A Demonstration of 4D Digital Avatar Infrastructure for Access of Complete Patient Information

Project acronym: MyHealthAvatar

Deliverable No. D3.5
Report on Local Cloud Infrastructure

Grant agreement no: 600929

Dissemination Level		
PU	Public	**X**
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

COVER AND CONTROL PAGE OF DOCUMENT	
Project Acronym:	MyHealthAvatar
Project Full Name:	A Demonstration of 4D Digital Avatar Infrastructure for Access of Complete Patient Information
Deliverable No.:	D3.5
Document name:	Report on local cloud infrastructure
Nature (R, P, D, O)[1]	R
Dissemination Level (PU, PP, RE, CO)[2]	PU
Version:	1
Actual Submission Date:	15/10/2014
Editor: Institution: E-Mail:	Xia Zhao University of Bedfordshire xia.zhao@beds.ac.uk

ABSTRACT:
This deliverable reports the investigation on deployment of a local cloud infrastructure for support data storage and computation requirement arising within the MyHealthAvatar project. It reviews the state-of-art local cloud deployment and provides the detail of MyHealthAvatar local cloud deployment using OpenStack. In addition, the document reports the exploitation of public cloud for supporting MyHealthAvatar demonstrator and investigates the possibility of using MapReduce on Apache Hadoop for the support of distributed computation.

KEYWORD LIST:
Cloud Computing, Private Cloud, OpenStack, Public Cloud, MapReduce

[1] **R**=Report, **P**=Prototype, **D**=Demonstrator, **O**=Other

[2] **PU**=Public, **PP**=Restricted to other programme participants (including the Commission Services), **RE**=Restricted to a group specified by the consortium (including the Commission Services), **CO**=Confidential, only for members of the consortium (including the Commission Services)

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 600929.

MODIFICATION CONTROL			
Version	**Date**	**Status**	**Author**
0.1	23/06/2014	structure	Xia Zhao
0.2	22/7/2014	draft	Emmanouil G. Spanakis, Manolis Tsiknakis, Georgios Zacharioudakis, Kostas Marias
0.3	24/09/2014	draft	Xia Zhao, Zhikun Deng
1.0		final	Xia Zhao, Zhikun Deng

List of contributors

- Xia Zhao, University of Bedfordshire
- Zhikun Deng, University of Bedfordshire
- Emmnaouil G. Spanakis, FORTH
- Georgios Zacharioudakis , FORTH
- Tsiknakis Manolis, TEI-C, FORTH
- Kostas Marias, FORTH
- Feng Dong, University of Bedfordshire

Contents

1 Executive Summary

MyHealthAvatar follows recommendations from relevant VPH activities on “Digital Patient”. MyHealthAvatar architectural platform will be designed as an integrated facility that allows multiple functionalities rather than just a data storage facility as in the previous attempts. Its distinctive features include:

- ICT utilities to support data collection with minimal user input, including web information extraction, mobile apps, etc.
- ICT toolbox to support clinical decisions by using simulation models and by using visual analytics.
- Data and model repositories to provide rich resources of data and models
- Ontology and RDF repositories to support data search and reasoning.
- A cloud based ICT architecture that allows the access of data from a range of different sources, and integration of the repositories, the toolbox and the ICT utilities.
- A local cloud solution to support the computing requirement for the avatars without remote data transfer.
- A proof of market on open sources for MyHealthAvatar APIs.

The 4D Avatar proposed within the MyHealthAvatar project is intended to be a citizen’s lifelong companion a complete digital representation of his/her health-related data, including both medical records and social or lifestyle information, to create a comprehensive picture of the person in the context of their individualised healthcare. All of the data will be useful at different levels in assisting health decisions to be made.

This deliverable reports the investigation on deployment of a local cloud infrastructure for support data storage and computation requirement arising within the MyHealthAvatar project. This document covers the activities for Tasks T3.5 and Milestone MS033 in respect to review of state-of-art cloud technologies, the selection of implementation technologies and the deployment of the local cloud for MyHealthAvatar. Although, the investigation task is completed by this deliverable, the support of using it for other work packages will be continued to the end of the project.

This report aims to investigate the potential of locally-deployed cloud infrastructure as well as public available cloud computing facilities to support data storage and computation requirements for scalable data processing within the MyHealthAvatar project. It studies the background and the state-of-art technologies of cloud computing, and focuses on deployment of a private cloud computing environment and exploitation of the public cloud.

2 Introduction

Cloud computing has been making radical changes to the way in which we undertake research, and the general citizen performs everyday tasks. An increasing number of organisations and research projects have either adopted cloud computing based solutions or considering a move to cloud computing. However, there are many concerns in the adoption of public clouding computing, such as data security and privacy as well as the cost of transferring and storing large volume of data. Hence, private cloud computing becomes attractive to many organisations.

Within MyHealthAvatar, collecting, accessing, managing and possibly sharing healthcare related data are not only important to individuals who can manage their own health, but also important for clinicians and other healthcare workers for patient monitoring and providing suitable in-time care. The consideration of cloud technology is vital to ensure the long term scalability and performance of such systems.

2.1 Structure of this document

This document analyses the cloud usage within the MyHealthAvatar project together with the background of cloud technologies (see Section 3). It reviews the state-of-art local cloud deployment (see Section 4) and provides the detail of MyHealthAvatar local cloud deployment using OpenStack (see Section 5). In addition, the document reports the exploitation of public cloud for supporting MyHealthAvatar demonstrator (see Section 6). The report also investigates the possibility of using MapReduce on Apache Hadoop for the support of distributed computation (see Section 7).

3 Analysis the Cloud Usage in MyHealthAvatar

3.1 Overview of Cloud deployment models

The term "Cloud Computing" implies the provisioning of computing resources, both hardware and software, in the form of a service [1]. Originated from Service Oriented Computing (SOC), services in cloud computing also refer to fine-grained reusable resources. These services are easy access, scale well and have fewer barriers for consumers. In principle, we distinguish between four service models [2]:

- Infrastructure as a Service (IaaS): IaaS is the supply of hardware as a service, that is, servers, net technology, storage or computation, as well as basic characteristics such as Operating Systems and virtualization of hardware (Amazon S3-EBS-EC2 [4,5], Eucalyptus[3], Oracle Coherence[4], OpenStack[6], RightScale[5], 3Tera App Logic [13], EnStratus[6], Flexiscale[7], GoGrid CloudStatus[8])
- Platform as a Service (PaaS): At the PaaS level, the provider supplies more than just infrastructure, i.e. an integrated set of software with all the stuff that a developer needs to build applications, both for the developing and for the execution stages. (Google App Engine [14], Microsoft Azure Services [15], Amazon SimpleDB [4], Oracle Higher Education Constituent Hub [16], Amazon SQS [17], , Force.com[9])
- Software as a Service (SaaS): Offer software as a service. This was one of the first implementations of Cloud services (Google App[10], Microsoft Dynamics CRM online[11], Microsoft Live@edu[12], Business Producivity Online Suite[13], CampusEAI[14], EducationERP.net[15])
- Network as a Service (NaaS) [3]: a category of cloud services where the capability provided to the cloud service user is to use network/transport connectivity services and/or inter-cloud network connectivity services. NaaS involves the optimization of resource allocations by considering network and computing resources as a unified whole.

These services are characterized by certain qualities, like, on-demand self-service, broad-band access, High Availability and rapid elasticity.

[3] https://www.eucalyptus.com/
[4] http://coherence.oracle.com/
[5] http://www.rightscale.com/
[6] http://www.enstratius.com/
[7] http://www.flexiscale.com/
[8] http://www.gogrid.com/
[9] http://www.force.com/
[10] http://www.google.com/enterprise/apps/business/
[11] http://www.microsoft.com/en-gb/dynamics/crm.aspx
[12] http://office.microsoft.com/en-gb/academic/
[13] http://www.microsoft.com/business/bpostestdrive/
[14] http://www.campuseai.org/
[15] http://educationerp.net/

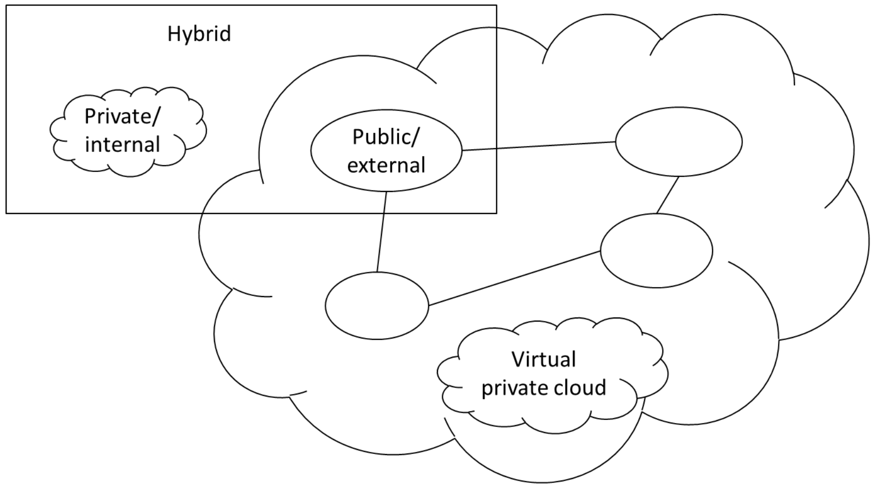

Figure 1. Deployment models of cloud computing.

As shown in Figure 1, currently there are mainly five deployment models of cloud: public cloud, private cloud, hybrid cloud and virtual private cloud. Community cloud is also commonly used in the research environment.

i. Private cloud: It is deployed within an organization's infrastructure and the resources are dedicated to the organization itself. Management and resource allocation is also controlled in-house by the organization.
ii. Public clouds: are open to public, but resources and infrastructure if owned by the organization providing the cloud service. Although public cloud providers often ensure client's data security and integrity, data control, especially for sensitive data, can always be an issue.
iii. Hybrid clouds: In this case, hybrid clouds combine both public and private clouds using data and application migration techniques.
iv. Virtual private cloud is a cloud computing infrastructure which runs on top of a public cloud which has secure and private settings for a particular organisation. By using virtual private cloud, organisations need less effort and maintenance for the cloud settings. SLAs are typically used by the public cloud provider and the consumers.
v. Community cloud refers to the cloud computing infrastructure that is set by a consortium with several organisations which likely have similar cloud computing requirements. This has been seen used in some research projects such as RESERVOIR[16].

3.2 Consideration of Cloud in MyHealthAvatar

The investigation of cloud infrastructure provides two key supports for the development of MyHealthAvatar demonstrator: scalable data storage and distributed data processing.

3.2.1 Cloud-Ready Data Storage

The design of repositories in MyHealthAvatar takes into account that the data will be large and likely to fall into the category of Big Data, bearing in mind the number of users and the quantity and types

[16] http://www.reservoir-fp7.eu/

of the information that will be stored on daily basis for the duration of a person's entire life. The use of cloud technology is promising in managing such data especially in terms of transparent database scalability, increases data availability, high ability of data distribution, and various data type support.

Cloud technology can offer the elasticity to scale well when the data storage/processing requirements as well as user traffic increase in MyHealthAvatar. We can start with 2-3 cloud nodes and modify later when the usage requirements change.

Cloud technology typically has the ability to distribute compute resources and data across different geographies or "zones" that the cloud provider makes available. The multi-geographic data distribution supplies the benefits of high availability if a particular region goes down, and simplified disaster recovery so the data is protected if a particular physical disaster befalls one of the cloud provider's data centres. This is critical to ensure the service available when MyHealthAvatar platform goes live to the public.

As proposed in D6.2 Design of Data and RDF Repositories, the NoSQL database Apache Cassandra is deployed as the main data repository for the MyHealthAvatar demonstrator. Cassandra offers a built-for-scale architecture which can smoothly scale across cloud nodes distributed at different data centres and can handle terabytes of information in various data types (unstructured, semi-structured, or structured) and thousands of concurrent users/operations per second.

3.2.2 Distributed Data Processing and Computation

Processing big longitudinal user data is one of the key challenges with MyHealthAvatar. The investigation of cloud computing technology also looks into how the deployment of the cloud infrastructure can enable the scalable data management and processing conducted in WP5/WP6, efficient visual analytics conducted in WP8. The current investigation has been focusing the cloud support for distributed data processing on heterogeneous user data, including the management of missing/noisy data. Further cloud computation investigation may require later with the continued development of the MyHealthAvatar.

MapReduce is a popular parallel computing method for cloud computing, which will be considered in MyHealthAvatar for computation and data processing tasks. MapReduce splits an input file into multiple chunks. In the first phase of MapReduce, a Mapper for each chunk reads the data, performs some computation, and outputs a list of key/value pairs. In the next phase, a Reducer combines the values belonging to each distinct key and outputs the result.

Together with Apache Cassandra, other data processing framework will also be considered, such as Apache Spark[17] for real-time in memory large scale data processing.

3.3 Existing Cloud Platform evaluation

For the purpose of this project and our deployment of a local cloud, we investigate four state-of-art private cloud implementations: OpenStack, Eucalyptus, CloudStack and VMWare vSphere. In order

[17] https://spark.apache.org/

to clarify the differences between them, we need to compare these products in terms of Storage, Virtualization, Networking, Management, Security and Community/Support criteria:

- *Storage criteria*: platforms must provide connectivity to either network block devices via iSCSI or FC, or via direct storage. Fault tolerance and HA configuration requirements must be met.
- *Virtualization criteria*: with this term, we define criteria concerning VM provisioning, resource management, migration, Hypervisor compatibility etc.
- *Network criteria*: network connectivity is essential to cloud computing. A cloud platform must be capable of either flat or VLAN networking.
- *Management criteria*: a cloud platform should provide effective management facilities in terms of logging, reporting, recovery mechanisms, high availability and hypervisors and guest OSs management and deployment.
- *Security criteria:* due to the scale and the sensitivity of data, a cloud platform should provide security mechanisms like user and key management, data encryption auditing and reporting.
- *Maturity and support criteria:* a cloud platform can be either developed in either a community based environment or inside a company. Active development and support are essential to its viability and especially for open source software, a large community of users along with the developers are critical for the welfare of the platform.
- *Performance criteria:* it is crucial that any cloud platform is evaluated for its performance, especially for storage, since it can usually be a bottleneck for the whole infrastructure.

Next, we provide a brief description of the basic concepts of four local cloud deployments, i.e. CloudStack, OpenStack, Eucalyptus and VMWare vSphere. We analyse the building blocks of each software and we focus in storage and public cloud interoperability options for each one. For public cloud interoperability, we use as metric the compatibility level they provide to the world market leader in public cloud platforms, Amazon EC2 and Amazon S3. EC2, is a web service that provides resizable compute capacity in the cloud [4] and S3, provide web services interface that can be used to store and retrieve data [5]. Both services are part of the Amazon web services facilities.

4 State-of-art Analysis of Local Cloud Deployment

The analysis of state-of-art local cloud deployment has been performed by two partners – BED and FORTH, on experiment of four development stacks: Eucalyptus, VMWare cloud, CloudStack and OpenStack. Among them at least two of the enterprise scale, open source, cloud platforms seem to be the most promising, OpenStack and Eucalyptus. OpenStack provides an API for applications built on top of the cloud to access a specialized “Object Storage”, which will be analysed later, and Eucalyptus is designed with Amazon Cloud interoperability in mind. Both cloud operating systems can be deployed in order to manage multiple cluster environments. Both platforms are enterprise scale but besides the similarities, they differ in many ways.

The remainder of this section provides the detail of each experiments and analysis of the results.

4.1 Eucalyptus

Overall Structure of Eucalyptus Cloud“: Eucalyptus is an open source private cloud software for building private and hybrid clouds that are compatible with AWS APIs. With AWS-compatibility, the open source software pools together existing virtualized infrastructure to create private or hybrid cloud resources for compute, network, and storage." [7]

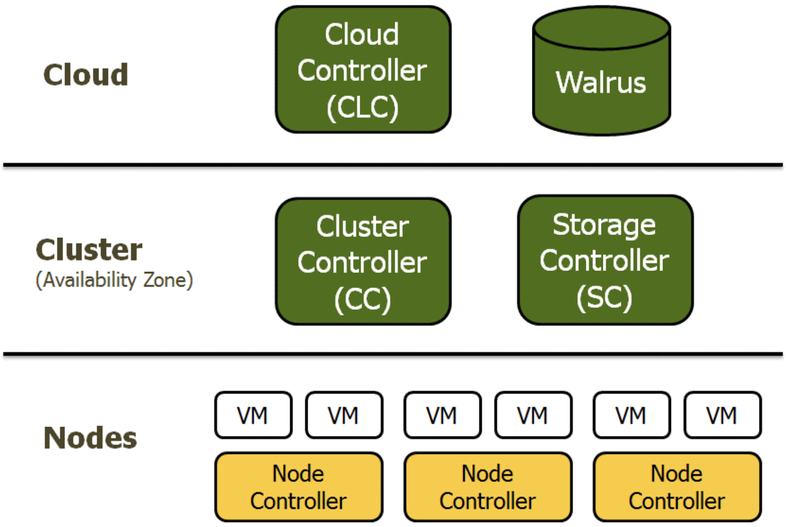

Figure 2. Overall structure of Eucalyptus cloud.

The structure of the Eucalyptus platform is presented by Figure 2. It consists of six different components:

Component	Abbreviation	Description
Cloud Controller	CLC	A web based interface and EC-2 interfaces. It handles incoming requests and provides administrative control and management of the infrastructure, as long as high-level resource scheduling and system accounting
Walrus	–	Offers persistent storage to VMs
Cluster Controller	CC	It acts as the front end for a cluster within a Eucalyptus cloud and communicates with the Storage Controller and Node Controller. It manages instance (i.e., virtual machines) execution and Service Level Agreements (SLAs) per cluster.
Storage Controller	SC	It communicates with the Cluster Controller and Node Controller and manages Eucalyptus block volumes and snapshots to the instances within its specific cluster.
VMware Broker	–	Provides an AWS-compatible interface for VMware environments and physically runs on the Cluster Controller
Node Controller	NC	Hosts the virtual machine instances and manages the virtual network endpoints

Next, we provide a small overview of four basic components of Eucalyptus: Cloud Controller, Walrus, Cluster Controller and Storage Controller.

4.1.1 Cloud Controller

Cloud Controller (CLC) offers EC2-compatible SOAP and Query interfaces, as well as a Web interface to the outside world. It acts as the administrative interface for cloud management and performs high-level resource scheduling and system accounting. CLC handles Authentication, Accounting, Reporting and Quota Management.

4.1.2 Walrus

Walrus is the equivalent to AWS Simple Storage Service (S3). It offers persistent storage to VMs and can be used as a simple HTTP put/get Storage-as-a-Service solution. It does not have any data restrictions.

4.1.3 Cluster Controller

Cluster Controller (CC) acts as the front end for a cluster within a Eucalyptus cloud and communicates with the Storage Controller (SC) and Node Controller (NC). The CC manages instance (i.e., virtual machines) execution and Service Level Agreements (SLAs) per cluster.

4.1.4 Storage Controller

Storage Controller (SC) is the equivalent to AWS Elastic Block Store (EBS). The SC communicates with the Cluster Controller (CC) and Node Controller (NC) and manages Eucalyptus block volumes and snapshots to the instances within its specific cluster.

4.1.5 Amazon EC2 and Amazon S3 support

Eucalyptus implements by design the Amazon specifications for EC2 and S3. It provides REST and SOAP interface compatible with Amazon AWS [8]. Moreover, Eucalyptus is part of the AWS partner network [9].

In brief, these are the AWS features supported by Eucalyptus

- Amazon Elastic Compute Cloud (EC2)
- Amazon Elastic Block Storage (EBS)
- Amazon Machine Image (AMI)
- Amazon Simple Storage Service (S3)
- Amazon Identity and Access Management (IAM)
- Auto Scaling
- Elastic Load Balancing
- Amazon CloudWatch

4.2 VMWare vSphere

VMware vSphere (formerly VMware Infrastucture 4) is VMware’s server virtualization platform for build cloud infrastructure. VMware vShpere platform is formed by a series products and components as shown in Figure 3, which includes core component - VMware ESXi Server, management – vCenter Server and vSphere Client. The following is a brief introduction to the main component of VMware vSphere.

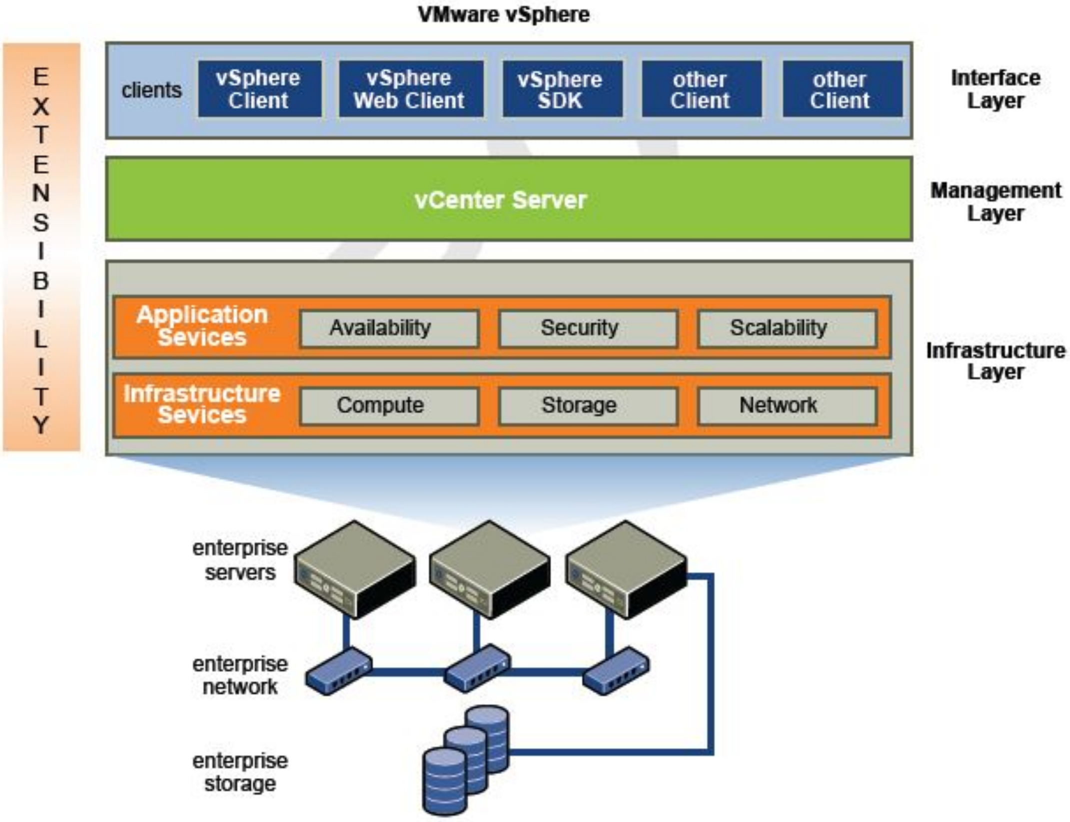

Figure 3. Overall structure of VMWare vSphere.

4.2.1 VMware Hypervisor (ESXi)

VMware Hypervisor is a free bare-metal hypervisor, it install/run directly from physical server and does not require a separate operation system (OS) to be installed. It has less overhead compared to other hypervisors which require a hosting OS. VMware vSphere used to support both ESX and ESXi, however from version 5, it only support ESXi as its hypervisor [20].

4.2.2 VMware vCenter Server

VMware vCenter server is the centralized platform for configuring and managing VMware vSphere environments. It provides core data center services including access control, performance monitoring and warning managements.

4.2.3 VMware vShpere Client/Web Client

VMware vSphere client is a windows application which connects to vCenter Server or EXSi user interface, and web client allows user to connect to vCenter server through browsers (with plugin installed).

4.2.4 vSphere Virtual Machine File System (VMFS)

VMFS is a high-performance cluster file system for VMware hypervisor (EXSi), and it is optimized for Virtual machines. VMFS allow multiple vSphere hosts to access the same storage concurrently by leverage shared storage.

4.2.5 vSphere Virtual Symmetric Multi-Processing (SMP)

Virtual SMP is a utility which allows a single virtual machine to use multiple physical processors simultaneously.

4.2.6 vSphere vMotion/Storage vMotion

VMware vMotion allows user to move an entire running virtual machine from one physical server to another one without downtime. It also allows automate and schedule migrations within the same datacentre. Storage vMotion allows live migration of virtual machine disk files within and across storage arrays, which is performed at zero-downtime and without service disruption.

4.2.7 vSphere High Availability (HA)

HA monitor vSphere hosts and virtual machines to detect hardware and OS failures, and restart virtual machine on other hosts in cluster automatically when a server outage is detected. HA reduces application downtime by automatic restart VM and delivers the availability.

4.2.8 vSphere Distributed Resource Scheduler (DRS) / Storage DRS

DRS groups vSphere hosts into resource clusters to segregate the computing demands of various business requirements. It provides HA resources for workloads, balance workloads for optimal performance and help scale computing resources.

4.2.9 vSphere Fault Tolerance (FT)

FT provides continuous availability for applications in the event of server failure. FT creates a live shadow instance of a VM which is always up-to-date with the primary VM, in the event of outage, FT automatically triggers failover – this ensures the shadow VM become live and prevent downtime or data lose.

4.2.10 vSphere Distributed Switch (VDS)

VDS provides a centralized interface for configuration, monitoring and administration of VMs across the whole data centre. It enables the same network identity when a VM is migrated between different hosts.

4.2.11 Amazon EC2 and Amazon S3 support

VMWare has two produces, which are vCloud Automation Center [18] and vFabric Application Director [19] that can support deployment on Amazon. Automation Center allows policy-based provisioning across VMware-based private and public clouds, physical infrastructure, multiple hypervisors, and Amazon Web Services. Application Director enables applications to be deployed across multiple virtual and hybrid cloud infrastructures, including Amazon EC2.

4.2.12 Environment Setup for Feasibility Test

During the cloud investigation, a small VMWare vSphere has been established in University of Bedfordshire as described below. The VMs themselves can be configured to be clustered computing, Hadoop cluster of two VMs are configured and tested. The findings are very promising and the performance is good. However, due to the licence cost, this cloud is not used for the consortium within MyHealthAvatar.

For MyHealthAvatar feasibility research, two Dell T7400 servers (with two Xeon E5410, 1T storage space, and 16G memory each) are connected to gigabit network switch. A windows PC with vSphere client installed is used to configure and install VMs. The basic layout the validation platform is configured as following graph.

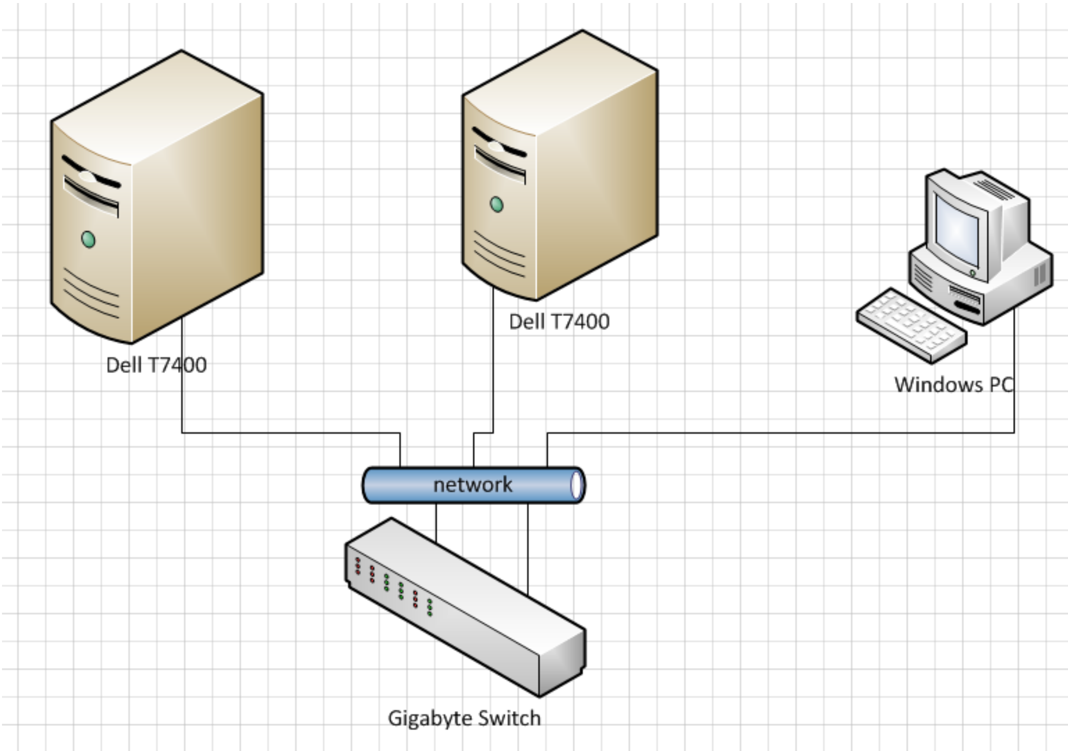

Figure 4. Feasibility setting of VMWare vSphere.

The configurations of servers and networks are for validation of the VMware vSphere for MyHealthAvatar, so all the resources are dedicated to MyHealthAvatar platform, which includes 32G full-buffered memory and 2T storage space and 16 cores of 2.33GHz CPU. For a normal VM, 4-6G memory and 1-2 cores are allocated, which gives us 6 powerful VMs to run MyHealthAvatar platform tests.

Regarding security, Https are always used for the vCenter server web interface access, strong password policy are enforced to administrators and users who have access. Least rights are assigned to various parties, this is to control the range of damage might be done with stolen credentials. For VMs access, Linux shell access is given with only private/public key authentication allowed, while windows are allowed to access through Remote Desktop Connections.

One of the servers has vSphere vCenter Server installed, which allows the management of the servers from various platforms through browsers. For the validation purpose, we have carried out following activities on above platform.

4.2.12.1 Install vSphere Hypervisor on both servers

Two 1T USB drives are used to install the EXSi on both servers, and the server then boot up from the USB drives. This allows all the hard drive capacities be used for VMs. USB access speed is not a concern here, since all content of USB is loaded into memory after initial boot process.

4.2.12.2 Connect to EXSi hosts and install vSphere vCenter Server

Windows PC is utilized to install vSphere client which connects to EXSi hosts, configure the hosts and install vSphere vCenter Server (VM) on one of the server, which further allows the clustering and management of hosts through web interface.

4.2.12.3 Configure both server as cloud and install Linux/Windows VMs

Configuration is done to make two server cluster together as local cloud environment, the vCenter server allows configuration, management and monitoring in one web UI. Ubuntu Linux 12.04 LTS and Windows Server 2012 R2 are installed as a test of the VM hosting ability.

4.2.12.4 Configure VM templates, create VM from template

To have a standard template for consortium members to work on is very important, it avoid the configuration caused deployment issues. VMware vSphere platform support the template and create new VM based on template very well, the interface is intuitive and friendly.

4.2.12.5 Clone and migrate VM

Clone and migrate VM within same hosts and between two different hosts are tested, the processes are completed without any problem.

4.2.12.6 Monitoring, disaster recovery

The monitoring of VM through vCenter interface, and tested disaster situations includes one hosts get offline and power cut of both hosts. The recovery from disaster as behaves as documented by VMware and the downtime depends on the situations (e.g. both hosts power cut will only go live when the power comes back again).

4.3 CloudStack

CloudStack architecture is based upon a hierarchical structure which provides IaaS cloud services for either private, public or hybrid configurations. It offers centralized management capabilities over all cloud components (network, storage and nodes) and it also provides on-demand access to infrastructure through a self-service portal.

Among other, CloudStack provides [10]:

- Multiple Hypervisor Support: CloudStack supports all major Hypervisor, like VMware, KVM and Xen.
- Massively Scalable Infrastructure Management: It is designed to support tens of thousands of servers installed in multiple geographically distributed datacenters.
- Automatic Configuration Management: Takes advantage of several software appliances offering DHCP, VPN firewall, routing, proxy and storage replication services.
- Graphical User Interface: A customizable UI is provides for both cloud administrators and end-users.
- API and Extensibility: CloudStack API gives programmatic access to all the management features available in the UI.
- High Availability: There is a number of HA options available to the cloud: management servers can be deployed in HA setup, supports NIC bonding on hosts and multipath configurations for storage access over iSCSI.

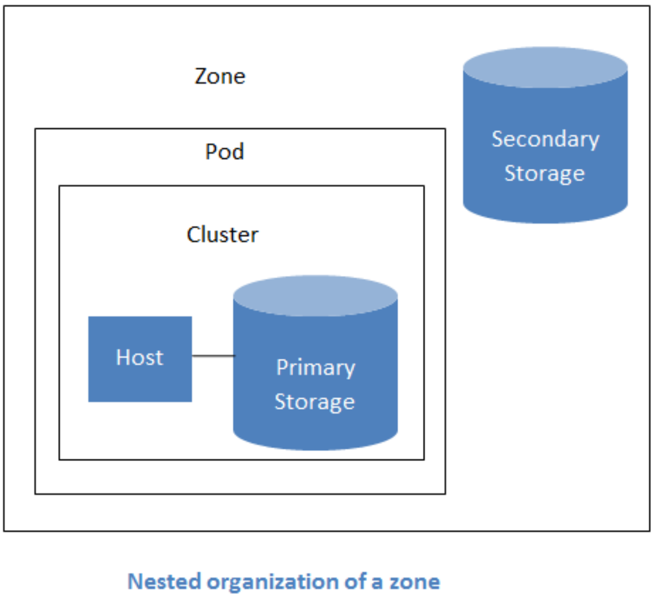

Figure 5. Overall structure of CloudStack.

CloudStack is organized in the form of nested components as shown in Figure 5. The primary organization structure of CloudStack is *Zone.* A Zone, is an organizational unit, which can be identified for example, as a single Datacenter. Each zone consists of *Pods.* Pods, can be a single rack of servers in a Datacenter. Hosts belonging to a Pod share the same subnet. Pods consist of *Clusters* and Clusters consist of Hosts and their affiliated primary storage. Clusters, are an organizational structure for grouping hosts together. It can be a pool of Xen, KVM or ESXi (VMware) hosts. Each Cluster has its own primary storage. There is also another type of storage, *Secondary Storage*, which is shared among by all the components of a Zone.

In brief, CloudStack has 9 components:

Component	Description
Hosts	Machines used for provisioning services
Cluster	A host group with their storage
Pod	A collection of clusters
Zone	A collection of pods.
Management Server Farm	Management nodes
Primary Storage	VM storage
Secondary Storage	Storage for support stuff, (templates, ISO images etc)
Network	A logical network structure associated with services

4.3.1.1 Amazon EC2 and Amazon S3 support

CloudStack includes, as part of its API, AWS services integration. It is compatible with Amazon's EC2 and S3 services and in the current version, the related service is fully integrated in CloudStack. CloudStack provides a translation service for AWS API calls, which are translated through this service into its own API calls. It supports both EC2 SOAP & Query API and S3 REST API [11].
There are also certain limitations:

- AWS is available only through zones with basic networking
- Elastic IP and Elastic Load Balancing are only available with Citrix NetScale device.

4.4 OpenStack

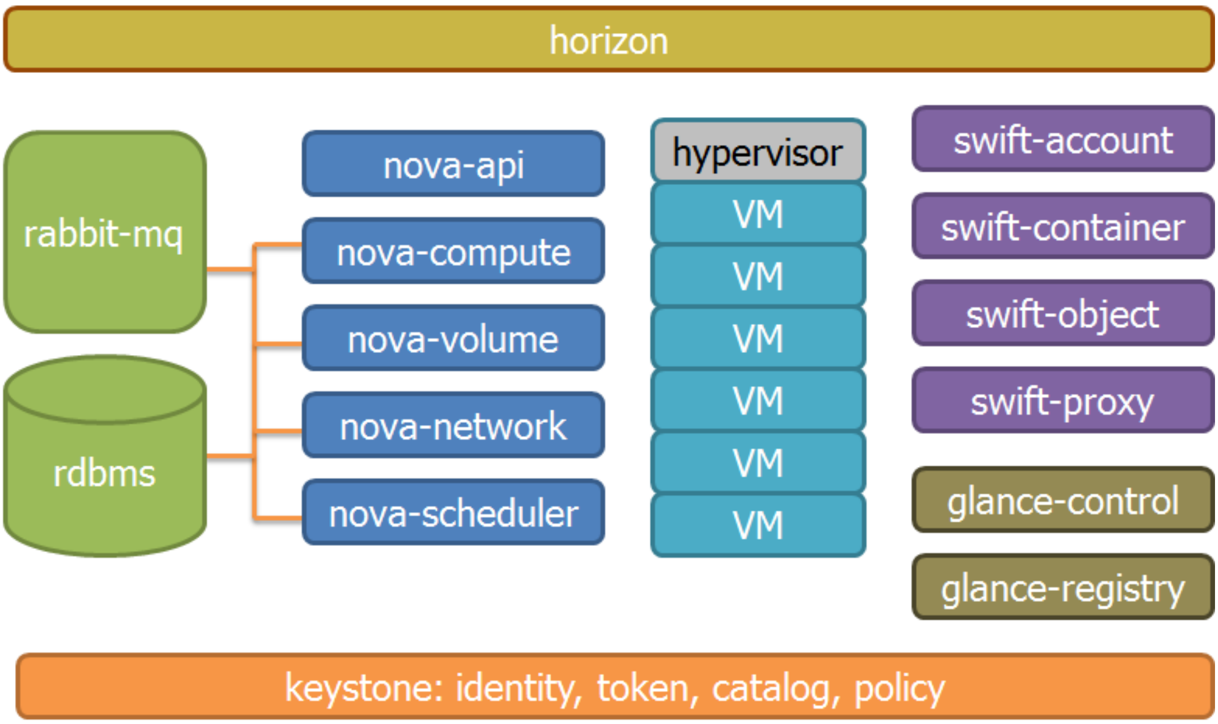

Figure 6. Overall structure of OpenStack.

In this section we describe the background of OpenStack and its general structure. OpenStack is a cloud operating system which basically, provides IaaS. It is actually designed to manage data centers and has no proprietary hardware or software requirements. Figure 6 depicts the overall structure of OpenStack.

OpenStack is built upon nine different services:

Service	Codename	Description
Dashboard	Horizon	A web based interface for OpenStack services
Compute	Nova	Provides virtual servers on demand
Identity	Keystone	A framework for authentication and authorization
Network	Neutron	It provides "*network as a service*" between interface devices (e.g., vNICs) managed by other OpenStack services (e.g., nova). The service works by allowing users to create their own networks and then attach interfaces to them. Quantum has a pluggable architecture to support many popular networking vendors and technologies
Image Service	Glance	Catalog and repository service for virtual disk images
Block Storage	Cinder	Traditional block storage services to VMs (iSCSI, FC etc.)
Object Storage	Swift	Object store allows you to store or retrieve files. **It provides a fully distributed, API-accessible storage platform that can be integrated directly into applications or used for backup, archiving and data retention.**
Monitoring	Ceilometer	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability, and statistics purposes.
Orchestration	Heat	Orchestrates multiple composite cloud applications by using the AWS CloudFormation template format, through both an OpenStack-native REST API and a CloudFormation-compatible Query API.

4.4.1 Storage

Storage options that OpenStack provides are of great significance. The object storage is exposed via the relevant API, so that it can be integrated into applications or even used as backup and data archiving. The OpenStack object storage, has several capabilities:

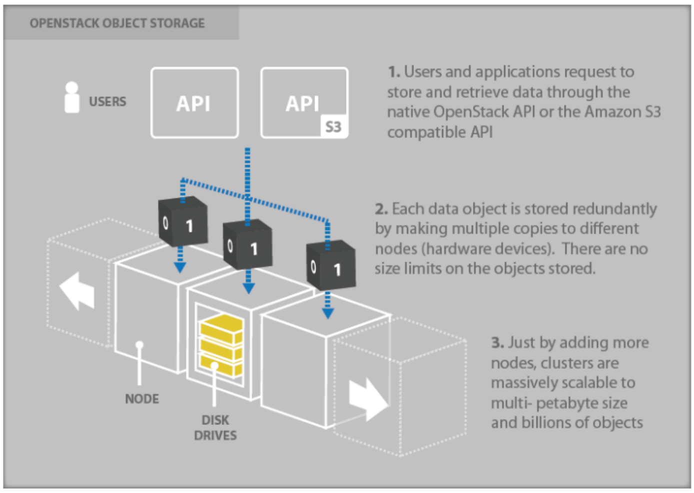

- Object storage is provided using clusters of servers capable of storing large amount of data of

data. The storage has characteristics of redundancy and scalability.
- It's a distributed storage system for static data such as virtual machine images, photo storage, email storage, backups and archives.
- Objects and files are written to multiple disk drives spread throughout servers in the data center
- Storage clusters scale horizontally simply by adding new servers.
- Should a server or hard drive fail, OpenStack replicates its content from other active nodes to new locations in the cluster.

On the other hand, Block storage uses traditional storage options over iSCSI or FC protocols. In brief:

- The block storage system manages the creation, attaching and detaching of the block devices to servers.
- It has unified storage support for numerous storage platforms including Ceph, NetApp, Nexenta, SolidFire, and Zadara.
- Block storage is appropriate for performance sensitive scenarios such as database storage, expandable file systems, or providing a server with access to raw block level storage.
- Snapshot management provides powerful functionality for backing up data stored on block storage volumes. Snapshots can be restored or used to create a new block storage volume.

4.4.2 Networking

OpenStack networking of VMs can be deployed with either as flat networking or VLAN networking.
To summarize:

- OpenStack Networking manages IP addresses, allowing for dedicated static IPs or DHCP. Floating IPs allow traffic to be dynamically rerouted to any of your compute resources, which allows you to redirect traffic during maintenance or in the case of failure.
- Users can create their own networks, control traffic and connect servers and devices to one or more networks.
- The pluggable backend architecture lets users take advantage of commodity gear or advanced networking services from supported vendors.
- Administrators can take advantage of software-defined networking (SDN) technology like OpenFlow to allow for high levels of multi-tenancy and massive scale.
- OpenStack Networking has an extension framework allowing additional network services, such as intrusion detection systems (IDS), load balancing, firewalls and virtual private networks (VPN) to be deployed and managed

4.4.3 Compute

OpenStack is designed to manage and automate pools of compute resources. It takes advantage of the popular KVM or Xen hypervisors and it also supports ARM and alternative hardware architectures

4.4.4 Amazon EC2 and Amazon S3 support

OpenStack provides API for both EC2 and S3. In the following matrices we present a view of these APIs and their compatibility level with Amazon services [6]:

OpenStack NOVA – EC2 API

FEATURE	SUPPORTED
EC2 QUERY API	Y
EC2 SOAP API	N
OPENSTACK API / RACKSPACE API	Y
SSL BETWEEN COMPONENTS	N
HORIZONTAL COMPONENT SCALABILITY	Y
WEB-BASED UI	Y
COMMAND LINE INTERFACE	Y

Open Stack Swift – S3 API Support

FEATURE	SUPPORTED
LIST BUCKET OBJECTS	Y
BUCKET ACLS	Y
BUCKET LIFECYCLE	N
BUCKET POLICY	N
BUCKET LOCATION	?
BUCKET LOGGING	N
BUCKET NOTIFICATION	N
BUCKET OBJECT VERSIONS	Y
BUCKET REQUEST PAYMENT	N
BUCKET VERSIONING	?
BUCKET WEBSITE	Y

4.5 Analysis of the Deployment Result

The following table summarizes the basic characteristics of OpenStack, Eucalyptus, CloudStack and VMWare vSphere.

	OpenStack	Eucalyptus	CloudStack	VMWare vSphere
	General			
Architecture	Fragmented into lots of pieces	Five main components. AWS clone	Nine main components in nested structure	Lots of components of various functionalities
Installation	RPM/DEB packages	RPM/DEB packages	Five main components. AWS clone	Standalone from USB drive, then vCenter Server VM
Administration	Web UI, euca2ools, native CLI	Strong CLI compatible with EC2 API	Customizable WEB UI and AWS compatible API	Windows Client, Web UI
Security	Baseline + Keystone	Baseline + component registration	SSO, user roles	SSO, Security Token Service and internal

				LDAP for Users and Groups
High Availability	Swift Ring, otherwise manual effort	Primary/secondary component failover	Load balancing and HA options	High Availability and Fault Tolerant components
Hypervisors	Xen, KVM, VMware	Xen, KVM, VMware	Xen, KVM, VMware	VMware Hypervisor (EXSi)
Language	Python	Java, C	Java	Assembly, C, C++
OS Support	Various Linux distributions, Windows and BSD systems	Various Linux distributions, Windows and BSD systems	Various Linux distributions, Windows and BSD systems	OS X Server 10.6+, Various Linux, MS-DOS, Windows, BSD, OS/2 Warp 4, NetWare 6.x, Solaris 10
		Storage		
Disk Images	Yes	Yes	Yes	Yes
Block devices	iSCSI, FC	iSCSI, FC	iSCSI, FC	iSCSI, FC
Object Storage	Yes	S3 API	Limited S3 API	Yes
Fault tolerance	Yes	Yes	YEs	Yes
		VM image facilities		
Image service	Yes	Yes	Yes	Yes
User VM images	Yes	Yes	Yes	Yes
Amazon API	Yes	Yes	Yes	No
		User UI and Management Facilities		
Web Interface	Yes	Yes	Yes	Yes
Users & Quotas	Yes	Yes	Yes	Yes
		Networking		
Floating IPs	Yes	Yes	Yes	Yes
L2 Support, VLANS	Yes	Yes	Yes	Yes
DHCP Support	Yes	Yes	Yes	Yes

5 Local Cloud Deployment for MyHealthAvatar

OpenStack is selected to be deployed in the premises of FORTH, in the form of a private computational and storage cloud for the consortium of MyHealthAvatar. In terms of hardware resources, the cloud infrastructure allows for maximum elasticity and flexibility by effectively adapting to the load of any given time. The current minimal specification includes:

- 300 GB of RAM
- 9TB of storage
- 16 cores Intel® Xeon® Processor E5-2690 and 4 cores Intel® Xeon® Processor E7520 (Dell PowerEdge R720 and SC 1425 Servers series)

In terms of the software the OpenStack open source cloud computing software has been installed on the machines using the Linux Ubuntu 12.04 operating system. The following lists the physical setup of the local cloud infrastructure.

- Physical security
  - Machine room in a restricted access area
  - Smartcard & PIN access control
  - Heavy duty air conditioning system
  - UPS (uninterruptible power supply)
  - Intrusion detection alarm
  - Temperature fluctuation alarm
  - Fire alarm
  - Automatic CO_2 fire extinguisher
  - Automatic power generator in case of power shortage
  - Administrative stuff available on nearby office
- Network security
  - Firewall
  - IP based access control
  - Certificate-based login (optional)
- Data security
  - RAID controllers (hardware)
  - Automatic replication of data (software -Openstack)
  - Tape backups

Figure 7-10 show the snapshots of the deployed OpenStack dashboard.

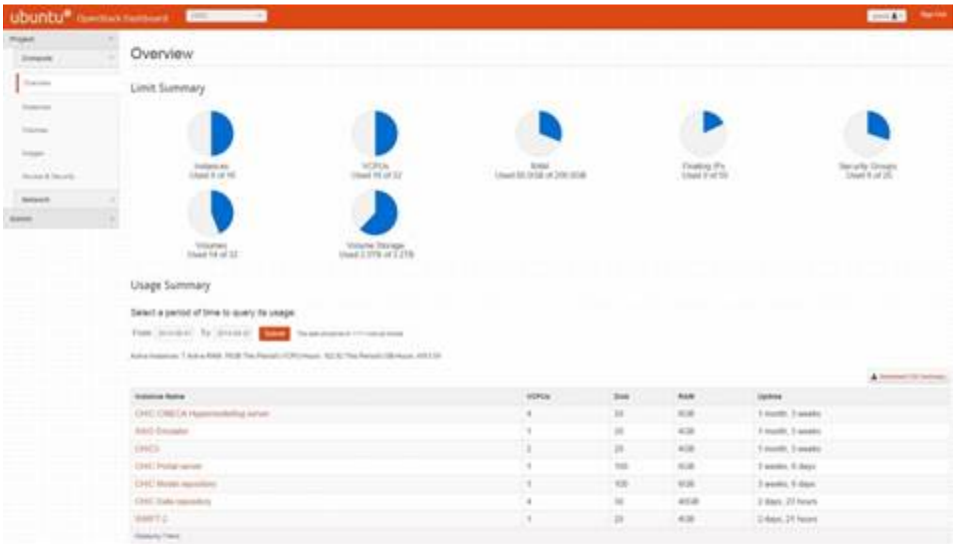

Figure 7. Snapshot of the OpenStack dashboard-limit summary.

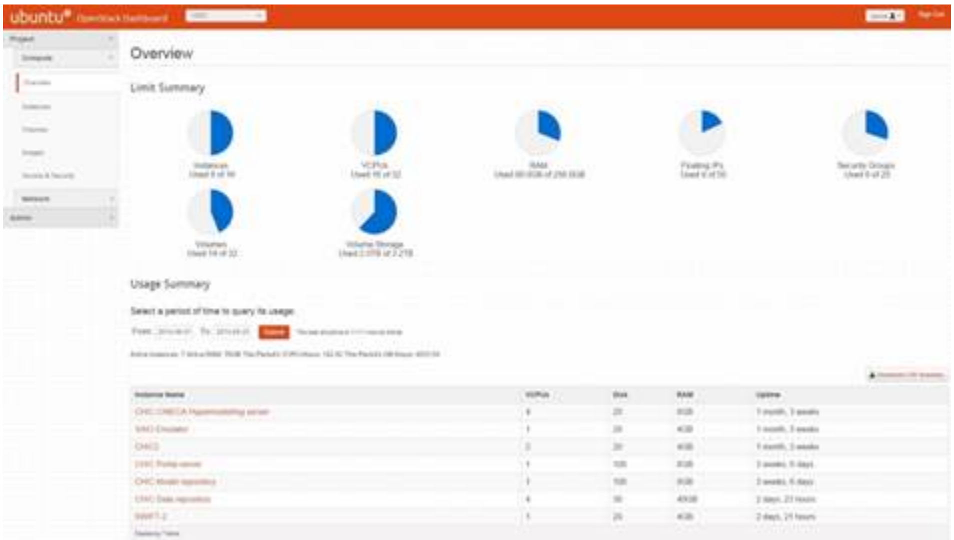

Figure 8. Snapshot of the OpenStack dashboard-usage summary.

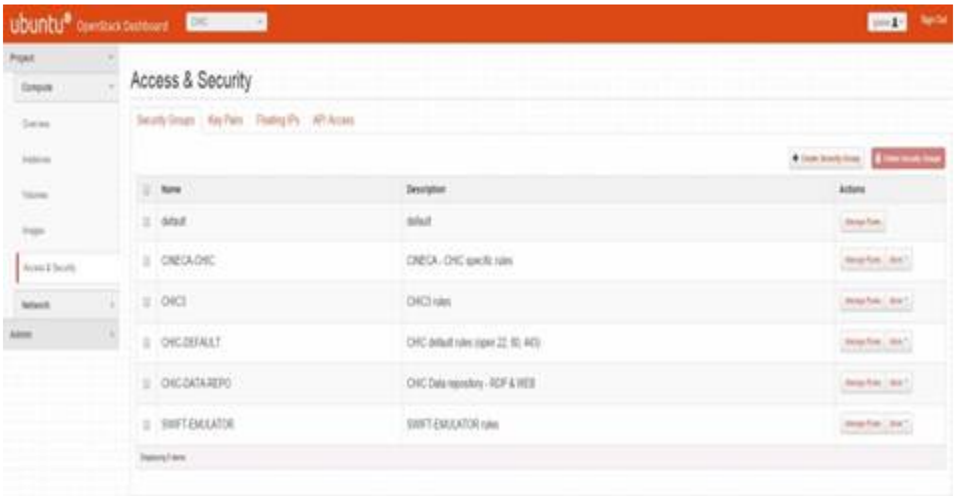

Figure 9. Snapshot of the OpenStack dashboard-access & security.

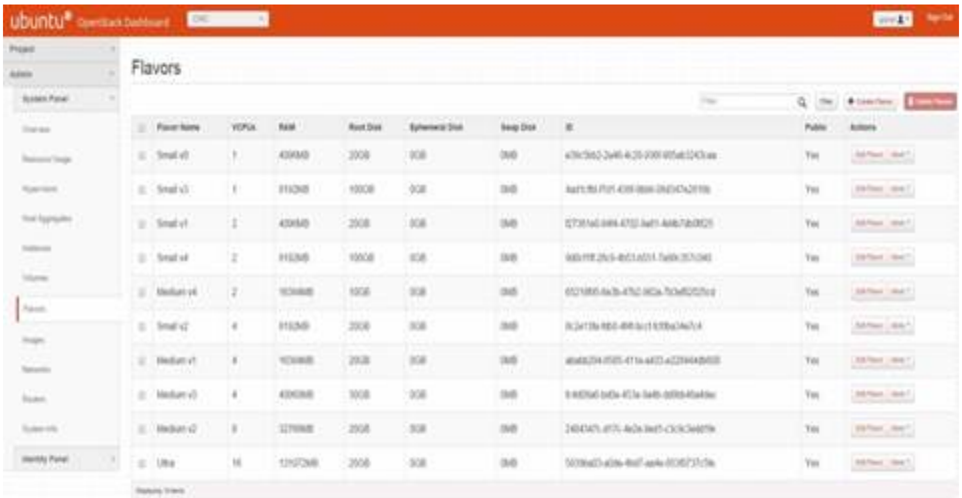

Figure 10. Snapshot of the OpenStack dashboard-instance list.

6 Support of Computation Requirements deployment for MyHealthAvatar

6.1 Hadoop MapReduce Setup

MapReduce is a programming paradigm that allows for massive scalability across hundreds and thousands machines in clusters. It has become popular for processing Big Data. As shown in Figure 11, a MapReduce program consists of two main procedures:

- For the map procedure, the master node takes an input, divided it into smaller sub inputs and distributes them to a number of work nodes. The map procedure performs sorting, shuffling and passes the results back to the master node.

- For the reduce procedure, the master node collects all results from the work nodes and summarise them in certain to form the outputs.

The application of the MapReduce paradigm is closely related to whether the problem can be partitioned into independent pieces which can be distributed to multiple mappers and reducers. Here, we list some examples of suitable and unsuitable MapReduce scenarios, respectively.

MapReduce is suitable for distributed scalable computation in the following examples.

- Log analysis: log entries can be mapped to <key, value> pair, and the reduce function can turn the data compatible to an entry of the system.
- Volume rendering: the work in [21] presents a multi-GPU parallel volume rendering implementation built using the MapReduce programming.

However, MapReduce is not suitable in the following situations:

- Computation depends on previous computed values, e.g. the Fibonacci series in which each value is the summarisation of previous two values. It is also not preferred for MapReduce if the dataset is small enough to be processed on a single machine.
- Algorithms depend on shared global space: map and reduce tasks run independently and in isolation, so it is difficult to implement when the algorithm requires access to shared global space.
- Real-time data processing: MapReduce jobs often assume the input as a static dataset; however, real-time data processing may see partial data or new datasets are introduced in an unexpected order.

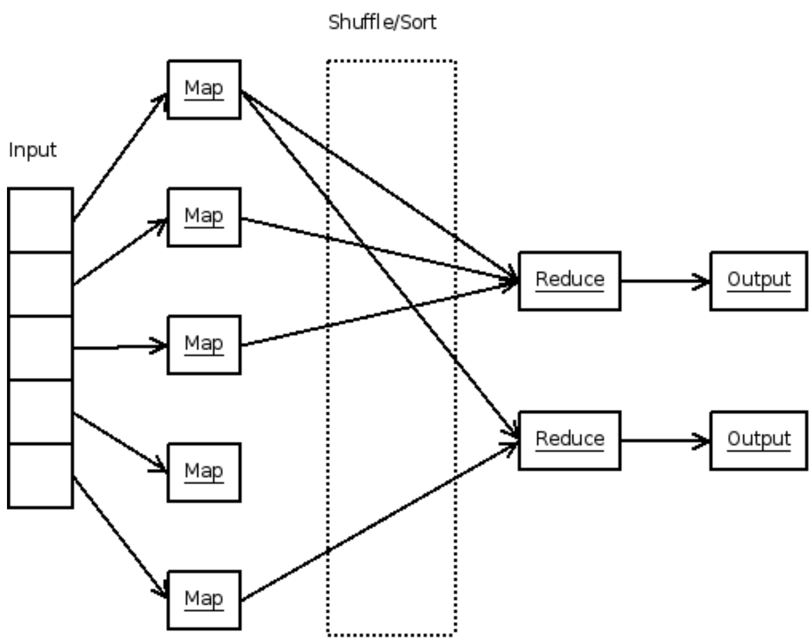

Figure 11. An overview of the Map/reduce programming paradigm.

7 Consideration of Public Cloud Services

Although the initial task of this report mainly focuses on the locally deployed cloud infrastructure, with the development of the project and the investigation of the cloud computing technologies, we also consider the use of publication cloud services for the MyHealthAvatar platform. In the process of investigation, the MyHealthAvatar demonstrator will be deployed on both private cloud at FORTH and commercial public cloud.

The major difference between locally deployed cloud (also known as private cloud) and public cloud is the control and access of the resource. In private cloud, resources are controlled and accessed by the premise only. On the other hand, in public cloud, resources are controlled by the cloud provider but are accessible for public users. Therefore, a hybrid cloud infrastructure is adopted in MyHealthAvatar with both public and private cloud facilities available.

Security is one primary concern when we choose the cloud service. A private cloud allows IT to control the perimeter and ensure the highly-secured data not to be transferred out to the premise; but it's also responsible for staying on top of a rapidly shifting security landscape and making all required fixes, updates, and upgrades. While by using public clouds, data is protected by both managed security on a software and physical level, since large-scale data centres like those used by public cloud providers have state-of-the-art security.

Service availability and disaster recovery is another key consideration. Public cloud typically operates in multiple data centres at multiple geo locations, which provides high level service availability especially in the scenarios of disaster recovery. Private cloud typically resides in one premise and requires expensive proposition in order to ensure the constant service availability.

8 Conclusion

This deliverable focused on the investigation of local cloud deployment for MyHealthAvatar. As pointed in the report, this investigation also introduces the feasibility use of public commercial cloud to ensure the high level of service availability.

The report analysed the storage and computation usage of the cloud for the MyHealthAvatar platform, and reviewed four state-of-art local cloud deployment technologies, namely Eucalyptus, OpenStack, CloudStack and VMWare vSphere. It reported the physical establishment of OpenStack at FORTH for support the deployment of MyHealthAvatar on local cloud. This report also discussed the possible adoption of MapReduce on cloud nodes for distributed data processing and computation.

Although the deployment of the local cloud task in concluded in this report, the participated partners (BED and FORTH) will continue to support the use of the cloud resources for WP5-7.

9 References

[1] "The NIST Definition of Cloud Computing," 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.

[2] Zhou, Minqi; Zhang, Rong; Zeng, Dadan; Qian, Weining;, "Services in the cloud computing era: a survay," in *4th International Universal Communication Symposium (IUCS), IEEE*, Shanghai, 2010.

[3] P. Costa, M. Migliavacca, P. Pietzuch and A. L. Wolf, "NaaS: network-as-a-service in the cloud," in *2nd USENIX conference on Hot Topics in Management of Internet, Cloud and Enterprise Networks and Services*, San Jose, CA, 2012.

[4] "Amazon Elastic Compute Cloud (Amazon EC2)," Amazon, [Online]. Available: http://aws.amazon.com/ec2/. [Accessed 09 10 2013].

[5] "Amazon Simple Storage Service (Amazon S3)," Amazon, [Online]. Available: http://aws.amazon.com/s3/. [Accessed 09 10 2013].

[6] OpenStack, [Online]. Available: https://wiki.openstack.org/wiki/Nova/APIFeatureComparison. [Accessed 09 10 2013].

[7] "Open Source Private Cloud Software," [Online]. Available: http://www.eucalyptus.com/eucalyptus-cloud/iaas. [Accessed 07 10 2013].

[8] "Eucalyptus FAQ," Eucalyptus, [Online]. Available: http://www.eucalyptus.com/faq#q4. [Accessed 09 10 2013].

[9] "Euclyptus," Amazon, [Online]. Available: http://www.aws-partner-directory.com/PartnerDirectory/PartnerDetail?id=3016. [Accessed 09 10 2013].

[10] "CloudStack Administrator's Guide," Apache Cloudstack, [Online]. Available: http://cloudstack.apache.org/docs/en-US/Apache_CloudStack/4.2.0/html/Admin_Guide/whatis.html. [Accessed 15 10 2013].

[11] "12.1. Amazon Web Services Compatible Interface," Apache Cloudstack, [Online]. Available: http://cloudstack.apache.org/docs/en-US/Apache_CloudStack/4.2.0/html/Installation_Guide/aws-ec2-introduction.html. [Accessed 15 10 2013].

[12] "CloudStack Administrator's Guide," Apache Cloudstack, [Online]. Available: http://cloudstack.apache.org/docs/en-US/Apache_CloudStack/4.0.0-incubating/html-single/Admin_Guide/#whatis. [Accessed 15 10 2013].

[13] CA 3Tera AppLogic for enterprise applications, Avaialble: http://www.scalematrix.com/wp-content/uploads/2010/12/applogic-enterprise-product-brief.pdf

[14] Google App Engine, Available: https://cloud.google.com/appengine/

[15] Microsoft Azure, Available: http://azure.microsoft.com/

[16] Oracle Higher Education Constituent Hub, Available: http://www.oracle.com/us/products/applications/master-data-management/education-

research-hech-ds-164687.pdf

[17] Amazon Simple Queue Service (SQS), Available: http://aws.amazon.com/sqs/

[18] VMWare vCloud Automation Center, Available: http://www.vmware.com/products/vcloud-automation-center

[19] VMWare vFabric Application Director, Available: https://www.vmware.com/files/pdf/vfabric/VMware-vFabric-Application-Director-Datasheet.pdf

[20] VMWare vSphere5, Available: https://www.vmware.com/support/vsphere5/doc/vsphere-esx-vcenter-server-50-new-features.html

[21] J. A. Stuart, C.-K. Chen, K.-L. Ma, and J. D. Owens, “Multi-GPU volume rendering using mapreduce," in Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, pp. 841-848, ACM, 2010.

Appendix 1 – Abbreviations and acronyms

AWS	Amazon Web Services
IaaS	Infrastructure as a Service
NaaS	Network as a Service
PaaS	Platform as a Service
RDF	Resource Description Framework
SaaS	Software as a Service
SOC	Service Oriented Computing
VM	Virtual Machine
VPH	Virtual Physiological Human