

IT University
of Copenhagen

Installing Ubuntu Enterprise Cloud in a Physical Environment

Supplement (2) to

“Guidelines for Building a Private Cloud
Infrastructure”

Zoran Pantić and Muhammad Ali Babar

Tech Report TR-155

ISBN: 978-7949-256-1

IT University of Copenhagen, Denmark, 2012

Summary

This document contains the supplemental material to “Guidelines for Building a Private Cloud Infrastructure.” This supplemental material provides guidance on how to install Ubuntu Enterprise Cloud in a physical environment. The purpose of this document is to provide a practical, step-by-step, detailed guide on how to pre-configure and install the machines and network. For more detailed description of the steps, a reader is advised to refer to another supplemental book named “*Installing and Scaling out Ubuntu Enterprise Cloud in Virtual Environment*.” There are a few more details, accompanied with screenshots.

The material included in this supplemental document is based on the installation that was performed in a physical environment based on HP ProLiant DL380 G4¹ servers, 100 Mbit/s switches, and a firewall/router that isolated the whole solution. The cloud installation was performed using the (for the time being) newest version of the UEC, based on the Ubuntu 11.04. The installation was performed using “CD Install” method² (otherwise, there was also a “Package Install” method³). This material in this document is based on Eucalyptus version installed is 2.0.2.

¹ http://h18000.www1.hp.com/products/quickspecs/12028_div/12028_div.html

² <https://help.ubuntu.com/community/UEC/CDInstall>

³ <https://help.ubuntu.com/community/UEC/PackageInstall>

Contents

| | |
|--|----|
| Summary..... | 1 |
| Installation of the cloud..... | 3 |
| Machine configuration – initial setup..... | 3 |
| Network configuration | 4 |
| Pre-configuring the servers | 5 |
| Installation of the servers..... | 5 |
| Installation of the firewall server | 5 |
| Installation of the administrative workstation..... | 7 |
| Installation of the first server (CLC, WS3, CC, SC)..... | 8 |
| Installation of second server (NC) | 14 |
| Verifying the cloud installation..... | 16 |
| Running the cloud..... | 16 |
| Initial steps | 16 |
| Cloud web interface | 17 |
| Downloading user credentials | 17 |
| Running an image using command line..... | 17 |
| Photos of the hardware infrastructure | 19 |

Installation of the cloud

The cloud installation described here is performed using the (for the time being) newest version of the UEC, based on the Ubuntu 11.04 (“ubuntu-11.04-server-amd64.iso” CD image). To get the most of the physical server, we choose the 64-bit architecture, although a 32-bit version is also available, if implementing on the older, x86 based architectures.

A “long-term support” (LTS) release is also available (10.04) if more stability needed, or for larger deployments.

Machine configuration – initial setup

Initial UEC setup is performed on 2 servers – first server being cloud front end and hosting the cluster related roles, and the second server being the compute node.

Also, a firewall server and an administrative workstation are installed, for the purpose of isolating the whole environment, and ease of control.

Here is how the machines should be configured:

UECserver01

Roles: CLC, WS3, CC, SC
HP ProLiant DL380 G4
Processor: 2 x 64 bit Intel Xeon 3.0 GHz, not VT enabled
RAM: 6 GB PC2-3200R 400MHz DDR2
Disk: integrated Smart Array 6i Ultra320 Array Controller, 3 spindles arranged in RAID 5
Network: public (eth0), private (eth1) - Embedded NC7782 PCI-X Gigabit Server Adapter

UECserver02

Roles: NC
HP ProLiant DL380 G4
Processor: 2 x 64 bit Intel Xeon 3.0 GHz, not VT enabled
RAM: 12 GB PC2-3200R 400MHz DDR2
Disk: integrated Smart Array 6i Ultra320 Array Controller, 6 spindles arranged in RAID10
Network: private (eth0) - Embedded NC7782 PCI-X Gigabit Server Adapter

UECcloudFW

Roles: firewall/router
HP ProLiant DL380 G4
Processor: 2 x 64 bit Intel Xeon 3.0 GHz, not VT enabled
RAM: 4 GB PC2-3200R 400MHz DDR2
Disk: integrated Smart Array 6i Ultra320 Array Controller, 3 spindles arranged in RAID1 (with spare disk)
Network: WAN (bge0), public (bge1) - Embedded NC7782 PCI-X Gigabit Server Adapter

UECworkstation01

Roles: administrative workstation
HP Compaq workstation, 1 processor, 1 GB RAM, 80 GB IDE hard disk
Network: private (eth0)

Here is the overview in a table:

| Machine name | Machine role | Processors | RAM (GB) | Hard disk (GB) | Eth0 | Eth1 |
|------------------|-------------------|------------|----------|----------------|---------|---------|
| UECserver01 | CLC, WS3, CC, SC | 2 | 6 | 293 | public | private |
| UECserver02 | NC | 2 | 12 | 218 | private | - |
| UECcloudFW | pfSense firewall | 2 | 4 | 72 | WAN | public |
| UECworkstation01 | Admin workstation | 1 | 1 | 80 | public | - |

Network configuration

All networks run on a class \C private network range (subnet mask 255.255.255.0).

The networking in our setup is implemented with three networks – private and public. The public network would be accessible to users and administrators from the “outside world”, and the private network is reserved for the backend communication, and is private for the cloud. The “outside world” network is on the WAN interface.

The solution can be implemented using only two networks, public and private, but then public IP addresses would have to be obtained for all the servers hosting CLC, WS3, CC and SC roles. That can be coordinated with the IT Department.

Here is the overview in a table:

| Name | Network | Subnet mask | Gateway | DNS1 | DNS2 |
|---------|---------------|---------------|---------------|-----------------|-----------------|
| Public | 192.168.1.0 | 255.255.255.0 | 192.168.1.1 | 130.225.127.212 | 130.225.127.213 |
| Private | 10.0.0.0 | 255.255.255.0 | 10.0.0.1 | 10.0.0.1 | - |
| WAN | 130.225.115.0 | 255.255.255.0 | 130.225.115.1 | - | - |

Public network:

- Range 192.168.1.1 - 49 is reserved for router, workstations, DHCP pool, etc.
- Range 192.168.1.50 - 99 is reserved for servers with cloud roles
- Range 192.168.1.100 - 192.168.1.254 is reserved for instances’ public addresses, possibly divided between several clusters, if present

Private network:

- Whole range 10.0.0.1 – 10.0.0.254 is reserved for servers with cloud roles

WAN network:

- Organization’s network, only one IP is needed, should be acquired from IT Department

Managed-NOVLAN networking mode is chosen, since it gives a variation of possibilities close to the ones that would be needed in production environments. NOVLAN is chosen simply for the reason that no V-LAN network components are available in the virtual environment. For more details on the 4 networking modes in UEC, please refer to the main document.

The public DNS servers are very connection specific, and should be obtained and configured in collaboration with you internet service provider, or internal IT department.

Pre-configuring the servers

Servers are preconfigured with 2 processors, RAM, hard disks, and redundant power supplies. Generally, for the local disk, RAID options should be configured on the controller level. That configuration is specific for each hardware supplier - please refer to the supplier's manuals. In our case, we used a special startup CD (Compaq SmartStart) delivered by HP with several configuring possibilities, among them disk array configuration utility.

RAID 0 (stripping) would give us best performance and most space, but no redundancy. RAID 1 (mirroring) is having a bit better read performance than just one disk, but gives redundancy. RAID 5 is giving better performance and redundancy (for missing just one disk). The best option if having local disks would be RAID 10 (that is, mirroring of the two stripped volumes) giving us both performance of RAID 0 and redundancy of RAID 1, but then we have to have at least 4 disks.

Cloud frontend server is configured with RAID 5 disk array configuration, since we are more interested in available space, than in performance. NC server is configured with RAID 10, since we assume that instances running on that server (being cached locally) would fit in the space available, and RAID 10 would give the necessary performance benefits.

We could also have chosen to separate the disk configuration, putting the operating system on the local disks (in some RAID configuration, preferably RAID 1 and up), and putting cloud specific storage on SAN (Storage Array Network), accessing it by either fiber net (optical, very quick but expensive) or iSCSI (using Ethernet infrastructure). That configuration would be preferred in the future, when scaling the cloud out.

Installation of the servers

For more details on installation of UEC, please refer to Ubuntu Documentation⁴.

Installation is performed using "ubuntu-11.04-server-amd64.iso" installation CD.

The installation could be divided in two parts: first part is a usual Ubuntu server installation, and second part is related to installing of cloud components.

As first step, we should install the firewall server, and the administrative workstation. Having that in place, we would install the UEC servers.

Installation of the firewall server

Installation is based on pfSense, which is a free, open source customized distribution of FreeBSD tailored for use as a firewall and router.

Fetch the installation media⁵ that should be burned on a CD. We used the stable version 1.2.3, with image "pfSense-1.2.3-RELEASE-LiveCD-Installer.iso.gz". The installation is 32-bit based, and that is the reason for giving the server 4 GB of RAM – 32-bit architecture cannot address more than 4 GB, and all memory space available above 4 GB is simply ignored.

⁴ <https://help.ubuntu.com/community/UEC>

⁵ <http://www.pfsense.org/mirror.php?section=downloads>

Mount the CD installation media, and boot from it. Choose to boot pfSense, and after booting and launching the init system, choose “I” for installer.

Accept the console video settings, and choose “Quick / easy install”. Choosing that, we are skipping configurations such as formatting, partitioning, and the system is installed automatically.

Choose “Symmetric multiprocessing kernel” if you have more than one processor, otherwise choose “Uniprocessor kernel”. After installation is done, the server reboots.

Note the logon information: default username: *admin*, default password: *pfSense*.

Take the installation CD out of CD-drive, and after the reboot, the initial configuration starts:

- choose not to set up VLANs;
- pfSense requires at least 2 assigned interfaces, and it lists the valid interfaces, in our case “bge0” and “bge1”;
- enter the LAN interface: bge1;
- enter the WAN interface: bge0;
- optional interface: just hit enter finishing the configuration, since there are no optional interfaces;
- confirm the assignment of WAN and LAN interfaces;
- system configuration starts, and can take up to several minutes;
- pfSense console setup menu appears.

Reaching to this point of installation, we are ready to continue with configuring using the web interface. From a machine on the LAN interface, open <http://192.168.1.1> in the browser, and log on with the default credentials mentioned above.

Run through the initial setup wizard, configuring as follows:

- hostname: UECcloudFW
- domain: privatecloud.itu.dk
- primary DNS server: 130.225.127.212
- secondary DNS server: 130.225.127.213
- time server: dk.pool.ntp.org
- time zone: Europe/Copenhagen

Next step is configuring WAN and LAN interfaces.

WAN configuration:

- choose static IP configuration instead of DHCP
- IP address: 130.225.115.199/24
- Gateway: 130.225.115.1
- remove: “Block RFC1918 Private Networks” and “Block bogon networks”

LAN configuration:

- IP address pool/network: 192.168.1.1/24

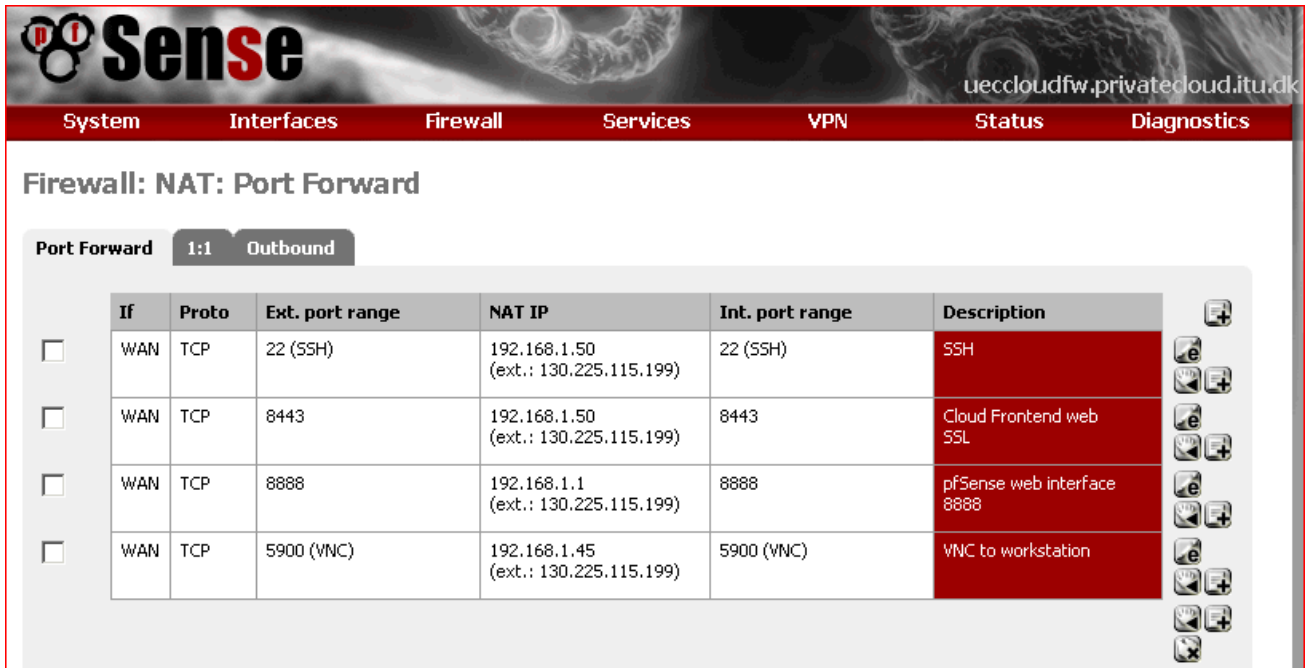
Set Admin WebGUI Password. Next, the system information appears.

There are a few more things that need to be configured:

- Set up DHCP server: make sure that the address would not collide with UEC server IP addresses, or pools for instance’s IPs (in our case we could put it at 192.168.1.20-192.168.1.45)
- Set up NAT (network address translation):

- SSH access to the CLC/WS3 server on port 22;
- SSL web access to the cloud web interface on port 8443;
- Enable configuring of firewall from WAN network: <https://130.225.115.199:8888>
- VNC access to the administrative workstation
(Remember to apply the changes!)

Configuring those forwarding rules would look like this in the pfSenses web interface:



Finally, test the SSH, VNC and pfSense WebGUI access from WAN network, and verify that they are working.

Installation of the administrative workstation

For the purpose of controlling the cloud, an administrative workstation would be needed, primarily for the purpose of having a GUI interface to administer the cloud and whole infrastructure. We could also use it for bundling and registering new Eucalyptus Machine Images (EMIs).

The installation is performed on a normal workstation, in our case using the latest Ubuntu 11.04. Make sure that you update your system, and that a browser is installed, in our case Mozilla Firefox browser. Also, make sure that you install the HybridFox6 for controlling the cloud via GUI.

Install KVM for the purpose of installing of images on KVM platform, and bundling them:

```
apt-get install qemu-kvm
```

Install euca2ools to be able to manage the cloud:

```
sudo apt-get install euca2ools
```

Now the workstation is ready for all the administrative tasks. The last step would be downloading the credentials from the cloud web interface, but that will be done after first installing the cloud servers.

⁶ <http://code.google.com/p/hybridfox>

Installation of the first server (CLC, WS3, CC, SC)

Mount the CD installation media, and boot from it. Choose the language – English.

Choose “Install Ubuntu Enterprise Cloud”.

Choose English as installation language, and the default language for the installed system.

Select your location – other. Select your location – Europe. Select your location – Denmark.

Configure locales – select United States. This comes only if no locale is defined for the combination of language and country of your selection. If we were choosing i.e. Portuguese and Brazil, everything would be fine, since Portuguese is defined for Brazil.

Configure your keyboard - choose no. If you want your keyboard to be detected, you could choose yes and go through a wizard of options, ending with your keyboard identified. Configure the keyboard for Denmark, and on the next screen choose Denmark.

If you would get troubles with getting some special characters, like “~” and “|”, as a quick workaround you could make a SSH connection to the machine from a machine where you positively know that you can type those characters, and run the commands from there, instead from the local host.

The installation starts loading additional components – this part should not take more than few minutes.

Next step is configuring network – choose the primary network interface for the installation: eth0.

The installation automatically tries to configure this network interface for DHCP – cancel this pressing ESC button, since we want to assign static IP address for the network interface. If you are not quick enough, pressing Esc later would just bring you to the menu of choosing the step for installation – here you should choose “Configure the network”.

Multiple choices for configuring the network are presented, where one of the options would be configuring the network manually – choose that option.

Machines IP address: 192.168.1.50

Netmask for the class C networks: 255.255.255.0

Gateway IP address: 192.168.1.1

Put all the DNS servers that are available, separated by blank space:

192.168.1.50 130.225.127.212 130.225.127.213

The first one to be contacted would be the first one from the list. In our case it is the local server’s address (192.168.1.50), since we will install DNS server on this server later on. The other servers are the network specific DNS servers, and their IP addresses should be obtained from the IT department.

Choose the name for the host (remember that hostnames are case sensitive in UNIX / Linux). In our case, we choose “UECserver01” to identify that the server is hosting UEC on it, but we don’t designate the server role in the name, since server roles could be changed. Choosing “UECserverNN” would allow us to scale out to up to 100 hosts without getting into naming problems.

Domain name – preferably your network domain obtained from your IT department, in our case it is organization’s domain name “privatecloud.itu.dk”.

Select cloud installation mode - for the initial install of the cloud, we just continue here, since there is no present Cloud Controller, and we are just installing one. If we were adding the server to an existing cloud, we would put the hostname or IP address of the Cloud Controller here (in case if it wasn't detected automatically, which it normally is).

Choose the components to be installed on the first server – we install Cloud controller, Walrus, Cluster, Storage and Node controller. This is a design choice to put all of the cloud and cluster related components on one server; for the purpose of initial private cloud implementation, we would group cloud and cluster roles on one server, and NC role on the other server, keeping the minimum installation on two servers.

Since we have two network interfaces in the system, we should designate one that is to be used for communication with computing nodes (NCs). Since our computing nodes are going to be placed on the private network, we should choose the secondary network interface (eth1) as the one used for communicating with nodes.

Confirm the time zone, in our case Europe/Copenhagen. If not correctly detected, choose No, and you will get the choice of setting up the time zone.

Now we have a choice to partition our disk(s). For the sake of simplicity, choose a guided method, using the entire disk, and setting up LVM. In our case we have “SCSI.CCIS (-,0,0) (cciss/c0d0) – 293.6 GB Compaq Smart Array” disk, and we would just partition it using default options. If having several disks, we could choose another way of partitioning, for the sake of performance or redundancy.

Logical Volume Manger (LVM), allows us to create logical volumes on physical hard disk(s). LVM volumes can be created on both software RAID partitions (that is, not having controller that supports RAID configuration, and doing it on the operating system level instead) and standard partitions residing on a single disk. Volumes can also be extended, giving greater flexibility to systems as requirements change. Choosing configuration using LVM has a greater degree of complication as side effect. Some of the terms used in the process are:

- Volume Group (VG): contains one or several Logical Volumes (LV);
- Logical Volume (LV): is similar to a partition in a non-LVM system; Multiple Physical Volumes (PV) can make up one LV, on top of which resides the actual EXT3, XFS, JFS, and other file system;
- Physical Volume (PV): physical hard disk or software RAID partition; the VG can be extended by adding more PVs.

In our simple case, choosing LVM setup doesn't give any extra work. If having another disk configuration, choosing setup using LVM is ensuring that the configuration is done the right way in the terms of performance and redundancy, and is recommended.

If the disk was used before, having some partitions on it, we would be prompted to overwrite the old configuration, warned that we would lose all the data saved on the disk.

Now, write the partition scheme to the disk choosing “Yes”. Choose the amount of VG to use: 293.4 GB. In our case, we use the whole disks capacity. Write the changes on the disk.

After the disk configuration, the base system installation starts. This could take from just a few minutes to tens of minutes, since the files are being copied from the installations media to the hard drive, and how quick that would be done is depending on the speed of reading the installation media, and writing it on the disk.

After installing the base system is finished, we should now set up the user. This user would be used to log on the server. Applying the principle of least privilege, we should not use administrative accounts (as root is) for non-administrative purposes. We would log on as this user, and when administrative work is needed to be done, we would elevate the privileges of this account using prefix “sudo”, thus running commands as root.

First, choose the full name for the user – that is display name: UEC Cloud. Second, choose the username for the account (bear in mind that usernames are case sensitive): ueccloud. At last, choose the password for the user, and repeat the password to make sure that it was typed correctly. In production environment it is recommended that the password is at least 16 characters long, and it should include at least one character from every of the four character groups (upper case letters, lower case letters, numbers, and special characters).

We chose not to encrypt our home directory, since it is putting some overhead on processing power and storage. In production, if you expect that your home directory would host sensitive data, you would consider encrypting it.

If your internal network configuration includes proxy server for accessing the internet, add the proxy's address at the next screen. If no proxies used, leave the field blank, and choose to continue.

Configuring of Advanced Packaging Tool (APT)⁷ starts. If the machine has access to the internet, packages are retrieved from there. After configuring APT, preparation for the software installation starts.

Choose the setting for updating of your system. We would choose not to update the system automatically, since we would like to have a full control over the packages on the system. Otherwise the automatic update is an available option. For production systems, choosing either no automatic updates, or managing the system with Landscape⁸ are valid options. Landscape is systems management and monitoring service provided by Canonical, and is a paid service for monitoring, managing and updating Ubuntu infrastructure using single interface, thus lowering management and administration costs.

Software installation starts. Again, this could take from just a few minutes to tens of minutes, since the files are being copied from the installations media to the hard drive, and how quick that would be done is depending on the speed of reading the installation media, and writing it on the disk.

Mail (Postfix) configuration step is next, choosing the system mail name. We will use registered domain name ("itu.dk") to ensure the mail flow, so the system mail name would be "privatecloud.itu.dk". Internal IT department would help you setting this up, and supply you with the required information.

The installation reaches the point of configuring the cluster. Name the cluster as "UECcluster01" to leave the naming place for scaling out the cloud later, adding another cluster.

Add the following pool of IP addresses: 192.168.1.100-192.168.1.199. We are not taking the whole class C subnet, since we will leave the rest of the addresses for another cluster that might come in the future.

Main installation starts, and again, this could take from just a few minutes to tens of minutes, since the files are being copied from the installations media to the hard drive, and how quick that would be done is depending on the speed of reading the installation media, and writing it on the disk.

Since the only operating system on our server would be Ubuntu server including UEC, we should choose to install boot loader, to make machine be able to boot from the disk: choose "Yes" to install GRUB boot loader to the master boot record.

After that, the system is finishing the installation. This includes several short steps, and should not take more than a few minutes. After that, the installation is completed – remove the installation CD, and choose "Continue" to reboot the server.

⁷ http://en.wikipedia.org/wiki/Advanced_Packaging_Tool

⁸ <https://landscape.canonical.com>

Make sure that you take out the installations media (CD), to prevent the system booting again into installation, if your system is set to boot from the CD first. If that happens, you just cancel the installation, take the media out, and recycle the machine – no damage to the installation would be done.

After restart, the machine boots for the first time, and after the first login, you will meet the welcome screen, listing some information about the system running on your machine. If the system indicates that *“System information disabled due to load higher than 1.0”*, don’t worry. This is normal behavior after the first logon on the newly installed system. Booting up the first time, a considerable load is put on the system, and thus no processing power is left for composing and showing detailed system configuration.

From now on, a SSH session to the machine(s) is suggested (from Windows, free SSH client “PuTTY” can be used), for several reasons. For the first, it is more convenient to work in just one interface (SSH terminal) than switching between machines. Next, you could use copy & paste option for the commands following the initial setup, during the post-installation steps. This way eventual misspelling would be avoided and configurations are performed faster. Be aware, though, that copying and pasting text with different coding (i.e. from Windows format to Linux) would eventually cause that some characters are not pasted, or pasted wrong. Please refer to the Appendix 4 about troubleshooting for more information.

Post installation steps

Set up static IP for both network interfaces. If the primary network interface was configured with DHCP during the installation, then set that interface to use static IP first. If you managed to set static IP during the installation, the configuration of only the secondary network interface is needed.

Public IP should be in 192.168.1.0 network, and private in 10.0.0.0 network.

Edit the network configuration of the server:

```
sudo vi /etc/network/interfaces
```

Under the primary network interface, remove the “dhcp” and configure it with static IP address:

```
auto eth0  
iface eth0 inet static  
    address 192.168.1.50  
    netmask 255.255.255.0  
    network 192.168.1.0  
    broadcast 192.168.1.255  
    gateway 192.168.1.1
```

Add the secondary network interface that is to be configured with static IP address too:

```
auto eth1  
iface eth1 inet static  
    address 10.0.0.1  
    netmask 255.255.255.0  
    network 10.0.0.0  
    broadcast 10.0.0.255
```

Restart the networking to apply the changes:

```
sudo /etc/init.d/networking restart
```

To be able to synchronize time between servers in the cloud, we will install NTP server on the first server, which will synchronize time with the external time source. All other servers in the cloud would synchronize the time with first server.

To set up NTP server, we have to install the NTP server first:

```
sudo apt-get install ntp  
(if error "unable to locate package" occurs, run sudo apt-get update)
```

To make sure that the server will continue to serve the time even when off-line, a setting that makes it possible for the server to use its own clock as source should be made. This is done by opening the file `/etc/ntp.conf` and adding the following after the "# Specify one or more NTP servers" line:

```
server 127.127.1.0  
fudge 127.127.1.0 stratum 10
```

Restart the NTP service afterwards to make the changes effective:

```
sudo /etc/init.d/ntp restart
```

Check server time:

```
date
```

Since only the first server has access to the internet through its public network, a router (Network Address Translation, NAT) will be installed⁹, to make it possible for the rest of the cloud servers being only on the private network to connect to the outside network, for the purpose of i.e. fetching updates.

To set-up router, run the following commands:

```
sudo iptables -A FORWARD -o eth0 -i eth1 -s 10.0.0.0/24 -m conntrack --ctstate NEW -j ACCEPT  
sudo iptables -A FORWARD -m conntrack --ctstate ESTABLISHED,RELATED -j ACCEPT  
sudo iptables -A POSTROUTING -t nat -j MASQUERADE  
sudo iptables-save | sudo tee /etc/iptables.sav  
sudo iptables-restore < /etc/iptables.sav  
sudo sh -c "echo 1 > /proc/sys/net/ipv4/ip_forward"
```

Edit the `/etc/sysctl.conf`:

```
uncomment net.ipv4.ip_forward=1
```

and add the following after the uncommented line:

```
net.ipv4.conf.default.forwarding=1  
net.ipv4.conf.all.forwarding=1
```

Now run:

```
sudo iptables-save | sudo tee /etc/iptables.sav
```

Edit `/etc/rc.local` and add the following lines before the "exit 0" line:

```
iptables-restore < /etc/iptables.sav
```

To configure the gateway for routing between two interfaces (by enabling IP forwarding):

```
sudo sh -c "echo 1 > /proc/sys/net/ipv4/ip_forward"
```

Restart the server to apply and test the changes:

```
sudo reboot
```

⁹ <http://ubuntuforums.org/showthread.php?t=713874>

To make it easier to resolve the DNS names for the servers on the private network, and also for the first server itself, DNS server (BIND) should be installed. To do that, with default options, run the following command:

```
sudo apt-get install bind9
```

Check the DNS configuration `/etc/resolv.conf`:

The file should look like:

```
domain UECcloud.private  
search UECcloud.private  
# first contact the local DNS server installed on the same server  
nameserver 10.0.0.1  
nameserver 192.168.1.99  
# if that fails, contact the secondary servers (ISP's DNS)  
nameserver 193.162.153.164  
nameserver 194.239.134.83
```

To test that the local DNS is serving the DNS queries.

```
nslookup dr.dk  
Server: 10.0.0.1  
Address: 10.0.0.1#53  
Non-authoritative answer:  
Name: dr.dk  
Address: 195.137.194.128
```

If having troubles with DNS, please refer to Appendix 4 – troubleshooting.

To enable the servers to send mails, such as user verification mails, SMTP server should be installed. The default SMTP server for Ubuntu is Postfix¹⁰, and it is installed during the initial installation. To check or adapt the configuration, run the following command:

```
sudo dpkg-reconfigure postfix
```

Our setup is:

- Internet Site (mail sent directly using SMTP)
- Privatecloud.itu.dk
- Enter admin mail address (zopa@itu.dk)
- Add eventual destinations, otherwise leave it like it is
- No
- Add local network 10.0.0.0/24 and 192.168.1.0/24
- Mailbox size limit – 0
- Leave “+”
- ipv4

The same could be done by editing `/etc/postfix/main.cf` (and others) as needed. After modifying `main.cf`, be sure to run `'sudo /etc/init.d/postfix reload'`.

Make sure that you have the latest version of Eucalyptus:

```
sudo apt-get upgrade eucalyptus
```

Depending on the version of your Eucalyptus and how many updates has been made since the release of your version, this process could take from several minutes to several tens of minutes, depending on the internet connection speed, and the speed of your hardware.

¹⁰ <https://help.ubuntu.com/community/Postfix>

To check the version of Eucalyptus, view the following file:
/etc/eucalyptus/eucalyptus-version

This is the last step of initial configuration for the first server in the cloud, and to make sure that the configuration is implemented, restart the machine (**sudo reboot**) and eventually check the installed features.

Installation of second server (NC)

Mount the installation media, and boot from it. Choose the language – English. Follow the same steps as installing the first server, until you reach to the point of configuring the network.

Choose to configure the network manually.

Set up static IP for the (private) network interface.

```
IP address 10.0.0.2  
netmask 255.255.255.0  
gateway 10.0.0.1  
name server addresses: 10.0.0.1
```

Write “UECserver02” for the hostname.

For the domain name, write “privatecloud.itu.dk”.

The installation can contact the CLC and get information about CLC, WS3, CC and SC components are installed in the cloud, and offers to install the NC component on the server – install the NC.

Configuring disks is done by following the same procedure as for the first server. Installing of the base system starts, after that the configuring of the APT, then installing the software.

If you get the error about hardware not supporting virtualization acceleration, please check if your processor is VT-enabled (having VT extensions), and remember to turn that option on in BIOS.

After the successful installation, the server reboots.

Post installation steps

As a first step, check the configuration of the network interface (eth0):

```
sudo vi /etc/network/interfaces
```

Networking is implemented using bridging.

To check network connectivity, try pinging an address that you positively know is giving response to ping requests.

To be able to synchronize the time on the server, we will install NTP:

```
sudo apt-get install ntp
```

To make the server synchronize the time with the first server, open the file **/etc/ntp.conf** and add:
server 10.0.0.1

Restart the NTP service afterwards to make the changes effective:
sudo /etc/init.d/ntp restart

Check server time:
date

Check the DNS configuration **/etc/resolv.conf** - it should look like:
search UECcloud.private
nameserver 10.0.0.1

Make sure that you have the latest version of Eucalyptus:
sudo apt-get upgrade eucalyptus

To check the version of Eucalyptus, view the following file:
/etc/eucalyptus/eucalyptus-version

Restart the NC:
sudo restart eucalyptus-nc

This is the last step of initial configuration for the third server in the cloud, and to make sure that the configuration is implemented, restart the machine (**sudo reboot**) and eventually check the installed features.

Verifying the cloud installation

There is a row of verifying steps that could be performed to make sure that the cloud components were registered correctly. As of Ubuntu 10.04 LTS, all component registration should be automatic:

- Public SSH keys have been exchanged properly;
- The services are configured properly;
- The services are publishing their existence;
- The appropriate uec-component-listener is running;
- Verify registration.

The steps mentioned above are needed if “Package Install” method is used. Since we are using “CD Install” method, we can skip those steps (or verify and perform a check if things are as they should be - please refer to the Appendix 4 about troubleshooting for more information on this).

Running the cloud

Now we have a cloud infrastructure installed on the server machines. Some initial administrative steps need to be performed to get the cloud up and running. For more detailed description of the steps accompanied with screenshots, please refer to the Appendix 2 – installation in virtual environment.

Initial steps

As a part of the initial setup via web interface, cloud admin should choose to perform the following steps:

- Change the password for the user “admin” (use strong password, and back it up at safe place)
- Check the cloud configuration found under “Configuration”, making sure that all the components are registered right
- Add users, both administrators and/or users
- Download and backup of credentials for user “admin” and other users
- Install image(s) available under “Store” – those images would appear under “Images” after downloading

Cloud web interface

Finishing installing the three servers, our cloud is set up. Next step is connecting to the web interface of the UEC, and that is done by opening the public address of our CLC in a browser: <https://192.168.1.50:8443> (ignore the certificate error after checking it out).

The default logon credentials are username “admin” and password “admin”. After logging on with those credentials, the next screen appears, where the password should be changed, administrator’s e-mail address should be given, and the CLC’s public IP address is guessed – make sure that the right public IP of the CLC is entered there. After submitting, an admin web interface opens, with the possibilities for managing credentials, images, users, configuration, extras and store.

Go to the “Credentials”, and download those credentials for the admin user, and save it on the client machine, and also make a backup of them. Click on “Download Credentials” button from the Credentials menu, and save the “euca2-admin-x509.zip” archive on your local machine.

Go to the “Store”, and download the standard images. (This could take some time, depending on your internet connection)

Downloading user credentials

We will obtain user credentials from command line. On the CLC, do the following:

```
mkdir -p ~/.euca  
chmod 700 ~/.euca  
cd ~/.euca  
sudo euca_conf --get-credentials mycreds.zip  
unzip mycreds.zip  
ln -s ~/.euca/eucarc ~/.eucarc
```

Running an image using command line

Create your key pair:

```
touch ~/.euca/mykey.priv  
chmod 0600 ~/.euca/mykey.priv  
euca-add-keypair mykey > ~/.euca/mykey.priv
```

Verify you keypair:

```
euca-describe-keypairs
```

That should list your “mykey” keypair that you just created.

Allow access via SSH (port 22) to the instances:

```
euca-authorize default -P tcp -p 22 -s 0.0.0.0/0
```

Instantiate an instance of your registered image using “euca-run-instances \$EMI -k your_keypair -t instance_size”, in our case it would be something like:

```
euca-run-instances emi-DF361069 -k mykey -t m1.xLarge
```

Generally, we could also run our other standard instances:

- Ubuntu 9.10 - Karmic Koala (i386)
euca-run-instances -k <your key pair> emi-DF361069
- M/DB Appliance
euca-run-instances -k <your key pair> emi-DEC91072
- Ubuntu 10.04 LTS - Lucid Lynx (amd64)
euca-run-instances -k <your key pair> emi-DEBC106F

To monitor the state of the instance, run:

watch -n5 euca-describe-instances

When the status is listed as “running”, and not pending, you can connect to your newly instantiated instance.

To determine which IP address the instance got:

IPADDR=\$(euca-describe-instances | grep \$EMI | grep running | tail -n1 | awk '{print \$4}')

Connect to your instance:

ssh -i ~/.euca/mykey.priv admin@\$IPADDR

That would be the installation and initial steps implementing private cloud on physical hardware.

After testing and initial implementation, running several machines a NC’s limit would be reached, and then we would need to scale out adding some more NC’s. When storage and cluster performance would be reached, we would need to scale out adding a whole cluster. Both of these scale-out possibilities are described in the Appendix 2 – installation in virtual environment.

Besides processing power, storage limits would soon be reached, and SAN (iSCSI) based storage should be added to all servers.

Photos of the hardware infrastructure

Here are some photos of the hardware units used to implement the solution.

This is the whole environment, with the reserve servers for scaling out the cloud, public network switch, and laptop computer serving as administrative workstation – all on the left side of the picture:



On the right side of the picture there are the 2 cloud servers and the firewall server. In the middle, behind the screen and the keyboard, there is a private network switch:



Acknowledgement

The work reported in this report has been partially funded by the Danish Agency for Science, Technology and Innovation under the project "Next Generation Technology for Global Software Development", #10-092313.