



**Haduzilla - Building hadoop cluster
with Debian preseed**
黑肚龍：無人值守自動安裝 Hadoop 叢集

Jazz Wang
Yao-Tsung Wang
jazz@nchc.org.tw



Powered by DRBL

WHO AM I? 這傢伙是誰啊? JAZZ?

- 講者介紹：
 - Jazz Yao-Tsung Wang @ NCHC / NCTU ECE Master
 - 國網中心 王耀聰 副研究員 / 交大電控八九級碩士
 - jazz@nchc.org.tw
- 所有投影片、參考資料與操作步驟均在網路上
- All the slides could be found at
 - <http://trac.nchc.org.tw/cloud>



FOSS Developer
行動力薄弱的開發者

TRTC WSU/
Hadop4Win /
Haduzilla / Ezilla


FOSS Promoter
自由軟體推廣者

DRBL/Clonezilla
Partclone/Tuxboot
Hadoop Ecosystem 2

FOSS End User

FOSS 使用者

Debian/Ubuntu
Access Grid
Motion/VLC
Red5
Debian Router
DRBL/Clonezilla
Hadoop



Let's start with
a story of
DATA!!

Data Explosion!! 始於 2007 的「資料大爆炸」時代

Information Versus Available Storage





Figure 1

Information Created, Captured and Replicated

6-Fold Growth
in Four Years



Source: IDC, 2007

2007 年，IDC 預估
2010 年會成長**六倍**！
(相較 2006 年)

Source: IDC, 2007

出處：The Expanding Digital Universe,
A Forecast of Worldwide Information Growth Through 2010,
March 2007, An IDC White Paper - sponsored by EMC

•<http://www.emc.com/collateral/analyst-reports/expanding-digital-idc-white-paper.pdf>

2006 161 EB
2010 988 EB (預測)

Digital Universe expanded 1.6x each year!! 每年約 1.6 倍



追蹤歷年的 IDC 數據：

2006 161 EB

2007 281 EB

2008 487 EB

2009 800 EB (0.8 ZB)

2010 988 EB (預測)

2010 1200 EB (1.2 ZB)

2011 1773 EB (預測)

2011 1800 EB (1.8 ZB)

景氣差而成長趨緩？
或受新技術抑制？

出處：[Extracting Value from Chaos](#),
June 2011, An IDC White Paper - sponsored by EMC

<http://www.emc.com/collateral/about/news/idc-emc-digital-universe-2011-infographic.pdf>

Now we all
need to store
and process

BIG DATA!!

What is Hadoop ?

用一句話解釋 Hadoop 是什麼 ??

Hadoop is a software platform that lets one easily write and run applications that process vast amounts of data.

Hadoop 是一個讓使用者簡易撰寫並執行處理海量資料應用程式的軟體平台。

亦可以想像成一個處理海量資料的生產線，只須學會定義 *map* 跟 *reduce* 工作站該做哪些事情。

Features of Hadoop ...

Hadoop 這套軟體的特色是 ...

- **海量 Vast Amounts of Data**
 - 擁有儲存與處理大量資料的能力
 - Capability to **STORE** and **PROCESS** vast amounts of data.
- **經濟 Cost Efficiency**
 - 可以用在由一般 PC 所架設的叢集環境內
 - Based on large clusters built of **commodity hardware**.
- **效率 Parallel Performance**
 - 透過分散式檔案系統的幫助，以致得到快速的回應
 - With the help of HDFS, Hadoop **have better performance**.
- **可靠 Robustness**
 - 當某節點發生錯誤，能即時自動取得備份資料及佈署運算資源
 - Robustness to add and remove computing and storage resource without shutdown entire system.

Which companies are powered by Hadoop ??

有哪些公司在用 Hadoop 這套軟體 ??

- **Yahoo** is the key contributor currently.
- **IBM** and **Google** teach Hadoop in universities ...
- http://www.google.com/intl/en/press/pressrel/20071008_ibm_univ.html
- **The New York Times** used **100 Amazon EC2 instances** and a Hadoop application to process **4TB of raw image TIFF data** (stored in S3) into **11 million finished PDFs** in the space of **24 hours** at a computation cost of about **\$240** (not including bandwidth)
 - from <http://en.wikipedia.org/wiki/Hadoop>
- <http://wiki.apache.org/hadoop/AmazonEC2>
- <http://wiki.apache.org/hadoop/PoweredBy>
 - **A9.com**
 - **ADSDAQ by Contextweb**
 - **EHarmony**
 - **Facebook** **Facebook**
 - **Fox Interactive Media**
 - **IBM**
 - **ImageShack**
 - **ISI**
 - **Joost**
 - **Last.fm**
 - **Powerset**
 - **The New York Times**
 - **Rackspace**
 - **Veoh**
 - **Metaweb**
- **Tweeter**

Hadoop in production run

商業運轉中的 Hadoop 應用

- February 19, 2008
- Yahoo! Launches World's Largest Hadoop Production Application
- <http://developer.yahoo.net/blogs/hadoop/2008/02/yahoo-worlds-largest-production-hadoop.html>

| | |
|---|--------------------------|
| Number of links between pages in the index | roughly 1 trillion links |
| Size of output | over 300 TB, compressed! |
| Number of cores used to run single Map-Reduce job | over 10,000 |
| Raw disk used in the production cluster | over 5 Petabytes |

You can store
and process
BIG DATA via
Large Cluster!!

Common method to deploy Cluster in Labs



1. Setup one
Template
machine

2. **Cloning**
to
multiple
machine



3. **Configure**
Settings



4. **Install**
Job
Scheduler



5. **Running**
Benchmark

Challenges of common method in Labs

Add New User Account ?

Upgrade Software ?

How to share user data ?

Configuration Synchronization

How to deploy 4000+ Nodes ?!

資料標題：Scaling Hadoop to 4000 nodes at Yahoo!

資料日期：September 30, 2008

| | |
|--------------------|--------------|
| Total Nodes | 4000 |
| Total cores | 30000 |
| Data | 16PB |

| | 500-node cluster | | 4000-node cluster | |
|-------------------------------|-------------------------|-------------|--------------------------|-------------|
| | write | read | write | read |
| number of files | 990 | 990 | 14,000 | 14,000 |
| file size (MB) | 320 | 320 | 360 | 360 |
| total MB processes | 316,800 | 316,800 | 5,040,000 | 5,040,000 |
| tasks per node | 2 | 2 | 4 | 4 |
| avg. throughput (MB/s) | 5.8 | 18 | 40 | 66 |



- Make Hadoop deployment *agile*
- Integrate with dynamic cluster deployments

Source: Deploying hadoop with smartfrog

http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf

SmartFrog - HPLabs' CM tool

- Language for describing systems to deploy
—everything from datacentres to test cases
 - Runtime to create *components* from the model
 - Components have a lifecycle
 - LGPL Licensed, Java 5+
- <http://smartfrog.org/>

Source: Deploying hadoop with smartfrog

12 | http://people.apache.org/~stevel/slides/deploying_hadoop_with_smartfrog.pdf



If you need to deploy in Cloud - try Puppet

如果要在 Amazon EC2 上佈署 Hadoop 等軟體，可以考慮 Puppet 因為作業系統已由虛擬機器的範本裝好了，只能用「有碟」的作法！

hstack

Blog

Hadoop/HBase automated deployment using Puppet

with 23 comments

Introduction

Deploying and configuring Hadoop and HBase across clusters is a complex task. In this article I will show what we do to make it easier, and share the deployment recipes that we use.

For the [tl;dr](#) crowd: go get the code [here](#).

Cool tools

Before going into how we do things, here is the list of tools that we are using, and which I will mention in this article. I will try to put a link next to any tool-specific term, but you can always refer to its specific home-page for further reference.

- [Hudson](#) – this is a great CI server, and we are using it to build Hadoop, HBase, Zookeeper and more
- The [Hudson Promoted Builds Plug-in](#) – allows defining operations that run after the build has finished, manually or automatically
- [Puppet](#) – configuration management tool We don't have a dedicated operations team to hand off a list of instructions on how we want our machines to look like. The operations team helping us just makes sure the servers are in the rack, networked and powered up, but once we have a set of IPs (usually from [IPMI](#) cards) we're good to go ourselves. We are our own [devops](#) team, and as such we try to automate as much as possible, where possible, and using the tools above helps a lot.

<https://github.com/hstack/puppet>

<http://hstack.org/hstack-automated-deployment-using-puppet/>

http://www.cioinsight.com/images/stories/slideshows/SS_142511_CIO_TechSkills/

Top 10 Tech Skills for 2012

Puppet

If you haven't ever heard of Puppet, it's time to check it out. The service is an IT infrastructure management solution that helps cut down on the amount of time it takes to handle simple tasks. Puppet is robust and useful, but companies need people who are skilled enough to harness its power.

CIO INSIGHT.

Can I install
ONE server to
deploy hadoop
cluster ?

Yes, use DRBL to deploy Hadoop

- Need to build new debian packages
- drbl-hadoop – Mounting local disk for HDFS and MapReduce
svn co <http://trac.nchc.org.tw/pub/grid/drbl-hadoop-0.1/>
- hadoop-register – for multiuser registration and ssh client
svn co <http://trac.nchc.org.tw/pub/cloud/hadoop-register>



root / **drbl-hadoop-0.1**

| Name ▲ |
|--------------------------|
| ↑ ../ |
| 📄 drbl-hadoop |
| 📄 drbl-hadoop-mount-disk |



root / **hadoop-register**

| Name ▲ | Size | Rev | Age | Last |
|---------------------------------|--------|------------|---------|------|
| ↑ ../ | | | | |
| ▶ 📁 etc | | 103 | 4 weeks | wa |
| 📄 adduser.php | 1.3 kB | 85 | 6 weeks | wa |
| 📄 check_activate_code.php | 2.2 kB | 85 | 6 weeks | wa |
| 📄 check_user_identification.php | 2.9 kB | 85 | 6 weeks | wa |

About hadoop.nchc.org.tw

- DRBL Server x 1 Node (hadoop)
- DRBL Client x 20 Nodes
(hadoop101~hadoop120)
- **Powered by Debian Squeeze 6.0.4**



使用者註冊頁面 Hadoop-Register

Hadoop 帳號申請

帳號:

密碼:

登入

重填

[新增帳號](#) [忘記密碼](#) [操作問題回報](#)


歡迎加入討論群組, 以利接收即時公告事宜

家目錄空間吃緊中, 請盡量上傳至HDFS後,
清除家目錄檔案, 謝謝!

註冊人數: 1460 / 1999 人

[MapReduce 狀態](#) | [HDFS 狀態](#)

過去 24 小時 CPU 負載 - [查詢完整系統負載](#):



Running Jobs

[Quick Links](#)

| Jobid | Priority | User | Name | Map % Complete | Map Total | Maps Completed | Reduce % Complete |
|-----------------------|----------|-------|--|----------------|-----------|----------------|-------------------|
| job_201104290234_0905 | NORMAL | h1196 | PA: Local Apriori over input: n/1mpy54 /input, with minSup: 15000, ep: 0.5 | 100.00% | 10 | 10 | 100.00% |
| | | | PA: Local Apriori over | | | | |

網站帳號 jazzwang E-mail [redacted] 姓名 王耀聰 電話 0 單位 0 用途 0 主機帳號 h998 主機密碼 [redacted] 登出

NameNode

[檔案\(F\)](#) [編輯\(E\)](#) [檢視\(V\)](#) [歷史\(Y\)](#) [工具\(T\)](#) [說明\(H\)](#)

1. hadoop.nchc.org.tw

| | |
|-----------|---|
| Started: | F |
| Version: | 0 |
| Compiled: | S |
| Upgrades: | T |

[Browse the filesystem](#)
[Namenode Logs](#)

Cluster Summary

2079646 files and

WARNING: There are

| | |
|----------------|--|
| Configured Cap | |
| DFS Used | |
| Non DFS Used | |

```
Linux hadoop 2.6.32-5-amd64 #1 SMP Wed Jan 12 03:40:32 UTC 2011 x86_64
```

```
The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.
```

```
Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
```


```
Last login: Tue Apr 26 15:45:44 2011 from nat235.dynamic.cs.nctu.edu.tw
h998@hadoop:~$
```

Powered by Zterm

<http://zhouer.org/ZTerm/>

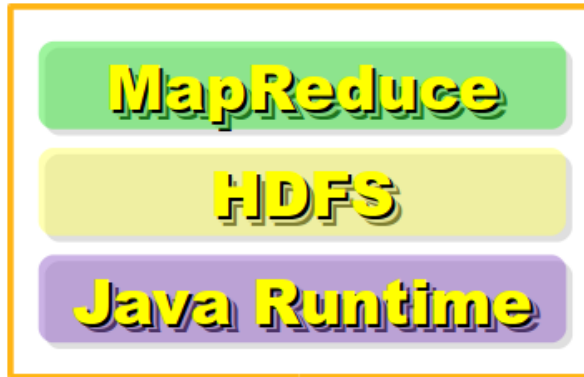
系統狀態監控 Ganglia

- 採用自由軟體 Ganglia 來蒐集電腦叢集的負載狀態
- <http://ganglia.sourceforge.net/>



DRBL+Hadoop=Haduzilla 黑肚龍系統架構

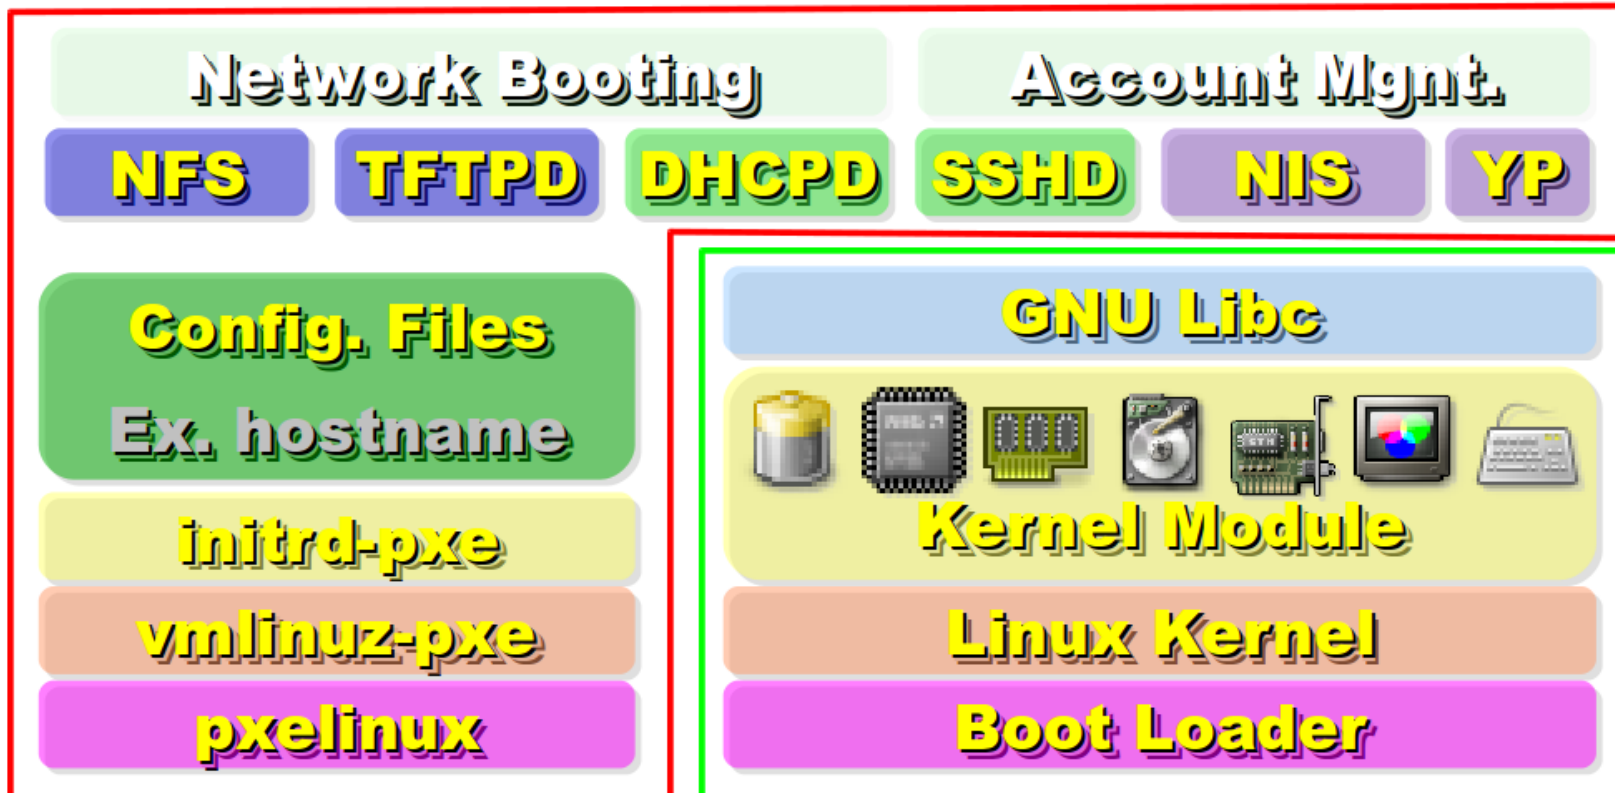
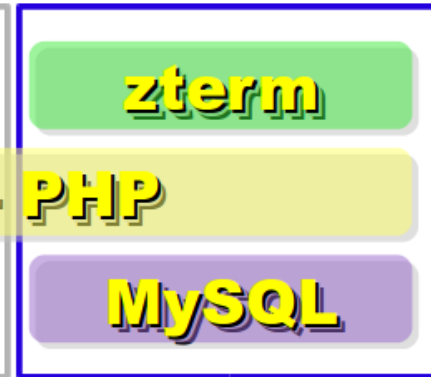
Hadoop



Ganglia



Register



DRBL


Linux

Can you help me to
deploy my own
multiuser hadoop
cluster like
hadoop.nchc.org.tw?

In 2009, DRBL-Hadoop Live CD released ...

舊影片：http://www.youtube.com/watch?hl=en&v=Ix4WigGvE_A

下載點：<http://drbl-hadoop.sf.net>



Running Hadoop on DRBL Live CD

jazzwbear 8 videos

The screenshot shows a VM console window titled 'Local host - VMware Server Console'. The main window displays the Hadoop NameNode status page for 'Hadoop NameNode (debian-eth1-9000 - (crownase))'. The page shows the following statistics:

- DFS Remaining : 244.61 MB
- DFS Used : 16 KB
- DFS Used% : 0.01 %
- Live Nodes : 2
- Dead Nodes : 0

Below the statistics, it indicates 'Live Datanodes : 2' and provides a table of details for each datanode.

| Node | Last Contact | Admin State | Size (GB) | Used (%) | Used (%) | Remaining (GB) | Blocks |
|-----------|--------------|-------------|-----------|----------|----------|----------------|--------|
| debian101 | 0 | In Service | 0.12 | 0.01 | | 0.12 | 0 |
| debian102 | 2 | In Service | 0.12 | 0.01 | | 0.12 | 0 |

The video player interface at the bottom shows a progress bar at 5:04 / 5:33 and a resolution of 360p.

But I want it
installed to disks
for production

...

What should I do ?

On 11 Feb 2011, 4\$ shared about preseed!

感謝 4\$ 大大分享 Debian 6.0 自動化安裝

The image shows a screenshot of a blog post on the 'FourDollars Blog'. The blog's header includes the title 'FourDollars Blog' and a subtitle 'GNU/Linux, Debian/Ubuntu, Mac OS X, Free Software/Open Source Software, Freeware, and Programming.' Below the header are navigation tabs for '首頁', 'Software Development', and 'Debian Package'. The main content area features a date '2011年2月11日星期五' and a title '\$4 的 Debian 6.0 自動化安裝'. The article text describes how to create a bootable USB stick for Debian 6.0, mentioning the download of 'debian-6.0.2.1-i386-businesscard.iso' and the use of the 'cat' command to write it to the USB device. A code block shows the command: '# cat debian-6.0.2.1-i386-businesscard.iso > /dev/sdc'. The article also notes that this feature is only available on i386/amd64 architectures and requires a network environment with DHCP. On the right side of the page, there are subscription buttons for '發表文章' and '留言', and a footer with the text '企鵝星球'.

FourDollars Blog
GNU/Linux, Debian/Ubuntu, Mac OS X, Free Software/Open Source Software, Freeware, and Programming.

首頁 Software Development Debian Package

2011年2月11日星期五

\$4 的 Debian 6.0 自動化安裝

製作可開機安裝 Debian 6.0 的 USB Stick

首先去下載 [debian-6.0.2.1-i386-businesscard.iso](#) (46M) 回來
然後再準備一個 USB Stick 接上電腦，如果 mount 在 Linux 系統上的話記得先 umount 掉
以 root 權限執行 (這裡假設 USB Stick 的 Device Node 位於 /dev/sdc)

```
# cat debian-6.0.2.1-i386-businesscard.iso > /dev/sdc
```

這樣就可以製作出可開機安裝 Debian 6.0 的 USB Stick 沒錯就是這樣製作的
這是 Debian 6.0 以後的新功能 hybrid CD/DVD image 不過只有在 i386/amd64 上面有作用

準備好要安裝的電腦跟可以透過有線網路 DHCP 上網的網路環境

現在這隻 USB Stick 可以接上任何可以使用 USB Stick 開機的電腦上開機就會看到：

訂閱
發表文章
留言
企鵝星球

Source: <http://fourdollars.blogspot.tw/2011/02/4-debian-60.html>

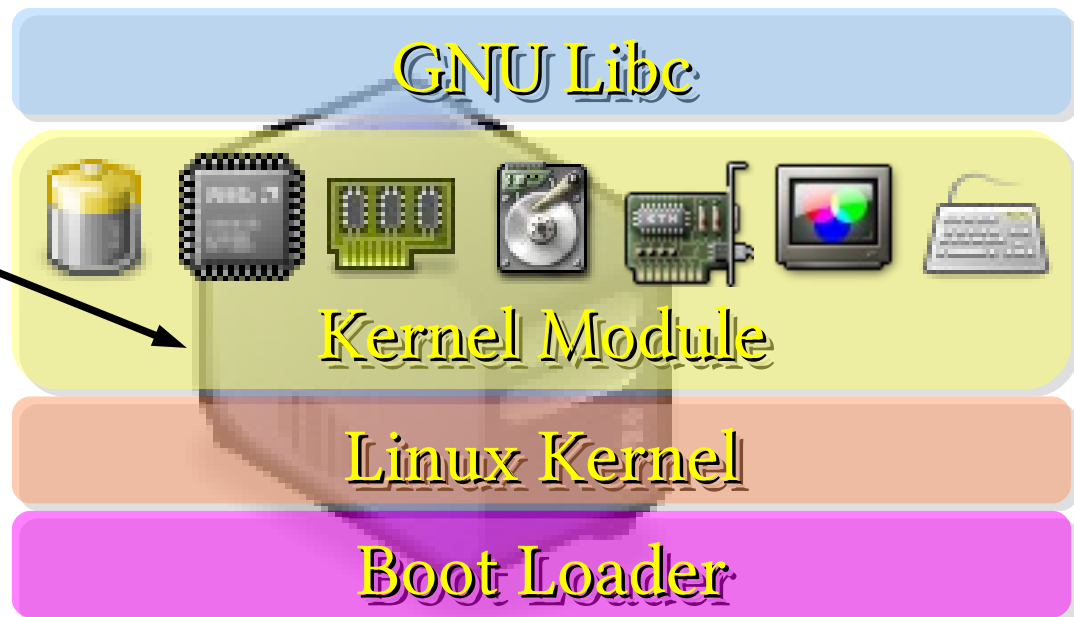
We install Base System of **GNU Debian Linux**
with Debian Installer and Preseed

According to <http://example.com/d-i/squeeze/preseed.cfg>

- It will install
- (1) Base Packages of Debian 6.0.4
 - (2) DRBL, JVM, Hadoop, etc....
 - (3) Run late_command script

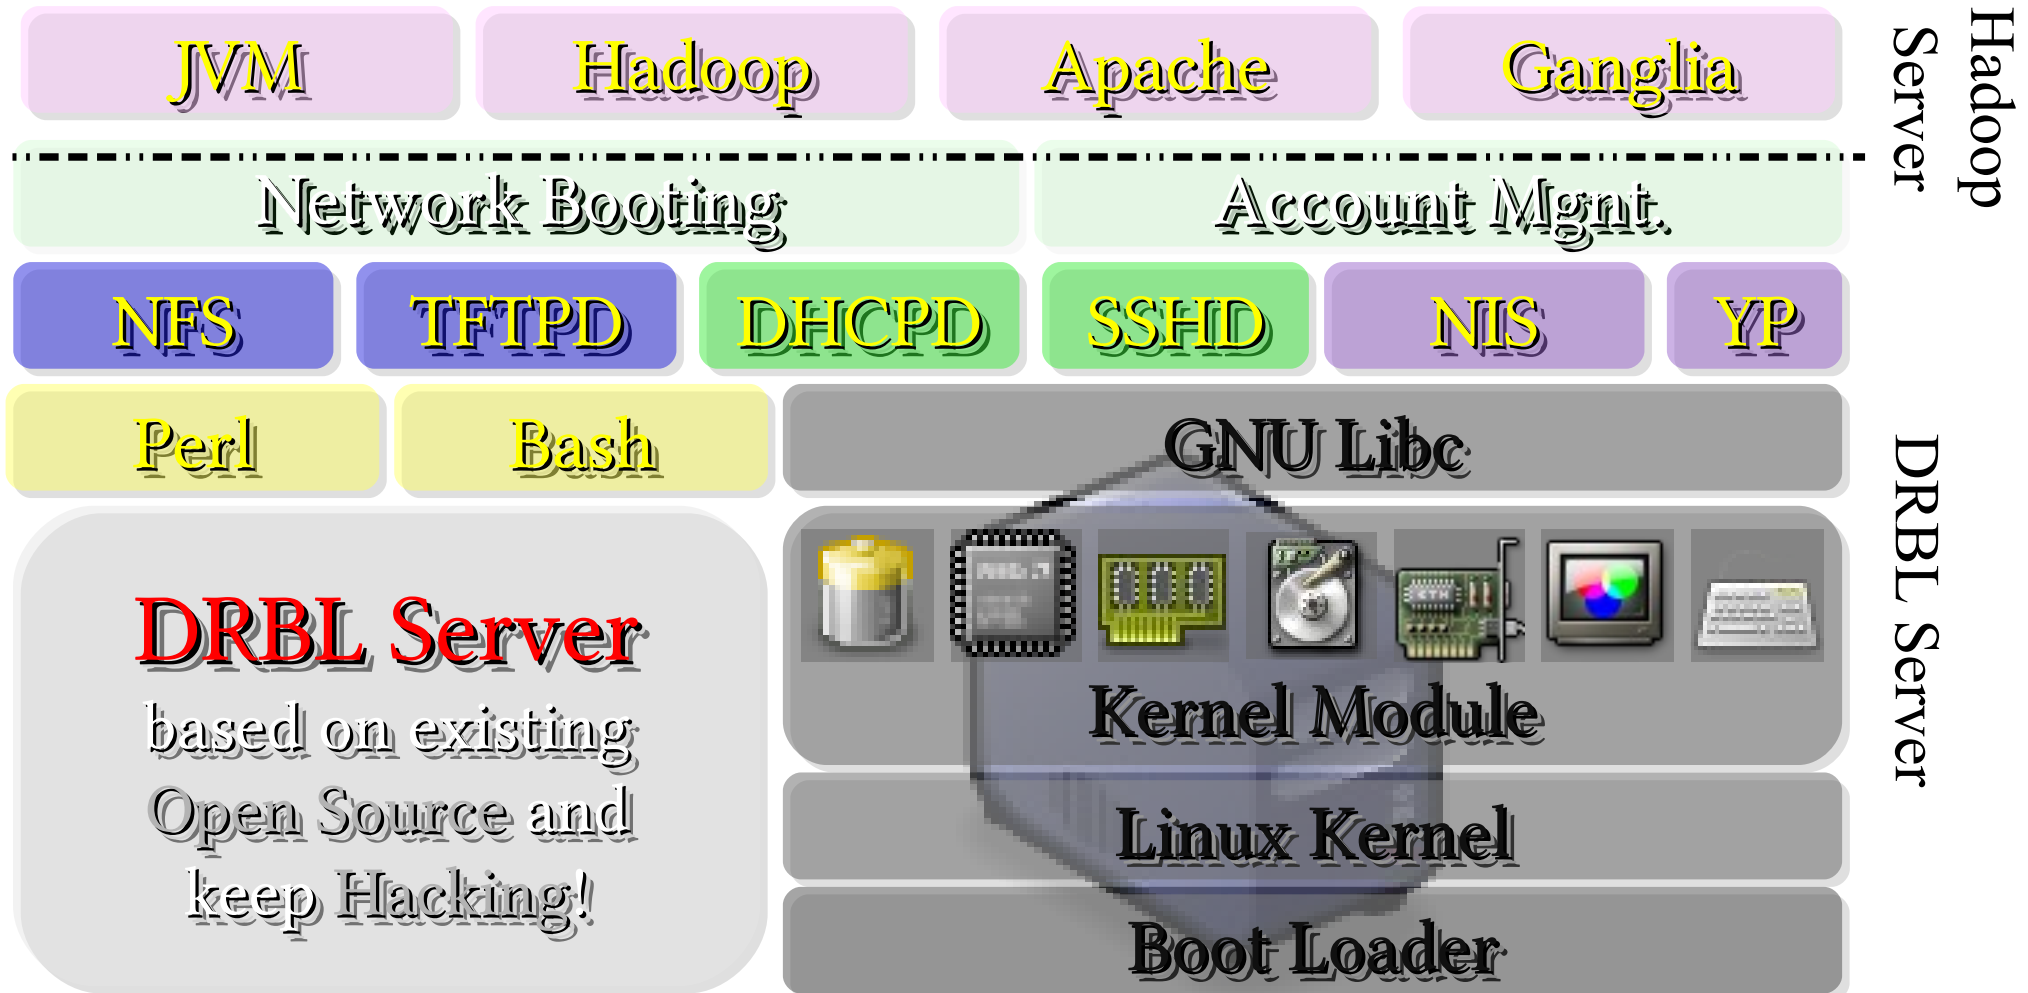


Debian Netinst CD



Reboot. It had installed **DRBL package** and **rc.local script** will configure it as **DRBL Server**.

There are lots of service needed:
SSH, DHCPD, TFTP, NFS Server, NIS Server, YP Server ...



DRBL Server
based on existing
Open Source and
keep Hacking!

rc.local script will run “drblsrv” & “drblpush”, there will be pxelinux, vmlinux-pxe, initrd-pxe in TFTPROOT, and different configuration files for each DRBL Client in NFSROOT

NFS

TFTPD

DHCPD

SSHD

NIS

YP

Config. Files
Ex. hostname

initrd-pxe

vmlinux-pxe

pxelinux

GNU Libc



Kernel Module

Linux Kernel

Boot Loader

Then, we enable **PXE** function in **BIOS** for each DRBL Client nodes.

BIOS PXE

BIOS PXE

BIOS PXE

BIOS PXE

NFS

TFTPD

DHCPD

SSHD

NIS

YP




Config. Files
Ex. hostname

initrd-pxe

vmlinuz-pxe

pxelinux

GNU Libc



Kernel Module

Linux Kernel

Boot Loader

While Booting, **PXE** will query IP address from **DHCPD**.

BIOS PXE

BIOS PXE

BIOS PXE

BIOS PXE

NFS

TFTPD

DHCPD

SSHD

NIS

YP




Config. Files
Ex. hostname

initrd-pxe

vmlinuz-pxe

pxelinux

GNU Libc



Kernel Module

Linux Kernel

Boot Loader

While Booting, **PXE** will query
IP address from **DHCPD**.

IP 1

IP 2

IP 3

IP 4

NFS

TFTPD

DHCPD

SSHD

NIS

YP

Config. Files
Ex. hostname

initrd-pxe

vmlinuz-pxe

pxelinux

GNU Libc



Kernel Module

Linux Kernel

Boot Loader

After PXE get its IP address, it will download booting files from TFTP.

IP 1

IP 2

IP 3

IP 4

NFS

TFTP

DHCP

SSH

NIS

YP




Config. Files
Ex. hostname

initrd-pxe

vmlinuz-pxe

pxelinux


GNU Libc



Kernel Module

Linux Kernel

Boot Loader



NFS **TFTPD** DHCPD SSHD NIS YP

Config. Files
Ex. hostname

initrd-pxe

vmlinuz-pxe

pxelinux


GNU Libc



Kernel Module

Linux Kernel

Boot Loader


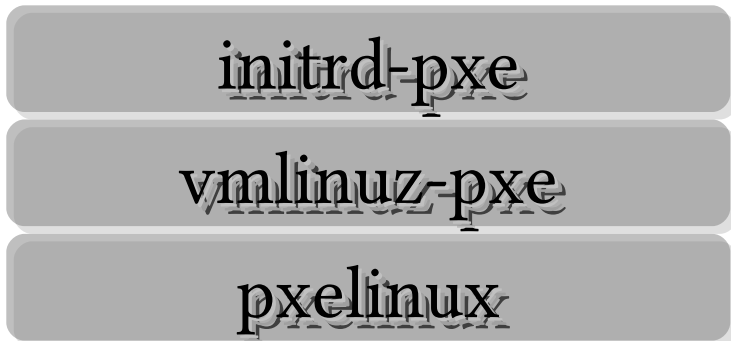





After downloading booting files, scripts in **initrd-pxe** will configure **NFSROOT** for each Compute Node.

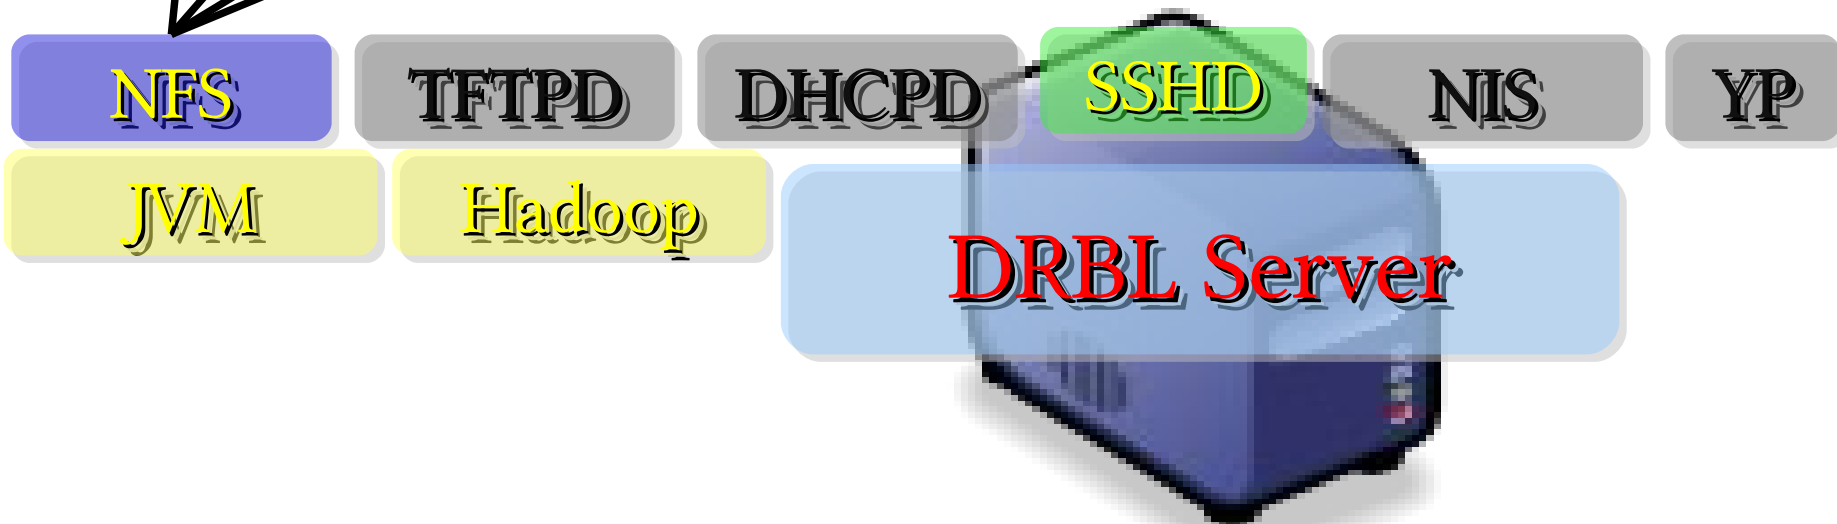
pxelinux

BOOT LOADER






Hadoop Services will also be deployed to each Compute Node via **NFS**






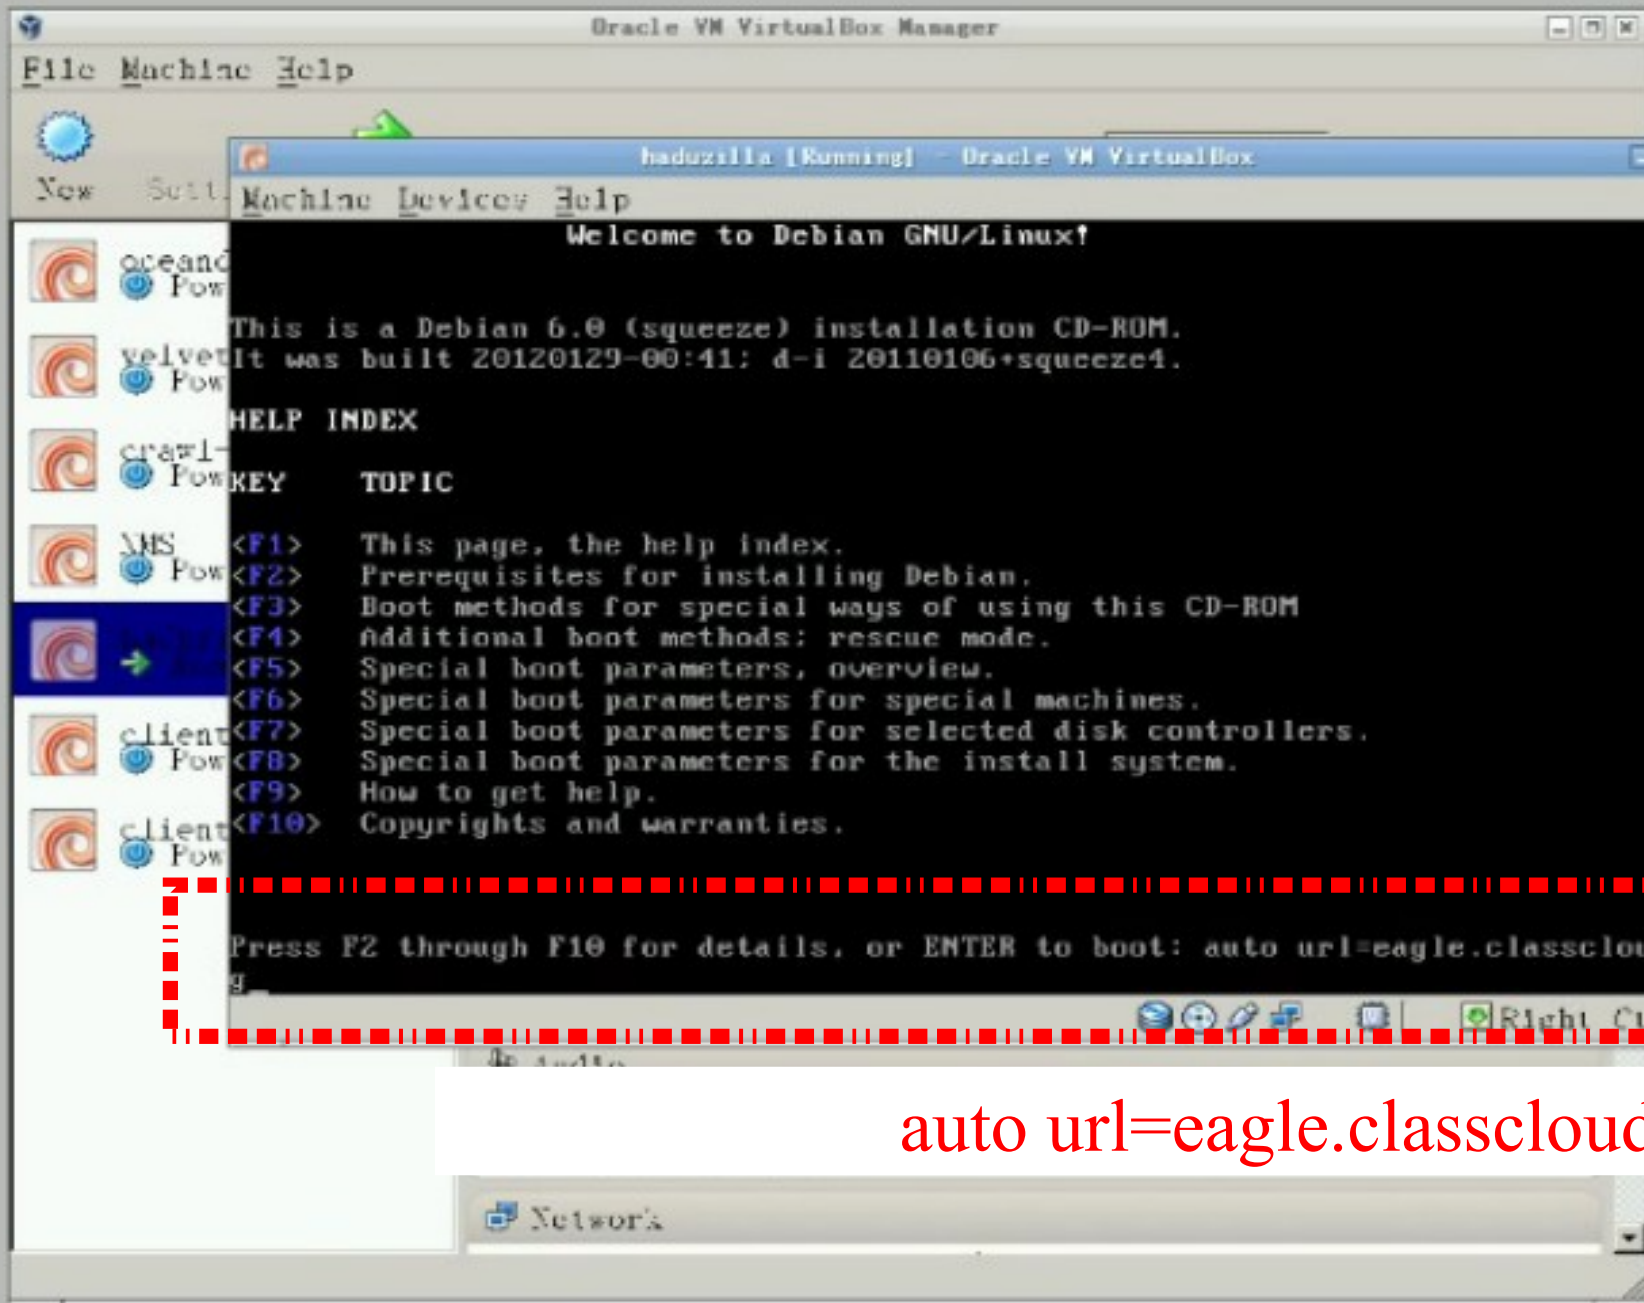
Demo

Jazz Wang
Yao-Tsung Wang
jazz@nchc.org.tw



Powered by **DRBL**





[auto url=eagle.classcloud.org](http://eagle.classcloud.org)

File Machine Help



New

Settings

Machine Devices Help

haduzilla [Running] - Oracle VM VirtualBox

oceanic
Powervelvet
Powercrawling
PowerMMS
Powerclient
Powerclient
Powerclient
Power

[!] Configure the network

From here you can choose to retry DHCP network autoconfiguration (which may succeed if your DHCP server takes a long time to respond) or to configure the network manually. Some DHCP servers require a DHCP hostname to be sent by the client, so you can also choose to retry DHCP network autoconfiguration with a hostname that you provide.

Network configuration method:

- Retry network autoconfiguration
- Retry network autoconfiguration with a DHCP hostname
- Configure network manually**

Do not configure the network at this time

<Go Back>

<Tab> moves; <Space> selects; <Enter> activates buttons



Right Ctrl

File Machine Help



New

Settings

Machine Devices Help

haduzilla [Running] - Oracle VM VirtualBox



ocean Power



velvet Power



crawling Power



MMS Power



Client Power



client Power



client Power

[!] Configure the network

The IP address is unique to your computer and consists of four numbers separated by periods. If you don't know what to use here, consult your network administrator.

IP address:

192.168.125.5

<Go Back>

<Continue>

<Tab> moves; <Space> selects; <Enter> activates buttons



Right Ctrl

File Machine Help



New

Settings

Machine Devices Help

haduzilla [Running] - Oracle VM VirtualBox



ocean Power



velvet Power



crawling Power



MMS Power



Client Power



client Power



client Power

[!] Configure the network

The netmask is used to determine which machines are local to your network. Consult your network administrator if you do not know the value. The netmask should be entered as four numbers separated by periods.

Netmask:

255.255.255.0

<Go Back>

<Continue>

<Tab> moves; <Space> selects; <Enter> activates buttons



Right Ctrl

File Machine Help



New

Settings

Machine Devices Help

haduzilla [Running] - Oracle VM VirtualBox



ocean Power



velvet Power



crawling Power



MMS Power



Client Power



client Power



client Power

[!] Configure the network

The gateway is an IP address (four numbers separated by periods) that indicates the gateway router, also known as the default router. All traffic that goes outside your LAN (for instance, to the Internet) is sent through this router. In rare circumstances, you may have no router; in that case, you can leave this blank. If you don't know the proper answer to this question, consult your network administrator.

Gateway:

192.168.125.254

<Go Back>

<Continue>

<Tab> moves; <Space> selects; <Enter> activates buttons



Right Ctrl

File Machine Help



New

Settings

ocean
Powervelvet
Powercrawling
PowerVMS
Powerclient
Powerclient
Powerclient
Power

haduzilla [Running] - Oracle VM VirtualBox



Machine Devices Help

[!] Configure the network

The name servers are used to look up host names on the network. Please enter the IP addresses (not host names) of up to 3 name servers, separated by spaces. Do not use commas. The first name server in the list will be the first to be queried. If you don't want to use any name server, just leave this field blank.

Name server addresses:

3.8.8.8

<Go Back>

<Continue>

<Tab> moves; <Space> selects; <Enter> activates buttons



Right Ctrl


```
haduzilla [Running] - Oracle VM VirtualBox
Machine Device Help
Checking initscripts... installed.
Checking tftpd-hpa... installed.
Checking nfs-kernel-server... installed.
Checking nis... installed.
Checking curl... installed.
Checking lftp... installed.
Checking iptables... installed.
Checking libdigest-sha1-perl... installed.
Checking memtest86+... installed.
Checking ntfs-3g... installed.
Checking clonezilla... installed.
Checking mkswap-uuid... installed.
Checking partclone... installed.
Checking drbl-chntpw... installed.
Checking mkpxenitrd-net... installed.
Checking ipxe... installed.
Checking freedos... installed.
Checking udev... installed.
*****
*****
*****
Installing kernel for clients... ..
The kernel for client is copied from server.
Installing kernel 2.6.32-5-686 for clients...
It might take several minutes to install this kernel, please be patient... ..
```

- File Machine Settings
- New Settings
- oceancb Powered Off
- velvet Powered Off
- crawl-dev Powered Off
- MMS Powered Off
- client1** →
- client1 Powered Off
- client2 Powered Off

| | |
|--------------------------------|-------------------------------------|
| (48.00 MB) | |
| SATA Controller | |
| SATA Port 0: | haduzilla.vdi (Normal, 10.00 GB) |
| Audio | |
| Host Driver: ALSA Audio Driver | |
| Controller: ICH AC97 | |
| Networks | |

haduzilla [Running] - Oracle VM VirtualBox

Machine Devices Help

File Machine Settings

New Settings

oceancb Powered Off

velvet Powered Off

crawi-dev Powered Off

MMS Powered Off

client1 Powered Off

client2 Powered Off

```

Syncing /lib/modules/2.6.32-5-686 to client's common root...
Syncing /boot/*-2.6.32-5-686* to client's common root...
Generating the /tftpboot/node_root/lib/modules/2.6.32-5-686/modules.dep
Syncing /lib/firmware/ to client's common root...
Copying the directory /etc/ to clients common root /tftpboot/node_root...
Cleaning the ssh key file ssh_host_dsa_key copied from server... done!
Cleaning the ssh key file ssh_host_dsa_key.pub copied from server... done!
Cleaning the ssh key file ssh_host_rsa_key copied from server... done!
Cleaning the ssh key file ssh_host_rsa_key.pub copied from server... done!
Commenting the TCPWrapper related file /tftpboot/node_root/etc/hosts.deny copied
from server... done!
Commenting the TCPWrapper related file /tftpboot/node_root/etc/hosts.allow copie
d from server... done!
Warning! Unable to find the fine-tune file /opt/drbl/setup/files/DBN/DBN6.0.5/fi
rstboot.DBN6.0.5.drbl, use /opt/drbl/setup/files/DBN/firstboot.default-DBN.drbl
as /etc/init.d/firstboot for DRBL clients!
This may cause some problems to DRBL clients!
The startup services for DRBL client are:
firstboot portmap nis nfs-common ssh dbus kbd acpid drblthincli mkswapfile arm-w
ol sendsigs umountfs
Using udev for clients... The default display manager is NOT found! We can NOT s
et text or graphic mode for Debian DRBL client.
Deleting the accounts (except root) in the clients common root template... done!
Enabling the NIS client in the common root template... done!
Creating some necessary files in the clients common root template.....
  
```

Right Ctrl

SATA Controller (48.00 MB)

SATA Port 0: haduzilla.vdi (Normal, 10.00 GB)

Audio

Host Driver: ALSA Audio Driver

Controller: ICH AC97

Networks

```
Machine Device Help
File Machine Settings
New Settings
oceancb Powered
velvet Powered
crawl-dev Powered
MMS Powered
client1 Powered
client2 Powered

Creating DRBL client: debian0-192 10.0.0.92... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-193 10.0.0.93... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-194 10.0.0.94... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-195 10.0.0.95... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-196 10.0.0.96... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-197 10.0.0.97... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-198 10.0.0.98... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-199 10.0.0.99... Pseudo client is created for DRBL
SSI or clonezilla box mode! done!
Creating DRBL client: debian0-1100 10.0.0.100... Pseudo client is created for DR
BL SSI or clonezilla box mode! done!
Template client for DRBL SSI, Clonezilla box mode or Clonezilla live client is 1
0.0.0.1
Using template host /tftpboot/nodes/10.0.0.1
Generating SSH host keys for client 10.0.0.1 if they do not exist... done!
Generating the files for DRBL single system image template... root... etc... var
```

(48.00 MB)

SATA Controller
SATA Port 0: haduzilla.vdi
(Normal, 10.00 GB)

Audio
Host Driver: ALSA Audio Driver
Controller: ICH AC97

Networks

- General
- System
- Display
- Storage
- Audio
- Network**
- Serial Ports
- USB
- Shared Folders

Network

Adapter 1 | Adapter 2 | Adapter 3 | Adapter 4

Enable Network Adapter

Attached to: Bridged Adapter

Name: tap0

Advanced

Adapter Type: PCnet-Fast III (Am79C973)

Mac Address: 086027F28189

Cable connected

Port Forwarding

Select a settings category from the list on the left-hand side and move the mouse over a settings item to get more information.

OK Cancel Help

File Machine Help

New Settings

- oceanic Power
- velvet Power
- crawling Power
- VMS Power
- haduzilla *Running*
- client1** *Running*
- client Power

Machine Devices Help

DRBL (<http://drbl.nchc.org.tw>, <http://drbl.sf.net>)

Debian 6.0.5 Linux (DRBL mode, mostly local resources)

Local operating system (if available)

Memory test using Memtest86+

Press [Tab] to edit options

Automatic boot in 7 seconds...

- DRBL version: 1.10.99-1drbl. (C) 2003-2012, NCHC, Taiwan
- Disclaimer: DRBL comes with ABSOLUTELY NO WARRANTY

DRBL

Free Software Labs NCHC, Taiwan
 自由軟體實驗室 國家高速網路與計算中心



Ctrl

Right Ctrl

USB

Show me the Source Code !

<http://eagle.classcloud.org/d-i/squeeze/>

New versions will be at github

<http://github.com/jazzwang>

Attribution-Noncommercial-Share Alike 3.0 Taiwan



姓名標示-非商業性-相同方式分享 3.0 台灣

您可自由：



分享 — 重製、散布及傳輸本著作



重混 — 修改本著作

惟需遵照下列條件：



姓名標示 — 您必須按照著作人或授權人所指定的方式，表彰其姓名（但不得以任何方式暗示其為您或您使用本著作的方式背書）。



非商業性 — 您不得為商業目的而使用本著作。



相同方式分享 — 若您變更、變形或修改本著作，您僅得依本授權條款或與本授權條款類似者來散布該衍生作品。

<http://creativecommons.org/licenses/by-nc-sa/3.0/tw/>

These slides could be distributed by Creative Commons License.



Questions?

Slides - <http://trac.nchc.org.tw/cloud>

Jazz Wang
Yao-Tsung Wang
jazz@nchc.org.tw



Powered by DRBL