Project Acronym: Open-DAI

Grant Agreement number: 297362

Project Title: Opening Data Architectures and Infrastructures of European Public Administrations

Work Package: System/Platform implementation

Deliverable Number: D4.2

Revision History

Revision Date	Author	Organisation	Description
28/9/2012	Luca Gioppo Mauro Campo Gianpiero Ardissono	CSI-Piemonte CSI-Piemonte CSI-Piemonte	Final version

Legal Discaimer

Statement of originality:
This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.

1 Introduction

The main purpose of this document is to explain how to install the hardware and software components of the Open-DAI Cloud infrastructure of the project architecture and how they work together. The document lists the technologies and configurations that have been chosen and the reasons for these choices.
For more details about the project you can read the following documents

- Call CIP-ICT-PSP-2011-5
- Document of the Work Open-DAI (297362)
- Open-DAI Description of General Architecture
- Open-DAI Data Virtualization Platform Installation Manual

2 Overview

The Open-DAI project deals with the following main issues:

- the opening of the huge amount of data stored in PA databases to the wide audience of potential users;
- the evolution toward an open architecture model for the PA information systems in order to overcome the monolithic and closed architecture models (silos);
- to facilitate software maintenance of existing silos, enabling PA to pace the evolution of legacy systems with Open Data initiatives.

The Open-DAI project aims to address the problems outlined above by testing the efficiency and added value of a SOA and Cloud-based architecture on several PA by:

- Implementing a data virtualization infrastructure deployed into a high availability infrastructure.
- Simplifying access to legacy vertical applications data, by providing a virtualized version of the data bases in the Cloud.
- Providing a new SOA data access layer, which could be combined in an appropriate manner in order to improve the products and services.
- Implementing the PA “open data” data hub, exposing it by using classic web services as well as other standard protocols.
- Assessing the business benefits for both PA and private organizations by developing new third-party added value services, focused on the following topics: transport and mobility, localization and geographic information, environment and pollution.

The project's functional servers have been hosted on a cloud virtualization infrastructure because working in cloud logic allows greater flexibility in provisioning needed computational capacity, lean management of the components and continuity of service.
Cloudstack and then its commercial evolution Citrix CloudPlatform has been chosen as the management tool of the virtualization infrastructure. Using this tool, described in detail in the following chapters, we have activated all the virtual machines needed to achieve the project objectives. The choice to use the commercial parallel solution of the selected cloud platform has been done for two main reasons:

- it was available in CSI-Piemonte's cloud laboratory infrastructure and it was fully compatible with the preliminary works developed on the free platform. It is still open source as Citrix has contributed its Cloud Platform to the Apache Foundation, the commercial part being the professional support by Citrix developers and specialists, therefore granting added professionalities to the projects' resources.
- What will be achieved on Cloud Platform could be in completely the same way replicated on the free platform because it is only added knowledge that has been purchased commercially.

It is one of the several Open Source business models, that of the “Product specialists”: companies that created, or maintain a specific software project, and use a Free Software license to distribute it. The main revenues are provided from services like training and consulting and follow the “best code here” and “best knowledge here” philosophy. It leverages the assumption, commonly held, that the most knowledgeable experts on a software are those that have developed it, and this way can provide services with a limited marketing effort, by leveraging the free redistribution of the code. The downside of the model is that there is a limited barrier of entry for potential competitors, as the only investment that is needed is in the acquisition of specific skills and expertise on the software itself. Most activities revolve around training, consulting, installation and configuration support, custom development and maintenance.

The safe access to existing data from legacy databases hosted in partners datacenters has been implemented with dedicated VPNs or private accesses.

3 Hardware infrastructure

This chapter describes the hardware components used to build the Open-DAI infrastructure, the charts and diagrams are at the end of the document.
In detail this chapter describes

- the cloud infrastructure management servers,
- the virtual machines' management hosts,
- the storage systems,
- the networks connecting all components,
- the WAN interconnections between the different partners,
- the security tools involved and
- the performance achieved in the early stages of implementation.

3.1 Servers

Open-DAI Project's infrastructure in the central Italian site consists of four physical servers dedicated to different tasks. The Cloud infrastructure has been deployed on two main virtualization servers that host all system's virtual machines (i.e. the cloud management node and all additional systems) in a duplicated virtual environment that allows prompt restoration of any system in case of a guest's failure. Two additional servers provide the hypervisor infrastructure based on a two nodes cluster.
Due to compatibility and Cloud Platform commercial support requirements the Centos 6.2 and RedHat 6.2 operative systems have been installed (and kept updated) on the servers.
The following table shows the main characteristics and task of the servers.

Open-DAI Project servers		
Server name	**Characteristic**	**Description**
csimgmtf1	***Model***	Dell PowerEdge R610
	Processors	4 CPU 6 Cores
	ram	72G RAM
	Disks configuration	2 disks 450G RAID 1
	NIC	4 NIC 10/100/1000 e 2 NIC 10G
	OS	Centos ver. 6.2
	Main task	Cloud management and base services
csimgmtf2	***Model***	Dell PowerEdge R610
	Processors	4 CPU 6 Cores
	ram	72G RAM
	Disks configuration	2 disks 450G RAID 1
	NIC	4 NIC 10/100/1000 e 2 NIC 10G
	OS	Centos ver. 6.2
	Main task	Cloud management and base services
kvmeudai1	***Model***	Dell PowerEdge R620
	Processors	4 CPU 8 Cores
	ram	96G RAM
	Disks configuration	2 disks 450G RAID 1
	Network	6 NIC 10/100/1000

	OS	RedHat ver. 6.2
	Main task	KVM hypervisor for VMs
kvmeudai2	***Model***	Dell PowerEdge R620
	Processors	4 CPU 8 Cores
	ram	96G RAM
	Disks configuration	2 disks 450G RAID 1
	Network	6 NIC 10/100/1000
	OS	RedHat ver. 6.2
	Main task	KVM hypervisor for VMs

The physical servers csimgmtf1 and csimgmtf2 are configured in QEMU-KVM pool mode sharing a filesystem via NFS protocol. The following virtual machines are hosted on these servers:

Virtual Machines running on csimgmtf1 and csimgmnt2 the KVM cluster		
Server name	**Characteristic**	**Description**
LDAP-Server	***OS***	Centos 5.8
	Processors	1 CPU
	Ram	1024 MB
	Disk configuration	10 GB
	Network	2 NIC
	Main task	directory information service
LDAP-Server-2	***OS***	Centos 5.8
	Processors	1 CPU
	Ram	512 MB
	Disk configuration	8 GB
	Network	2 NIC
	Main task	directory information service
CloudMgmt-ODAI	***OS***	Centos 6.3
	Processors	2 CPU
	Ram	4096 MB
	Disk configuration	60 GB
	Network	2 NIC
	Main task	CloudPlatform management
Munisquid	***OS***	Centos 5.8
	Processors	1 CPU
	Ram	2048 MB
	Disk configuration	16 GB
	Network	2 NIC
	Main task	http proxy, socks server (dante)
Pfsense	***OS***	FreeBSD/pfsense 2.0.1
	Processors	1 CPU
	Ram	2048 MB
	Disk configuration	8 GB
	Network	7 NIC
	Main task	Firewalling and vpns
MailServer	***OS***	Debian 5.0
	Processors	1 CPU
	Ram	1024 MB
	Disk configuration	10 GB

	Network	2 NIC
	Main task	e-Mail service
MailServer-2	***OS***	Debian 5.0
	Processors	1
	Ram	512 MB
	Disk configuration	8 GB
	Network	2 NIC
	Main task	e-Mail service
Dnsmicrosoft	***OS***	Windows 2003 Server
	Processors	1 CPU
	Ram	1024 MB
	Disk configuration	10 GB
	Network	2 NIC
	Main task	Domain name service
Dnsmicrosoft2	***OS***	Windows 2003 Server
	Processors	1 CPU
	Ram	1024 MB
	Disk configuration	10 GB
	Network	2 NIC
	Main task	Domain name service
Vulnserv	***OS***	Ubuntu 10.04
	Processors	2 CPU
	Ram	2048 MB
	Disk configuration	15 GB
	Network	2 NIC
	Main task	Vulterabilities assessment

The virtual machines listed in the previous table provide availability of all the base services to manage the cloud architecture, such as firewalling, mailing, authentication/authorization, address resolution, vulnerability management.
The **kvmeudai1**, **kvmeudai2** and **kvmeudai3** physical servers are controlled by CloudPlatform with an agent installed and configured on board. These servers act as a virtualization cluster that hosts the virtual machines of the Open-DAI project. These servers have RedHat 6.2 installed as OS because this is a CITRIX support requirement. KVM has been chosen as hypervisor because it's open source, it supports native virtualization on processors and it is directly developed by Red Hat Group.

The following picture shows the whole cloud infrastructure with the two zones (Italian and Turkish) and the logical position of the above described servers.

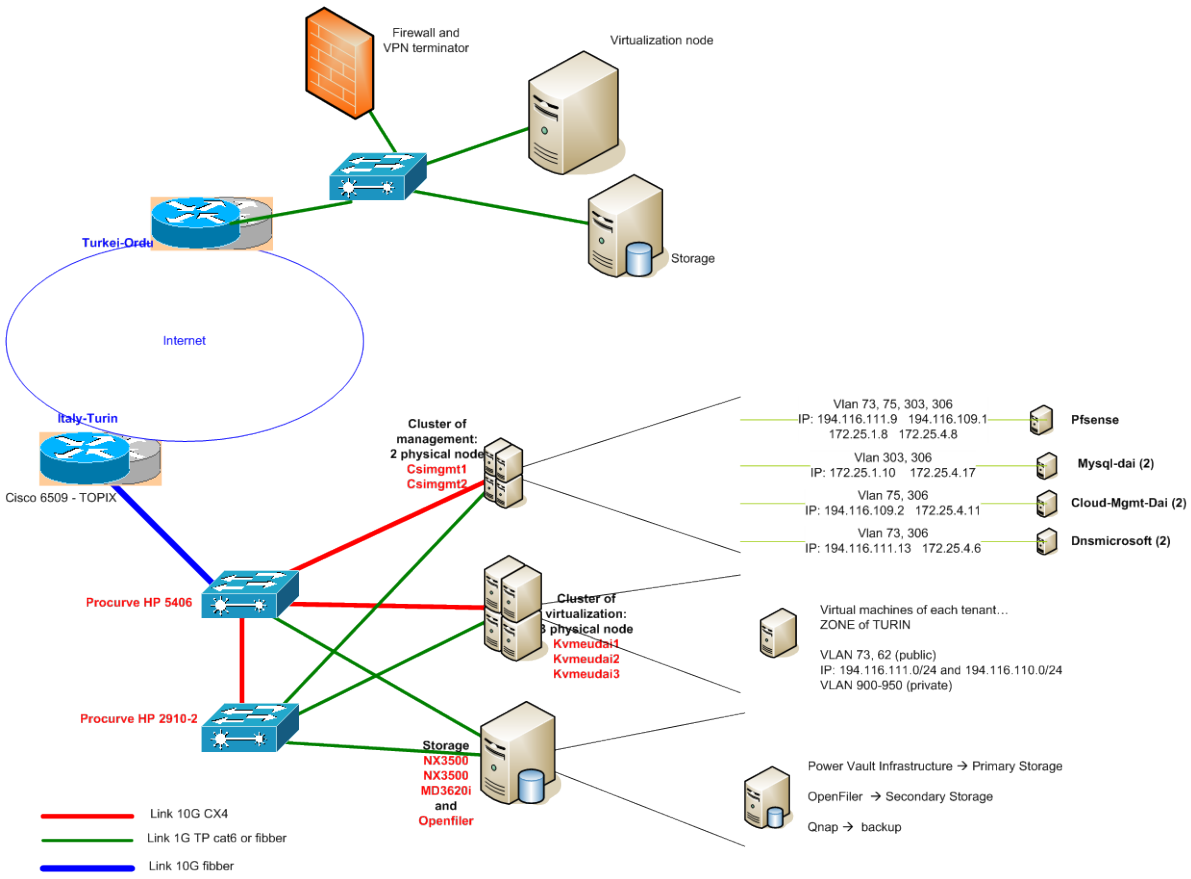

3.2 Storage

The storage plays an essential role in this IT project and it has to be highly reliable and efficient.
In OpenDAI cloud environment the selection of the storage system solutions was based on the following considerations:

1) Massive scalability and elasticity: the cloud storage solution should be able to scale to accommodate both growing (scalability) and decreasing amounts (elasticity). This feature allows users to choose and change in time their data requirements.
2) Different level of performance and reliability are needed in OpenDai storage systems due to the different roles: the Primary storage system supplies the disk space for the guest VM's, therefore this system needs the best performance parameters, high reliability and previous point characteristics; the Secondary storage system is less dinamically used but stores large quantities with a growing trend, therefore it needs reliability, scalability and throughput when needed; the backup system has a very predictable and planned use, therefore its main characteristics are scalability and reliability.
3) Different solutions, both software based and appliances have been selected based on what already available (the appliances Qnap and PowerVault that presented optimal characteristics for Primary and backup systems) and what needed with the given specifications (the secondary storage, set up with the leading Open Source NAS solution in its two-nodes cluster configuration, OpenFiler).

The following table lists the storage characteristics and the technology choices made to implement the requirements.

Storage requirements		
Requirement	**Description**	**Space (TB)**
Virtual machine images	The virtual machines listed in table 1 use a shared filesystem. Both the physical management servers of the Qemu-KVM pool can access this filesystem where the disk images of the virtual machines are stored. The images are saved in qcow2 format, the standard format for the KVM hypervisor.	1,5TB
Primary Storage	"Primary Storage" (PS) is the name CloudPlatform gives to its filesystems containing the disks of the virtual machines on KVM servers. The Open-DAI primary storage is statically mounted on each server in the KVM cluster. It is assumed that the project would include more than 20 user machines for each partner for a total of 80 to 100 virtual machines instanced, each with a local disk of average 20GB. Under this hypothesis the project would need at least 2TB. The space is divided in an Italian "zone" and a Turkish "zone". The Italian zone will host virtual machines for the Italian, Swedish and Spanish partners. The Turkish zone will host its own machines.	2,4TB
Secondary Storage	"Secondary Storage" (SS) is the name CloudPlatform gives to the disk space where it keeps distributions ISOs to install user VMs, templates to instance preconfigured user VMs, snapshots to backup running VMs and the templates of service VMs (virtual routers and secondary storage machine). For the SS the same assumptions and the same choices have been made as for sizing the PS	2,4TB
Backup	Data is backupped on a remote infrastructure located separately from Virtual machine images, primary storage and secondary storage. The saved data comprises support VM images, their configurations and user VM images.	6,8TB

Virtual machine images, Primary Storage and Secondary Storage are hosted in distinct physical storage systems for debugging and performance reasons.

Storage systems map		
Requirement	**hw/sw**	**Configuration**
Virtual machine images	Linux Centos 6.2 server with NFS	• 2 drives of 900GB mirrored in a RAID 1 cluster, for the operative system • 4 drives of 900GB in RAID 10 for a 1.8TB filesytem to export via NFS
Primary Storage	Dell PowerVault NX3500 and Dell MD3620i	• RAID 50 • 2 logical volume of 2.5TB each, only one is dedicated to the project
Secondary Storage	OpenFiler 2.9.9 server	• 2 drives of 900GB mirrored in a RAID 1 cluster, for the operative system • 6 drives of 900GB in RAID 10 for a 2.4TB filesytem to export via NFS
Backup	Qnap TS-EC1279U-RP storage system	• 8TB with and array 16 drives of 2TB each in mirror mode

3.3 Network

The Open-DAI cloud-based architecture comprises a physical network's and an overlayed virtual one's infrastructure.
The physical network provides a performing Internet access through a 10G direct peering with an Italian Internet exchange, TOPIX, directly connected to the project's network rack. Both the public network infrastructure and the private one is equipped with network devices (up to the servers' network boards) with full 10G support.
IP address ranges offer three full public C classes and part of a fourth either for direct connections or for protected ones. On the private side, standard private ip address ranges compliant with RFC 1918.
All these abundant resources have been made available by CSI-Piemonte because it is and it has been in the past decade the engine of the regional ICT development enacting the regional, national and European policies for:

- reducing digital divide. Therefore coordinating a huge joint private-public effort for developing ICT operators and providers and deploying network infrastructures both wide band in urban areas and for alternative connectivity in rural and mountain areas (satellite, wireless operators);
- bringing international level Internet peering services in the Piemonte region by supporting the birth of the Internet Exchange **TOP-IX (TOrino Piemonte Internet eXchange),** a non-profit consortium set up in 2002 with the aim of creating and managing a NAP (Neutral Access Point, otherwise known as Internet Exchange – IX) for the exchange of Internet traffic in North West Italy. The consortium has today 65 members including main national and international network operators;
- enabling and developing the e-Government services of the regional and local Public Administrations.

CSI-Piemonte has therefore accumulated a wealth of know-how, competences and resources that can be shared and made available for the projects it is involved in. In the present case, specifically through its Cloud Laboratory infrastructure where innovative solutions for the future of the networks of the regional PA's are tested and evaluated.

3.3.1 Wide Area Network (WAN)

Access to the public Internet is provided by means of a 10G fiber optic connection to the Italian north-western internet exchange TOP-IX.
This fiber optic is physically connected to the CSI datacenter where a TOP-IX's Cisco 5609 is housed, directly connected to the primary switch of the Open-DAI project (HP 5406).
The public interconnection supports all network traffic to and from the cloud's hosted VMs. On this connection goes also all traffic to and from the Turkish “zone”, protected by a VPN that allows the secure connection between remote cloud's zones.

3.3.2 LAN – Italian Zone

The Italian zone hosts the Italian, Swedish and Spanish domains on the project's dedicated cluster.
The KVM nodes where the project's virtual machines are hosted are equipped with three logical network connections (physically three couples of bonded network boards).
One connection is devoted to the storage VLAN, which hosts the primary storage systems; one connection hosts the public VLAN that allows Internet access (in and out) for the virtual machines; the third connection is dedicated for the cloud's internal services and traffic: this interface supports the “trunks” bearing the management VLAN, all internal networks' VLANs (see CloudPlatform documentation) and all public or specific VLANs on which specific network offerings are provided, i.e. for direct public IPs acquisition or for private IP addressing plans with specific traffic policies for legacy databases hosted elsewhere and accessible through virtual private networks (VPNs).

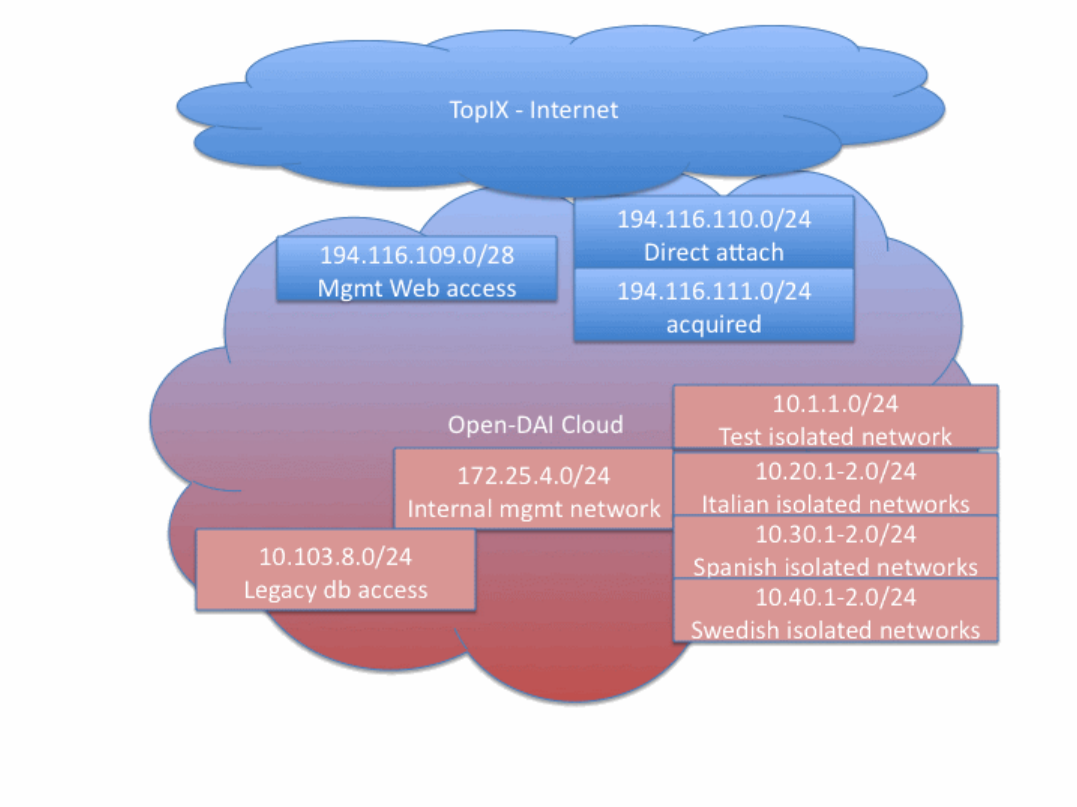

The following image shows the physical layout of networks and hosts of the Italian zone.

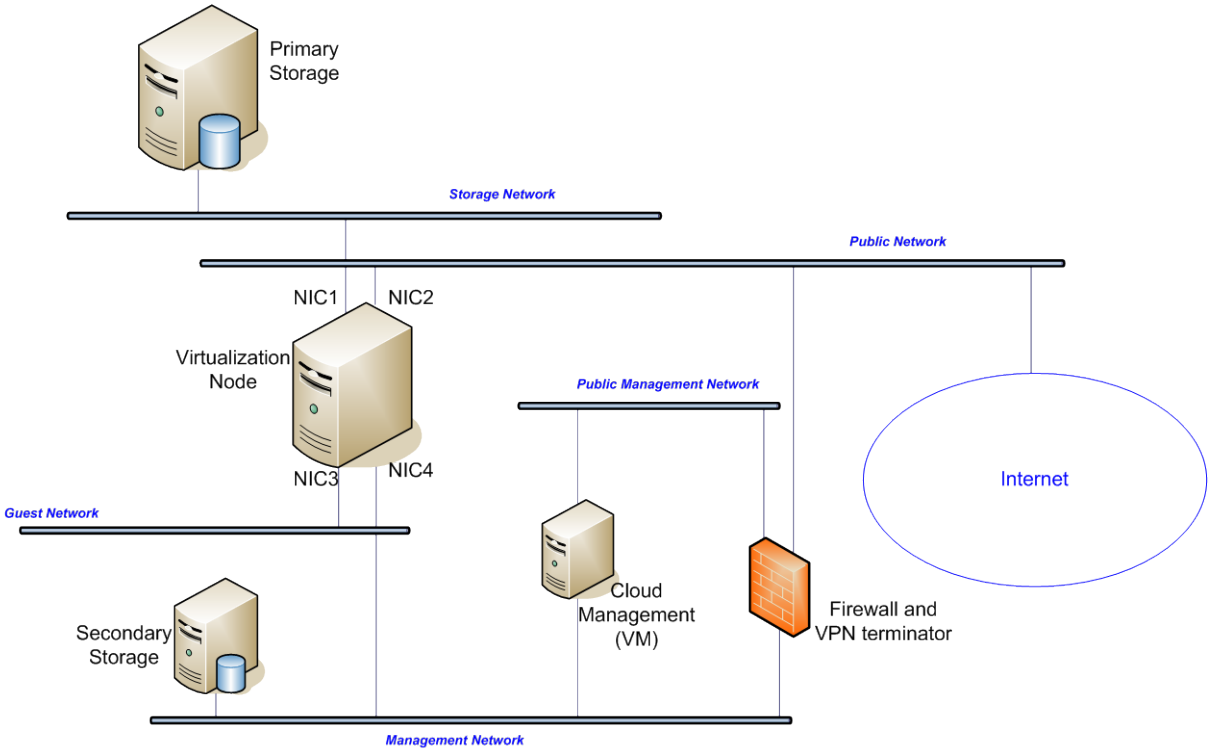

3.3.3 LAN – Turkish Zone

The Turkish zone has to be configured like the Italian one.
The management network is connected through a VPN between the two zones.
In the Turkish zone a single virtualization node will be deployed with a qemu/KVM hypervisor and the necessary primary and secondary storage systems. The technology selected for the network storage systems is Fiber Channel for best access performance.

4 Software tools

CloudPlatform 3.0.4. is the main tool the cloud is built with. For more details about CloudPlatform see the references.

Both in the Turkish and in the Italian zone the choice of the hypervisor has been the open source solution Qemu/KVM on the RedHat/CentOS Linux distribution.

The operative system on the management machine is Centos 6.2, while the virtualization cluster uses RedHat 6.2.
Particular care has been paid to the choice and configuration of the firewall isolating the network management of the cloud and protecting the management interface. CloudPlatform suggests a Juniper appliance, but in order to promote open source software and to virtualize the perimetral security component, the open source PfSense application version 2.0.1 (www.pfsense.org) has been selected. PfSense is a free, open source customized distribution of FreeBSD tailored for use as a firewall and router. In addition to being a powerful, flexible firewalling and routing platform, it includes a long list of related features and a package system allowing further expandability without adding bloat and potential security vulnerabilities to the base distribution.
PfSense is also used as the italian zone VPN end-point with the Turkish zone where the tunnel is terminated on a Zywall USG 2000 firewall, the existing security appliance available at the Turkish Ordu site.

5 Additional systems and services

This chapter describes the integrated systems that complement the IaaS architecture offered by CloudPlatform. They are the backup policy, the performance and fault monitoring system, the vulnerability assessment of systems exposed to public networks and the naming service for the users' machines.
The additional services listed are functional to the deployment of the infrastructure supporting the Cloud and, while it could have been possible to deploy them inside the Cloud itself, it has been deemed useful and a logical choice to supply them with external systems. They are a permanent feature of the system's architecture, not directly connected to the cloud-based services.

5.1 Backup

The backup policy has been defined based on the characteristics of the selected tool and on the different levels of importance and criticality of the systems involved. Different backup rules have been defined for the infrastructural systems and the hosted ones.
The critical systems are:

- management host (does not vary in time except for updates)
- management db for the cloud' infrastructure (does not vary much in time)
- network configuration of the virtualization nodes (not changing)
- application systems hosted on the hypervisors (may vary significantly in time)

The management host and the network configuration of the hypervisors are saved on the Qnap backup system and they are updated manually when (seldom) a change occurs. The application hosts are backed up daily. The db with the cloud's infrastructure configuration is frozen as a dump file daily and then it is saved on the Qnap backup storage.
The application hosts have been configured to be saved daily through CloudPlatform's internal snapshot mechanism that stores copies of the virtual machines on the secondary storage. The snapshots are then also backed up on the Qnap system. This policy would allow the prompt restoration of any virtual machine from inside the Cloud and its recovery even in case of problems with the cloud-based infrastructure. The two last snapshots are kept available. Alas, the snapshot mechanism on the present version of CloudPlatform is not working properly: the “rotation” of the snapshots that should delete progressively older ones to keep available the two more recent is broken and, once a snapshot is created it won't be deleted. While waiting for this bug to be solved by the commercial support, the application hosts are backed up manually.
The daily backup operations also involve the most important additional systems such as the pfSense firewall protecting the management host and providing the VPN infrastructure, the Racktables db where all information about the physical cabling and network connections for all physical hosts are stored and the vulnerabilities' monitoring system (openvas).

5.2 Monitoring

The project has a physical machine with all the monitoring tools on board. This choice has been made in order to have a dedicated physical machine separated from the systems (both virtual and physical) being monitored. The monitoring tools used are comprised in the munin 2.0.6 platform (www.munin-monitoring.org).
The log files are collected from all the physical machines of the project, from all the operative systems of both virtual and real machines and from all the support applications.
All the major problems and all the high priority warnings are automatically sent by e-mail to a support team. The support team is organized to check the logs regularly and act ASAP in case of problem.

5.3 Vulnerability Assessment

The project includes a number of machines that can be contacted directly on the internet.
The cloud system is extremely flexible and try to activate only the necessary services.

However having the opportunity to acquire public addresses, essentially unfiltered by the cloud platform, administrators must enable all the necessary countermeasures to prevent attacks from the INTERNET.
In order to harden the project systems, there is OpenVas 1.0 (www.openvas.org) that periodically analyzes the vulnerabilities on the machine exposed to the INTERNET and produces detailed reports giving indication of possible countermeasures.
The reports are sent by e-mail to the administrators of the machines with found vulnerabilities.

5.4 DNS

Two DNS servers on Microsoft systems support the fully qualified domain naming of all hosts in the physical and virtual infrastructure, they will be soon substituted by 2 open source Bind servers.
The DNS servers manage the public domain of the project dedicated to the virtual infrastructure, ***dev-cloud-open-dai.eu*** enabling the proper addressing from the Internet of machines, services and complex systems. They also manage the internal name and addressing system on the domains ***cloudlabcsi.eu*** and ***cloudlabcsi.local*** used for consistent address resolution in the management and monitoring systems and for virtual machines that do not need to be exposed on the public network.

6 Installation and Configuration of the Cloud

Here there are the steps to install and configure CloudPlatform infrastructure for the OpenDAI project.

6.1.1 Checks

The first thing to do installing the cloud is to check the resources available, then you can choose and design how to allocate them.
The tables “Storage requirements” and “Open-DAI Project servers” show the storage infrastructure and the servers bought for the Open-DAI Project. 300 public IP addresses (divided in 2 different networks) have been allocated for the requirements of access of the virtual machines.

6.1.2 Installation of the CloudPlatform (3.0.4) package

CloudPlatform comes as a software package available for download on Citrix's website. There is a single package containing both the server software to install onto the management server and the agent software to install on the virtualization nodes.
This chapter describes the main steps to install CloudPlatform and for the creation of the virtual machines for the Open-DAI Project. See ***“CloudPlatform 3.0.3 - 3.0.5 Installation Guide”*** for details.
The first things to check are that all of the components have a FQDN, selinux is turned off, the systems time is synchronized with NTP, RPC and NFS are installed.
Although not required during normal operation, the installation procedure requires the management machine to be temporarily connected to the secondary storage to install the system virtual machines from templates downloaded there. These system virtual machines are specific to the hypervisors used in the zone, so only the KVM template has been installed in the Italian Zone while in the Turkish Zone the VMWare one has been installed.
After installing the management software and the agent software, only the manager host ip has been written in the properties files of the virtualization nodes. All the other information are written by the manager itself while configuring.
Logging in web portal for the first time the manager realizes that the database is empty and starts a wizard to create the first zone, the pod and the cluster with the first KVM node.
The first zone has been created using the web user interface on the port 8080. Then following the procedure at http://docs.cloudstack.org/Knowledge_Base/Enable_HTTPS_for_CloudStack_Web_Interface the user interface has been migrated in https.

6.1.3 Early (Cloud environment) setup

6.1.3.1 Operative systems

While CloudPlatform is an open-source platform that can be set up on different Linux distributions, to warrant its commercial support Citrix requires RedHat 6.2 for the KVM virtualization nodes.
The management server has been set up with Centos 6.2.
It is important to note here that the virtualization nodes are not going to be connected to the Internet after the initial setup, therefore all system software upgrades have to be performed during installation. Further updates will have to be deployed by indirect or time limited and scheduled Internet connectivity.

6.1.3.2 Initial Cloud configuration

In the following paragraphs the description will reference or detail specific deviations from the standard configuration described in the ***“CloudPlatform 3.0.3 - 3.0.5 Installation Guide”***.
It is here nonetheless necessary to understand the basic mandatory steps to setup the initial cloud configuration.

To provision the cloud infrastructure for the first time it is necessary to follow precise steps. This is the most delicate phase that has to be accomplished without errors because the Cloud is a complex system,

quite robust once set up but most sensitive to the lack of an ordinate and complete initial setup. Any error in this phase results in the need to restart the process from scratch:

1. Add zones and pods: a zone is the largest organizational unit within a CloudPlatform deployment. A zone typically corresponds to a single datacenter, although it is permissible to have multiple zones in a datacenter. The benefit of organizing infrastructure into

zones is to provide physical isolation and redundancy. A zone consists of:

- One or more pods. A pod often represents a single rack. Hosts in the same pod are in the same subnet.

A pod is the second-largest organizational unit within a CloudPlatform deployment. Pods are contained within zones. Each zone can contain one or more pods.

A pod consists of one or more clusters of hosts and one or more primary storage servers.

- Secondary storage, which is shared by all the pods in the zone.

2. Configure the physical network defining all networks, virtual and physical that will support the Cloud's infrastructure. A basic configuration needs at least three categories, linked to the use type, Guest, Management and Storage networks. Advanced configurations may add new categories such as Public, DMZ, etc.
3. Add clusters. A cluster provides a way to group hosts. To be precise, a cluster is a XenServer server pool, a set of KVM servers, or a VMware cluster preconfigured in vCenter. The hosts in a cluster all have identical hardware, run the same hypervisor, are on the same subnet, and access the same shared primary storage. Virtual machine instances (VMs) can be live-migrated from one host to another within the same cluster, without interrupting service to the user.

A cluster is the third-largest organizational unit within a CloudPlatform deployment. Clusters are contained within pods, and pods are contained within zones. The size of the cluster is limited by the underlying hypervisor.

A cluster consists of one or more hosts and one or more primary storage servers.

4. Add hosts. A host is a single computer. Hosts provide the computing resources that run the guest virtual machines. Each host has hypervisor software installed on it to manage the guest VMs. For example, a Linux KVM-enabled server, a Citrix XenServer server, and an ESXi server are hosts.

The host is the smallest organizational unit within a CloudPlatform deployment. Hosts are contained within clusters. Hosts in a CloudPlatform deployment

- provide the CPU, memory, storage, and networking resources needed to host the virtual machines,
- interconnect using a high bandwidth TCP/IP network and connect to the Internet,
- may reside in multiple data centers across different geographic locations,
- may have different capacities (different CPU speeds, different amounts of RAM, etc.), although the hosts within a cluster must all be homogeneous

- may be designated for use only as a restart location for HA-enabled virtual machines.

Hosts may be designated as dedicated HA restart nodes only if the Dedicated HA Hosts feature is enabled in the global configuration of the cloud.

CloudPlatform automatically detects the amount of CPU and memory resources provided by the hosts. Hosts are not visible to the end user. An end user cannot determine which host their guest has been assigned to.

5. Add primary storage and secondary storage. Primary storage is associated with a cluster, and it stores the disk volumes for all the VMs running on hosts in that cluster. You can add multiple primary storage servers to a cluster. ***At least one is required***. It is typically located close to the hosts for increased performance. CloudPlatform manages the allocation of guest virtual disks to particular primary storage devices.

CloudPlatform is designed to work with all standards-compliant iSCSI and NFS servers that are supported by the underlying hypervisor. Secondary storage is associated with a zone, and it stores the following:

- Templates – OS images that can be used to boot VMs and can include additional configuration information, such as installed applications
- ISO images – disc images containing data or bootable media for operating systems
- Disk volume snapshots – saved copies of VM data which can be used for data recovery or to create new templates.

The items in secondary storage are available to all hosts in the zone.

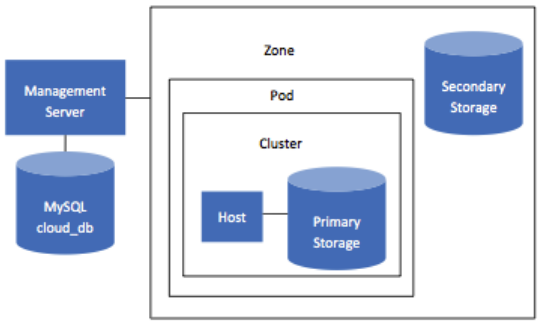

Conceptual view of a basic deployment

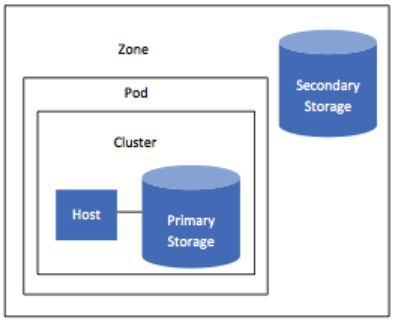

A simple zone

6.1.3.3 Zone specifics

In OpenDAI two Zones have been setup: an Italian and a Turkish one. These two geographically separated physical parts of the Cloud have been joined through an IPSec Virtual Private Network in a single logical and managed Cloud architecture. The simplest structure has been chosen for each zone regarding the hosts, i.e. one pod and one cluster per Zone, with two KVM hosts in the Italian cluster and one in the Turkish one.

6.1.3.4 Network configuration

Network setup implies configuring the network switches to support the traffic for the hypervisor nodes and for the management VM. Also, the hypervisors need to be properly configured to host all the virtual networks needed by the guests VMs. The networks' address spaces and routing and switching infrastructure is shown in the images in **paragraph 3.3.2**.
The network categories deployed in the Cloud for the OpenDAI project are Management, Storage, Guest and Public. These categories correspond, on the hypervisor nodes, to specific configurations of the network interfaces. The configurations of the Italian KVM hypervisors' network interfaces are shown below. Each host has to be accurately configured in the same way to warrant interoperability and the possibility for the management node to move the guest VMs across the virtualization environment.

ifcfg-brprivate
DEVICE=brprivate
ONBOOT=yes
IPADDR=172.25.4.45
BOOTPROTO=none

NETMASK=255.255.255.0
DNS2=172.25.4.19
TYPE=Bridge
GATEWAY=172.25.4.8
DNS1=172.25.4.6
IPV6INIT=no
USERCTL=no

ifcfg-brpublic
DEVICE=brpublic
ONBOOT=yes
BOOTPROTO=none
TYPE=Bridge
IPV6INIT=no
USERCTL=no

ifcfg-storage
DEVICE=em3
HWADDR=D4:AE:52:A0:FB:D3
NM_CONTROLLED=no
ONBOOT=yes
BOOTPROTO=none
TYPE=Ethernet
IPV6INIT=no
USERCTL=no
IPADDR=192.168.1.35
NETMASK=255.255.255.0

ifcfg-em1
DEVICE=em1
HWADDR=d4:ae:52:a0:fb:d1
NM_CONTROLLED=no
ONBOOT=yes
TYPE=Ethernet
IPV6INIT=no
USERCTL=no
BRIDGE=brprivate

ifcfg-em4
DEVICE=em4
HWADDR=D4:AE:52:A0:FB:D4
NM_CONTROLLED=no
ONBOOT=yes
BOOTPROTO=none
TYPE=Ethernet
IPV6INIT=no
USERCTL=no
BRIDGE=brpublic

Both the virtualization nodes and the management server need to access a private network that should allow the management traffic between them and with the secondary storage system. Moreover traffic has to be allowed to and from remote CloudPlatform's zones.

The management server has to be reachable from the Internet only on a specific network port where an https service for the web interface is offered. This network, firewall protected, must be isolated from all other public networks used, included the one used to provide the address ranges for the users' VMs.

The virtualization nodes should be able to access public networks, both for “acquired” (i.e. not directly assigned to VMs' virtual network interfaces) addresses and for “directly attached” ones. The nodes need to be connected to the isolated networks of the users' VMs as well.

For all these purposes the physical and virtual network infrastructure has to support 802.1Q VLAN tagging and trunking. The network design phase has taken care of defining the VLANs' ranges for public, private and isolated VLANs and their virtualization rules.

7 Additional Configurations

In this chapter there are all the configurations applied after the initial CloudPlatform installation and setup, i.e. the global settings used and the configurations given to template and ISO images made available to users.

7.1 Service offering

Service offerings in CloudPlatform define the virtual hardware that user will be able to choose. They have been defined by the CloudStack administrator and provide a choice of CPU speed, number of CPUs, RAM size, tags on the root disk, and other choices.
The following tables show the “Compute offering” and the “Network offering” added to the installation default offering of OpenDAI's CloudPlatform installation.
The compute offering's name is descriptive and specifies the number of cpus, RAM size available and presence of the High Availability option.

Compute offering	
2cpu2ram	2cpu2ram
2cpu2ramHA	2cpu2ramHA
2cpuH4ram	2cpuH4ram
1cpu1ram	1 cpu (1300) 1 Giga no HA
1cpu1ramHA	1 cpu (1300) 1 Giga HA
2cpuH2ramHA	2cpuH2ramHA
2cpuH2ram	2cpuH2ram
3cpuH6ramHA	3cpuH6ramHA
3cpuH6ram	3cpuH6ram
2cpuH4ramHA	2cpuH4ramHA

Network offering		
NetOff-1	Description	Network Offering isolato
	State	Enabled
	Guest Type	Isolated
	Availability	Optional
	Created by system	No
	Specify VLAN	No
	Specify IP ranges	No
	Conserve mode	No
	Network Rate	200 Mb/s
	Traffic Type	Guest
	Supported Services	Vpn, PortForwarding, SourceNat, Dns, Firewall, StaticNat, Dhcp
	Service Capabilities	SupportedSourceNatTypes: peraccount, RedundantRouter: false, ElasticIp: false
	Tags	

Open-DAI-offering	Description	Network Offering for Open-DAI
	State	Enabled
	Guest Type	Isolated
	Availability	Optional
	Created by system	No
	Specify VLAN	No
	Specify IP ranges	No
	Conserve mode	No
	Network Rate	200 Mb/s
	Traffic Type	Guest
	Supported Services	UserData, Lb, Vpn, PortForwarding, SourceNat, Dns, Firewall, StaticNat, Dhcp
	Service Capabilities	SupportedLBIsolation: dedicated, ElasticLb: false, SupportedSourceNatTypes: peraccount, RedundantRouter: false, ElasticIp: false
	Tags	

7.2 Available ISOs

The ISO images available to generate standard virtual machines are only for the CentOS Linux distribution:

- The Centos 6.3 complete install that offers the full scope of applications and services of the distribution and a GUI
- The Centos 6.3 minimal install that is the essential operating system with basic services

The choice to not use in OpenDAI more operating systems is due to the necessity of developing and adapting the applications and services on a uniform technological base for ease of debugging, interoperability and documentation.

7.3 Available Templates

The following table shows the list of the templates prepared to create standard virtual machines with pre-installad applications and services.

Available templates
COS-DAI-TMPL(BASE)
COS-DAI-TMPL-(FULL)
WIP-4
Autoinstall
BareMinimal
open-dai-bootstrap

7.4 Global settings

This is the table of the configurations listing Cloudplatform's global settings.

host	The ip address of management server	172.25.4.36
management.network.cidr	The cidr of management server network	172.25.0.0/16
max.account.public.ips	The default maximum number of public IPs that can be consumed by an account	50
max.account.snapshots	The default maximum number of snapshots that can be created for an account	2000
max.account.user.vms	The default maximum number of user VMs that can be deployed for an account	50
max.account.volumes	The default maximum number of volumes that can be created for an account	50

7.5 Additional Remote Zone setup

To add the remote Turkish Zone to the openDAI Cloud the following requisites are mandatory:

- a set of public networks or subnetworks (not single sparse public IP addresses);
- an L2/L3 network infrastructure with VLAN (IEEE 802.1Q) support;
- a storage infrastructure with at least two different physical/logical volumes;
- a virtualization infrastructure with hypervisors compatible with CloudPlatform, in this case the choice has been again KVM;
- a VPN terminator.

Zone creation happens by adding a new Cluster/Host (only one host in the Turkish cluster) managed by the Cloud. The configuration of the first Host of the new Zone has to comply with the usual requirements for a KVM host in CloudPlatform and with the following OpenDAI specific requirements:

- one NIC (Network Interface Card) has to be devoted to the public network (Internet) without firewall protection, configured as a bridge and without an IP address (like in the “brpublic” configuration of the interfaces of the Italian hosts;
- another NIC has to be configured for supporting all private traffic between guest VMs and the management traffic. Again, it will be a bridge (in this case like the “brprivate” interfaces of the Italian hosts.

7.6 Other configurations

Only the GUI timeout has been changed from 5 minutes to 300, because the default time was too short for debugging purposes and required repeated logins during extended sessions of configuration and testing of the Cloud's infrastructures and services.