) [online]*.
https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/5/html/Getting_Started_Guide/The_PackStack_Deployment_Utility1.html, 2015 [quoted 2015-12-09].
- [13] *Packstack quickstart: Proof of concept for single node [online]*.
<https://www.rdoproject.org/install/quickstart/>, 2015 [quoted 2015-12-09].
- [14] Ansible, I.: *Ansible Documentation [online]*.
<http://docs.ansible.com/ansible/index.html>, 2015 [quoted 2015-11-14].
- [15] Fifield, T.; Fleming, D.; Gentle, A.; aj.: *OpenStack Operations Guide*. O'Reilly Media, 2014, ISBN 978-1-4919-4694-7.
- [16] Hochstein, L.: *Ansible: Up and Running*. O'Reilly Media, 2014, ISBN 978-1-4919-1532-5.
- [17] Taylor F., V. A., Mahroua R.: *Red Hat Enterprise Linux OpenStack Platform 6.0 CL210, Red Hat OpenStack Administration*. Red Hat, 2015, edition 2 20150415.

Appendices

List of Appendices

A	CD Contents	47
B	Manual	48
B.1	Using Roles with Variables	48
B.2	Variables of Individual Roles	49
B.2.1	Role <code>keystone</code>	49
B.2.2	Role <code>glance</code>	49
B.2.3	Role <code>nova-controller</code>	50
B.2.4	Role <code>nova-compute</code>	50
B.2.5	Role <code>neutron-controller</code>	51
B.2.6	Role <code>neutron-compute</code>	52
B.2.7	Role <code>cinder-controller</code>	52
B.2.8	Role <code>cinder-storage</code>	53
B.2.9	Role <code>dashboard</code>	53
B.2.10	Role <code>rabbit</code>	53
B.2.11	Role <code>sql-database</code>	54

Appendix A

CD Contents

The CD attached includes the Ansible playbook with default configuration, a `README` file with instructions, and a `Vagrantfile` to provision testing environment.

Appendix B

Manual

B.1 Using Roles with Variables

When deploying your own OpenStack cloud, you will need to set your own passwords and configure the right network interfaces to match your infrastructure. All roles designed in this Thesis implement variables which allow you to do the configuration. See the following example to learn how to use roles and variables:

```
- name: compute node
  hosts: compute1
  sudo: yes
  roles:
    - compute-basic
    - nova-compute
    - neutron-compute
  controller_hostname: controller.example.com
  public_interface_name: eth2
```

The example above applies three roles, `compute-basic`, `nova-compute`, and `neutron-compute` to the compute node with hostname `compute1`. It also sets two variables `controller_hostname` and `public_interface_name` to respective values.

See the following section for list and descriptions of all variables used by each role.

B.2 Variables of Individual Roles

This section describes all variables used by each role. All variables are mandatory.

B.2.1 Role keystone

The role `keystone` installs the OpenStack Identity service. Variables used in this role are described in table [B.1](#).

Variable Name	Default Value	Description
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>keystone_db_password</code>	<code>redhat</code>	Database password for the OpenStack Keystone service.
<code>keystone_admin_token</code>	<code>keystone_admin_token</code>	Admin token for the OpenStack Keystone service.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>admin_password</code>	<code>redhat</code>	Password for the default <code>admin</code> user.

Table B.1: Variables of the `keystone` role

B.2.2 Role glance

The role `glance` installs the OpenStack Image service. Variables used in this role are described in table [B.2](#).

Variable Name	Default Value	Description
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>keystone_admin_token</code>	<code>keystone_admin_token</code>	Admin token for the OpenStack Keystone service.
<code>glance_db_password</code>	<code>redhat</code>	Database password for the OpenStack Image service.
<code>glance_password</code>	<code>redhat</code>	Keystone password for the OpenStack Image service.

Table B.2: Variables of the `glance` role

B.2.3 Role nova-controller

The role `nova-controller` installs the controller services of the OpenStack Compute service. Variables used in this role are described in table B.3.

Variable Name	Default Value	Description
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>keystone_admin_token</code>	<code>keystone_admin_token</code>	Admin token for the OpenStack Keystone service.
<code>nova_db_password</code>	<code>redhat</code>	Database password for the OpenStack Compute service.
<code>nova_password</code>	<code>redhat</code>	Keystone password for the OpenStack Compute service.
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.
<code>controller_ip</code>	<code>10.0.0.11</code>	IP address on the interface attached to the Management network.

Table B.3: Variables of the `nova-controller` role

B.2.4 Role nova-compute

The role `nova-compute` installs the `nova-compute` component of the OpenStack Compute service. Variables used in this role are described in table B.4.

Variable Name	Default Value	Description
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.
<code>compute_ip</code>	<code>10.0.0.31</code>	IP address on the interface attached to the Management network.
<code>nova_password</code>	<code>redhat</code>	Keystone password for the OpenStack Compute service.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.

Table B.4: Variables of the `nova-compute` role

B.2.5 Role neutron-controller

The role `neutron-controller` installs the controller services of the OpenStack Networking service. Variables used in this role are described in table B.5.

Variable Name	Default Value	Description
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>keystone_admin_token</code>	<code>keystone_admin_token</code>	Admin token for the OpenStack Keystone service.
<code>neutron_db_password</code>	<code>redhat</code>	Database password for the OpenStack Networking service.
<code>neutron_password</code>	<code>redhat</code>	Keystone password for the OpenStack Networking service.
<code>nova_password</code>	<code>redhat</code>	Keystone password for the OpenStack Compute service.
<code>management_interface_name</code>	<code>eth1</code>	Name of the interface attached to the Management network.
<code>public_interface_name</code>	<code>eth2</code>	Name of the interface attached to the VM network.
<code>overlay_interface_ip</code>	<code>10.0.0.11</code>	IP address on the interface attached to the Management network.
<code>metadata_secret</code>	<code>metadata_secret</code>	A secret for the metadata agent.
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.

Table B.5: Variables of the `neutron-controller` role

B.2.6 Role neutron-compute

The role `neutron-compute` installs the compute services of the OpenStack Networking service. Variables used in this role are described in table B.6.

Variable Name	Default Value	Description
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.
<code>neutron_password</code>	<code>redhat</code>	Keystone password for the OpenStack Networking service.
<code>public_interface_name</code>	<code>eth2</code>	Name of the interface attached to the VM network.
<code>overlay_interface_ip</code>	<code>10.0.0.31</code>	IP address on the interface attached to the Management network.

Table B.6: Variables of the `neutron-compute` role

B.2.7 Role cinder-controller

The role `cinder-controller` installs the controller services of the OpenStack Block Storage service. Variables used in this role are described in table B.7.

Variable Name	Default Value	Description
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>controller_ip</code>	<code>10.0.0.11</code>	IP address on the interface attached to the Management network.
<code>cinder_password</code>	<code>redhat</code>	Keystone password for the OpenStack Block Storage service
<code>cinder_db_password</code>	<code>redhat</code>	Database password for the OpenStack Block Storage service
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.

Table B.7: Variables of the `cinder-controller` role

B.2.8 Role `cinder-storage`

The role `cinder-storage` installs the storage services of the OpenStack Block Storage service. Variables used in this role are described in table B.8.

Variable Name	Default Value	Description
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>cinder_db_password</code>	<code>redhat</code>	Database password for the OpenStack Block Storage service
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.
<code>cinder_password</code>	<code>redhat</code>	Keystone password for the OpenStack Block Storage service
<code>management_ip</code>	<code>10.0.0.41</code>	IP address on the interface attached to the Management network.
<code>cinder_volume_group</code>	<code>cinder-volumes</code>	Name of the volume group used by the OpenStack Block Storage service.
<code>cinder_partition</code>	<code>/dev/sdb</code>	Partition used by the OpenStack Block Storage service.

Table B.8: Variables of the `cinder-storage` role

B.2.9 Role `dashboard`

The role `dashboard` installs the OpenStack Dashboard service. Variables used in this role are described in table B.9.

Variable Name	Default Value	Description
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.

Table B.9: Variables of the `dashboard` role

B.2.10 Role `rabbit`

The role `rabbit` installs the RabbitMQ message bus. Variables used in this role are described in table B.10.

Variable Name	Default Value	Description
<code>rabbit_password</code>	<code>redhat</code>	Password for the AMQP message bus.

Table B.10: Variables of the `rabbit` role

B.2.11 Role `sql-database`

The role `sql-database` installs the MariaDB database. Variables used in this role are described in table [B.11](#).

Variable Name	Default Value	Description
<code>mysql_root_password</code>	<code>redhat</code>	Root password for the MariaDB database.
<code>controller_hostname</code>	<code>controller</code>	Hostname of the controller node.
<code>controller_management_ip</code>	<code>10.0.0.11</code>	IP address on the interface attached to the Management network.

Table B.11: Variables of the `sql-database` role