




# Hadoop 叢集設定解說

*Setup Hadoop Fully Distributed Mode*

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



# Yahoo's Hadoop Cluster

## 雅虎的大象軍團

- ~10,000 machines running Hadoop in US
- The largest cluster is currently 2000 nodes
- Nearly 1 petabyte of user data (compressed, unreplicated)
- Running roughly 10,000 research jobs / week



# Hadoop Pseudo-Distributed Mode

## 我們已經實作過單機模式

- Step 1: Setup SSH key exchange
- Step 2: Install Java
- Step 3: Download Hadoop Source Package
- Step 4: Configure `hadoop-env.sh`
  - `export JAVA_HOME=/usr/lib/jvm/java-6-sun`
- Step 5: Configure `*-site.xml`
  - Set Namenode to `hdfs://localhost:9000`
  - Set Jobtracker to `localhost:9001`
  - `bin/hadoop namenode -format`
- Step 6: Format HDFS
- Step 7: Start Hadoop
  - `bin/start-all.sh`
- Step 8: Complete!! Let's check the status of Hadoop
  - Job admin <http://localhost:50030/> HDFS <http://localhost:50070/>



# Diagram of Pseudo-Distributed Mode

## Hadoop 單機環境示意圖

Node 1

conf / core-site.xml:

fs.default.name ->  
hdfs://localhost:9000

conf / mapred-site.xml:

mapred.job.tracker ->  
localhost:9001

conf/slaves:

localhost

Localhost

Namenode


JobTracker

Datanode

Tasktracker

localhost:50070

localhost:50030



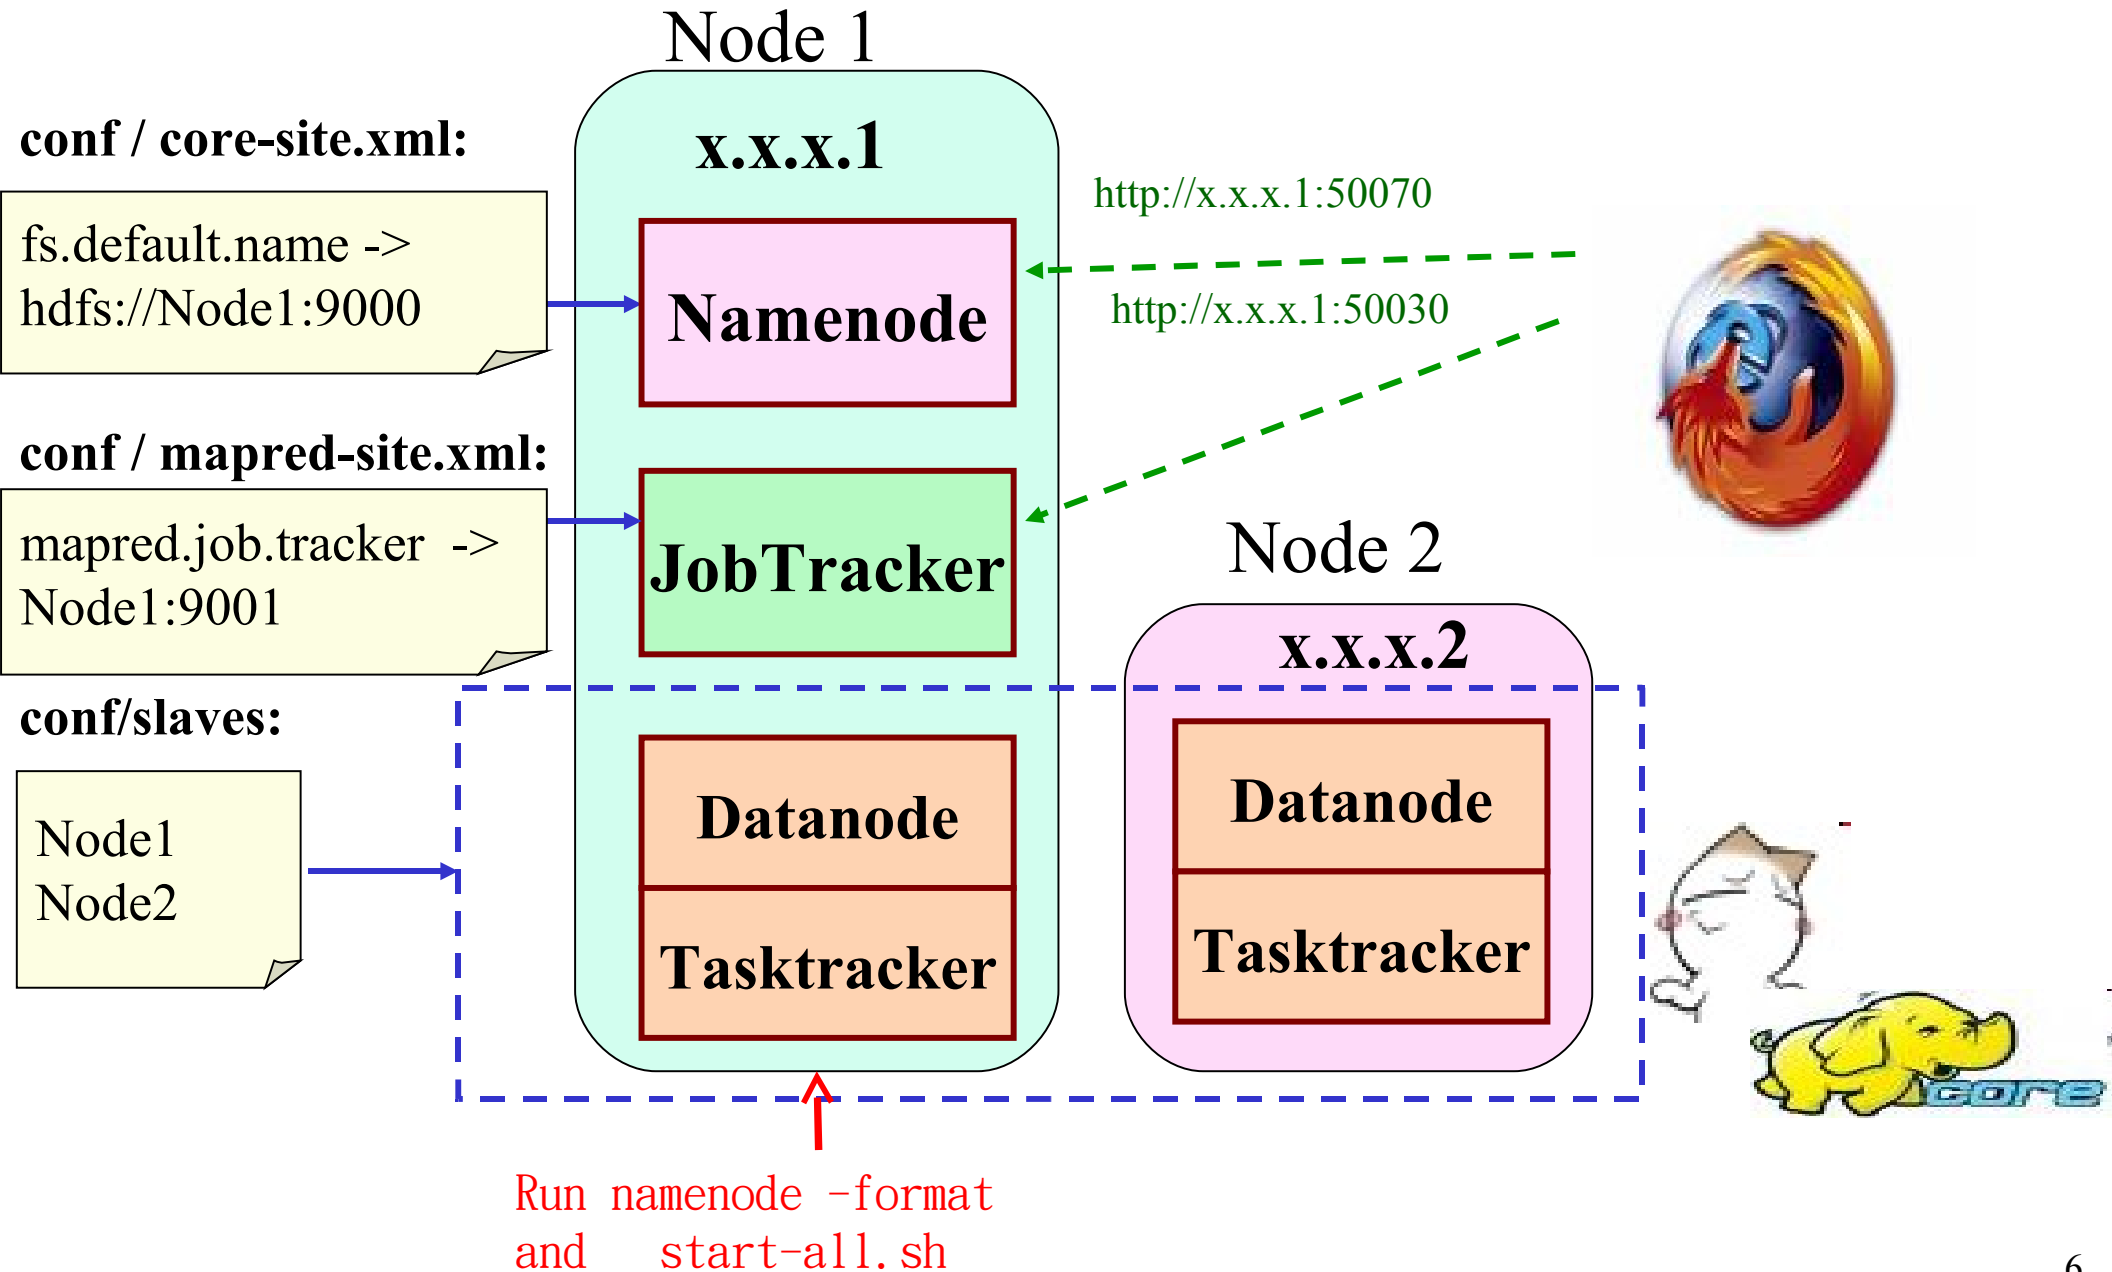
# Hadoop Fully-Distributed Mode

## 我們接著要用兩台電腦實作叢集模式

- Step 1: Setup SSH key exchange
- Step 2: Install Java
- Step 3: Download Hadoop Source Package
- Step 4: Configure hadoop-env.sh
  - export JAVA\_HOME=/usr/lib/jvm/java-6-sun
- Step 5: Configure \*-site.xml
  - Set Namenode to hdfs://x.x.x.1:9000
  - Set Jobtracker to x.x.x.2:9001
- Step 6: Configure Slaves
- Step 7: Synchronization of all slaves
- Step 8: Format HDFS
  - bin/hadoop namenode -format
- Step 9: Start Hadoop
  - On NameNode : bin/start-dfs.sh
  - On JobTracker : bin/start-mapred.sh
- Step 10: Complete!! Let's check the status of Hadoop
  - Job admin <http://x.x.x.2:50030/> HDFS <http://x.x.x.1:50070/>

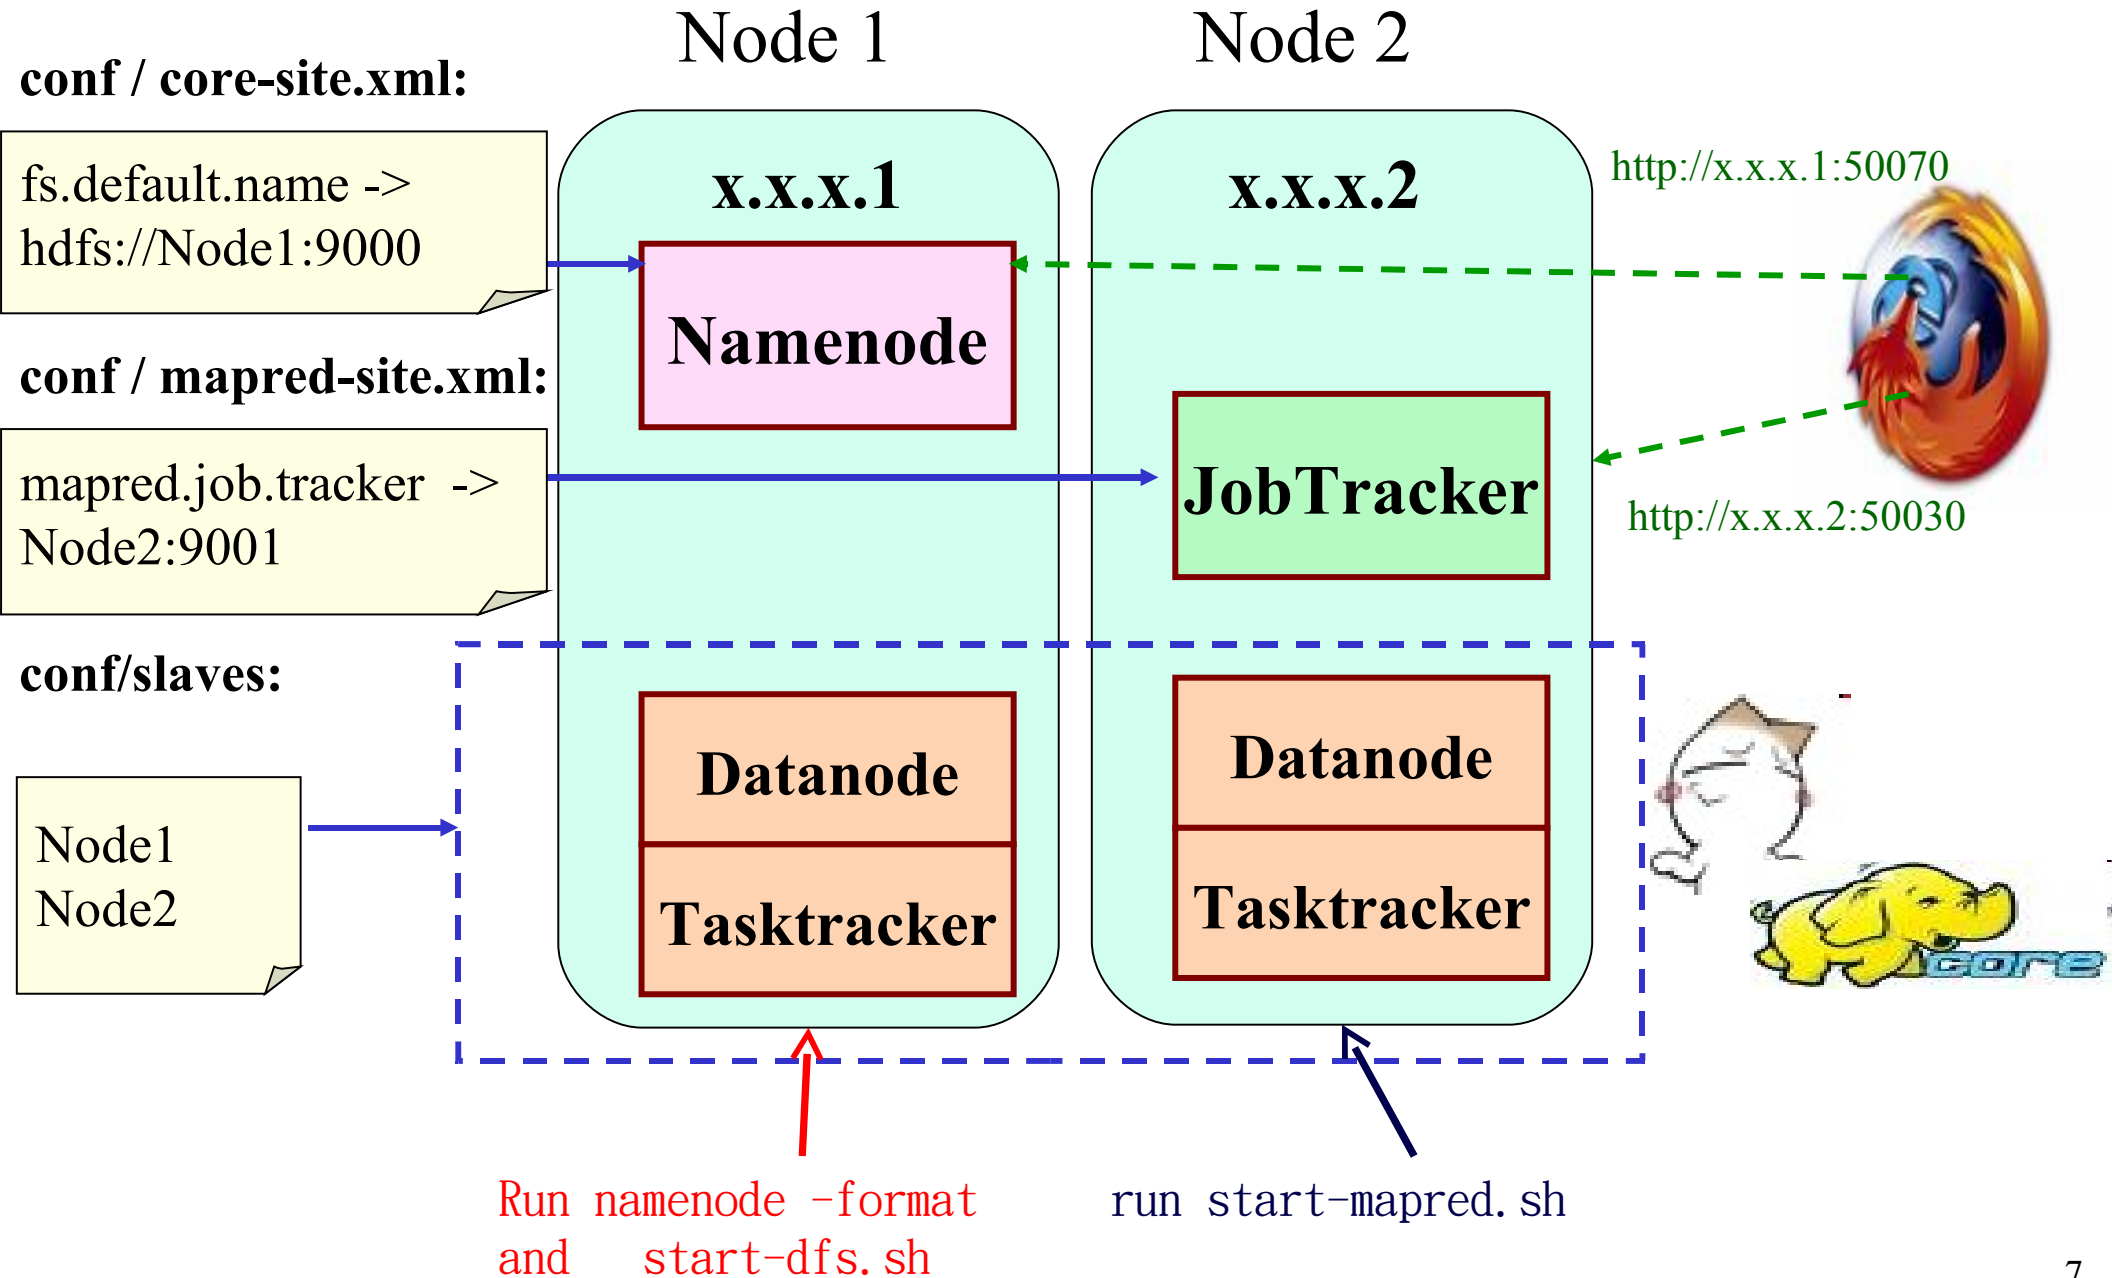
# Use case #1

## 設定情境一



# Use case #2

## 設定情境二



# Use case #3

## 設定情境三


conf / core-site.xml:

fs.default.name ->  
hdfs://Node1:9000

conf / mapred-site.xml:

mapred.job.tracker ->  
Node1:9001

Node 1



http://x.x.x.1:50070

http://x.x.x.1:50030



conf/slaves:

Node2  
.....  
NodeN

Node 2

x.x.x.2

Datanode

Tasktracker


Node N

x.x.x.n

Datanode

Tasktracker

...






# Use case #4

## 設定情境四

### conf / core-site.xml:

fs.default.name ->  
hdfs://Node1:9000

Client




http://x.x.x.2:50030

### conf / mapred-site.xml:

mapred.job.tracker ->  
Node2:9001

G





HTTP Monitoring UI

http://x.x.x.1:50070

### conf/slaves:

Node3  
.....  
NodeN






# Hadoop 叢集佈署工具

*Hadoop Deployment Tool : SmartFog and DRBL*

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



# **Programmer v.s. System Admin.**



Source: <http://www.funnyjunksite.com/wp-content/uploads/2007/08/programmer.jpg>



Source: <http://www.sysadminday.com/images/people/136-3697.JPG>



**PART 1 :**

# ***PC Cluster 101***

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



Powered by **DRBL**



*At First, We have "4 + 1" PC Cluster*

*It'd better be*  
**2<sup>n</sup>**




*Manage*  
**Scheduler**

*Then, We connect 5 PCs with  
**Gigabit Ethernet Switch***



***GiE Switch***



*10/100/1000  
Mbps*



***Add 1 NIC  
for WAN***



**Compute Nodes**

**4 Compute Nodes will communicate via LAN Switch. Only Manage Node have Internet Access for Security!**



**WAN**



**Manage Node**



# **Compute Nodes**

## **Basic System Setup for Cluster**

*Messaging*

**MPICH**

*Account Mgmt.*

**SSHD**

**NIS**

**YP**

**GCC**

**GNU Libc**

**Bash**

**Perl**



**Kernel Module**

**Linux Kernel**

**Boot Loader**



On **Manage Node**,  
We need to install **Scheduler** and  
**Network File System** for sharing  
Files with **Compute Node**

*Job Mgmt.*

**OpenPBS**

*File Sharing*

**NFS**

**Extra**

*Messaging*

**MPICH**

**GCC**

**Bash**

**Perl**

*Account Mgmt.*

**SSHD**

**NIS**

**YP**

**GNU Libc**



**Kernel Module**

**Linux Kernel**

**Boot Loader**

# Challenges of Cluster Computing

- **Hardware**

- **Ethernet Speed | PC Density**
- **Power | Cooling | Heat**
- **Network and Storage Architecture**

- **Software**

- **Job Scheduler ( Cluster level )**
- **Account Management**
- **File Sharing | Package Management**

- **Limitation**

- **Shared Memory**
- **Global Memory Management**

# *Common Method to deploy Cluster*



**1. Setup one  
Template  
machine**

**2. Cloning  
to  
multiple  
machine**



**3. Configure  
Settings**



**4. Install  
Job  
Scheduler**



**5. Running  
Benchmark**

# **Challenges of Common Method**

**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

<b>Total Nodes</b>	<b>4000</b>
<b>Total cores</b>	<b>30000</b>
<b>Data</b>	<b>16PB</b>

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

# Advanced Methods to deploy Cluster

- **SSI ( Single System Image )**
  - **Multiple PCs as Single Computing Resources**
  - **Image-based**
    - **homogeneous**
    - **ex. SystemImager, OSCAR, Kadeploy**
  - **Package-based**
    - **heterogeneous**
    - **easy update and modify packages**
    - **ex. FAI, DRBL**
- **Other deploy tools**
  - **Rocks : RPM only**
  - **cfengine : configuration engine**

# Comparison of Cluster Deploy Tools

	<i>Distribution</i>	<i>Support Diskless/ Sysmless</i>	<i>Type</i>	<i>Node configuration tools</i>	<i>Cluster management tools</i>	<i>Database installation</i>
<i>System Imager</i>	<i>ALL</i>	<i>Yes</i>	<i>Image</i>	<i>Yes</i>	<i>No</i>	<i>No</i>
<i>OSCAR</i>	<i>RPM- based</i>	<i>Yes</i>	<i>Image</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>
<i>Kadeploy</i>	<i>ALL</i>	<i>No</i>	<i>Image</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
<i>DRBL</i>	<i>ALL</i>	<i>Yes</i>	<i>Package</i>	<i>Yes</i>	<i>Yes</i>	<i>No</i>
<i>FAI</i>	<i>Debian- Based</i>	<i>Yes</i>	<i>Package</i>	<i>Yes</i>	<i>No</i>	<i>No</i>



**PART 2-1 :**

# ***Hadoop Deployment Tool***

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



Powered by **DRBL**





- 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](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)

12 June 2008

# SmartFrog - HPLabs' CM tool

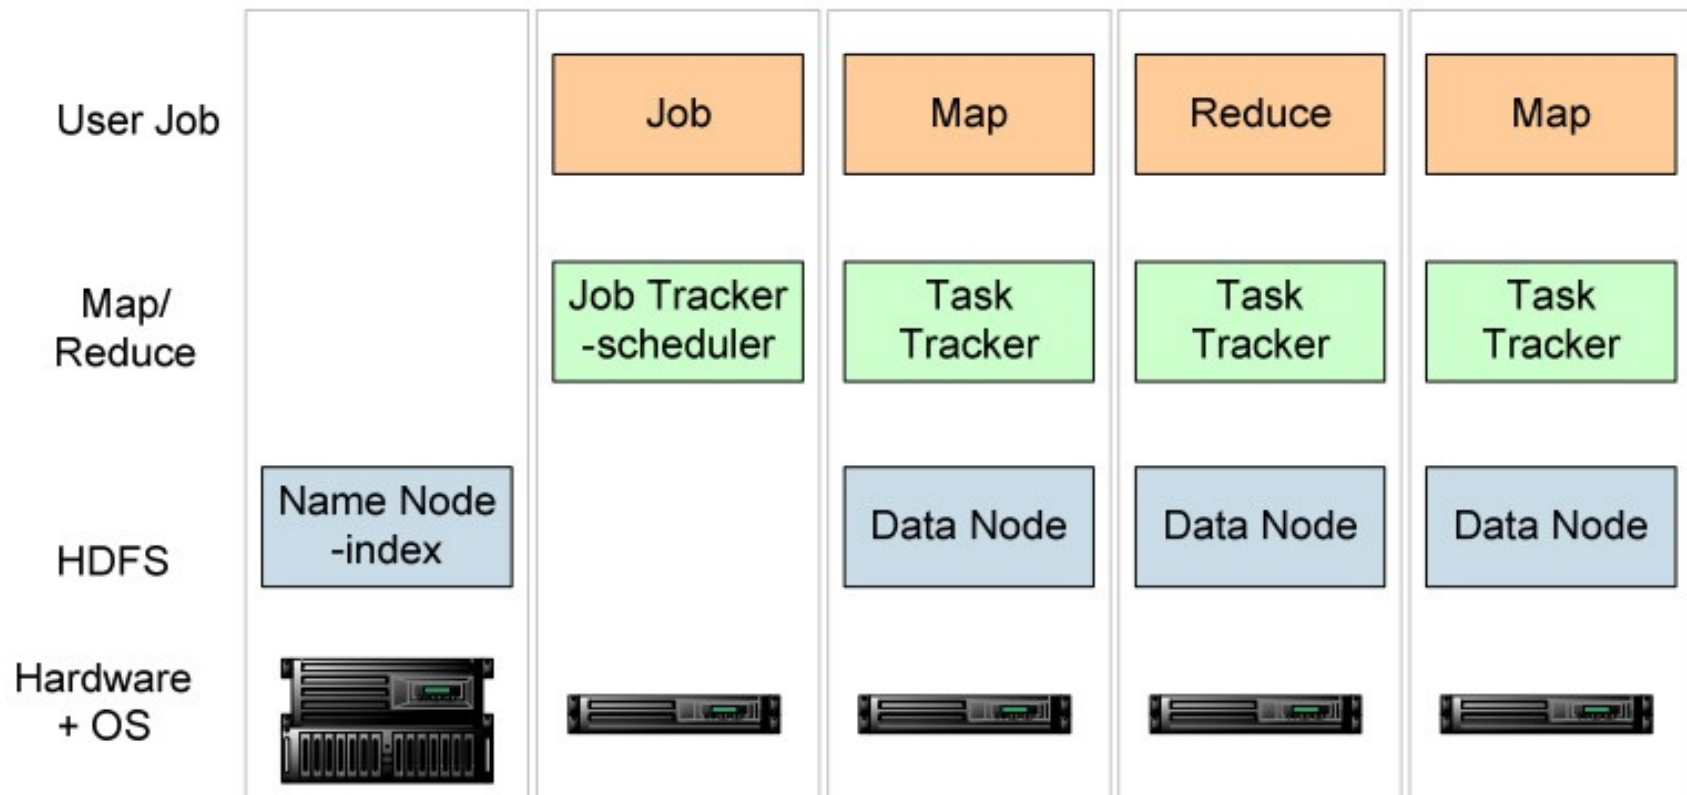
- Language for describing systems to deploy  
—everything from datacentres to test cases
  - Runtime to create *components* from the model
  - Components have a lifecycle
  - LGPL Licensed, Java 5+
- <http://smartfrog.org/>

Source: Deploying hadoop with smartfrog

12 [http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# Basic problem: deploying Hadoop



*one namenode, 1+ Job Tracker, many data nodes and task trackers*

Source: Deploying hadoop with smartfrog

12 [http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)

# The hand-managed cluster

- Manual install onto machines
- SCP/FTP in Hadoop zip
- copy out hadoop-site.xml and other files
- edit /etc/hosts, /etc/rc5.d, SSH keys ...
- Installation scales  $O(N)$
- Maintenance, debugging scales worse

Source: Deploying hadoop with smartfrog

12 [http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# The locked-down cluster

- PXE Preboot of OS images
- RedHat Kickstart to serve up (see [instalinux.com](http://instalinux.com))
- Maybe: LDAP to manage state, or custom RPMs

## Requires:

uniform images, central LDAP service, good ops team, stable configurations, home-rolled RPMs

**Source: Deploying hadoop with smartfrog**

[http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# CM-tool managed cluster

## Configuration Management tools

- State Driven: observe system state, push it back into the desired state
- Workflow: apply a sequence of operations to change a machine's state
- Centralized: central DB in charge
- Decentralized: machines look after themselves

CM tools are the only way to manage big clusters

Source: [Deploying hadoop with smartfrog](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)

12 [http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# Model the system in the SmartFrog language

```
TwoNodeHDFS extends OneNodeHDFS {  
  
    localDataDir2 extends TempDirwithCleanup {  
  
    }  
  
    datanode2 extends datanode {  
        dataDirectories [LAZY localDataDir2];  
        dfs.datanode.https.address "https://localhost:0";  
    }  
}
```

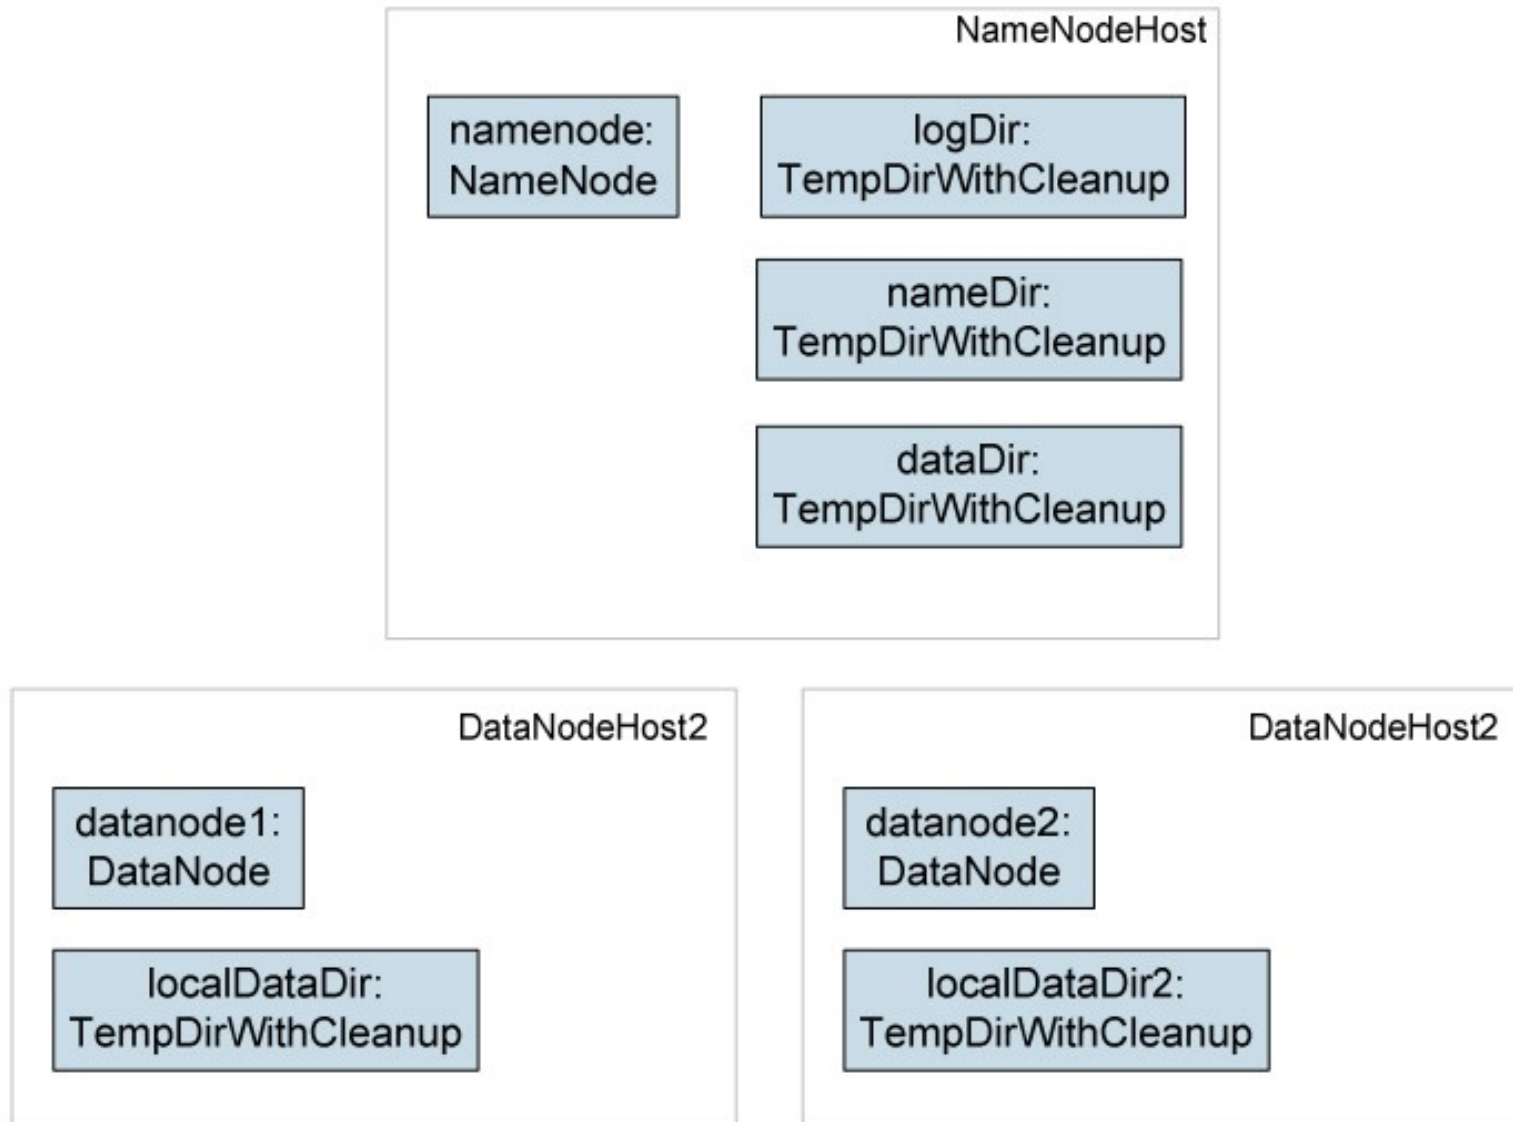
Inheritance, cross-referencing, templating

Source: [Deploying hadoop with smartfrog](#)

12 [http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# The runtime deploys the model



Source: Deploying hadoop with smartfrog

[http://people.apache.org/~stevell/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevell/slides/deploying_hadoop_with_smartfrog.pdf)



# Steps to deployability

1. Configure Hadoop from an SmartFrog description
2. Write components for the Hadoop nodes
3. Write the functional tests
4. Add *workflow* components to work with the filesystem; submit jobs
5. Get the tests to pass

Source: Deploying hadoop with smartfrog

12 [http://people.apache.org/~stevel/slides/deploying\\_hadoop\\_with\\_smartfrog.pdf](http://people.apache.org/~stevel/slides/deploying_hadoop_with_smartfrog.pdf)





**PART 2-2 :**

# ***Introduction to DRBL***

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



Powered by **DRBL**

# What is DRBL ??

- **Diskless Remote Boot in Linux**
- Network is cheap, and our time is expansive
- In simple words, DRBL is .....
  - Replace IDE/SATA cable with network cable
  - 40+ student PCs connected to one DRBL server



**Diskfull  
PC**



=



+



+




**Diskless  
PC**



Server

**1st, We install Base System of **GNU/Linux** on **Management Node**. You can choose:**

**Redhat, Fedora, CentOS, Mandriva, Ubuntu, Debian, ...**



*2nd, We install **DRBL package** and  
configure it as **DRBL Server**.*

*There are lots of service needed:  
**SSHD, DHCPD, TFTPD, NFS Server,  
NIS Server, YP Server ...***

*Network Booting*

*Account Mgmt.*

**NFS**

**TFTPD**

**DHCPD**

**SSHD**

**NIS**

**YP**


**Perl**

**Bash**

**GNU Libc**

**DRBL Server**

*based on existing  
Open Source and  
keep Hacking!*



**Kernel Module**

**Linux Kernel**

**Boot Loader**

*After running “**drblsrv -i**” & “**drblpush -i**”, there will be **pxelinux**, **vmlinux-pex**, **initrd-pxe** in TFTPROOT, and different **configuration files** for each Compute Node 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**

**3rd, We enable *PXE* function in *BIOS* configuration.**

***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 **TFTPD**.**

**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**



**NFS**   **TFTPD**   **DHCPD**   **SSHD**   **NIS**   **YP**

**Config. Files**  
**Ex. hostname**

**initrd-pxe**  
**vmlinuz-pxe**  
**pxelinux**





**NFS**   **TFTPD**   **DHCPD**   **SSHD**   **NIS**   **YP**

*Config. Files*   *GNU Libc*

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

**pxelinux**

**Boot Loader**



- NFS**
- TFTPD**
- DHCPD**
- SSHD**
- NIS**
- YP**


**Config. Files**  
 Ex. hostname

- initrd-pxe
- vmlinuz-pxe
- pxelinux





**Applications and Services will also  
deployed to each Compute Node  
via *NFS* ....**






*With the help of **NIS** and **YP**,  
You can login each Compute Node  
with the **Same ID | PASSWORD**  
stored in DRBL Server!*

**SSH Client**



**DRBL Server**





***Part 3 : How we use DRBL  
to deploy Cloud Testbed ?***

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



Powered by **DRBL**



# Building IaaS using DRBL-Xen

## Application

Social Computing, Enterprise, ISV, ...

eyeOS, Nutch, ICAS,  
X-RIME, ...

## Programming

Web 2.0, Mashups, Workflows, ...

Hadoop (MapReduce),  
Sector/Sphere, AppScale

## Management

Qos Negotiation, Admission Control,  
Pricing, SLA Management, Metering...

OpenNebula, Enomaly,  
Eucalyptus, OpenQRM, ...

## Virtualization

VM, VM management and Deployment


Xen, KVM, VirtualBox,  
QEMU, OpenVZ, ...

## Physical Hardware

Infrastructure: Computer, Storage, Network

# Virtualization ?? Emulator ??

## Virtual Hardware / OS



mame4iphone

Latest Version:

MAME 0.138 15 May 10

— DOWNLOAD NOW —

[Download source updates to MAME 0.138](#)

Mac4Lin

QEMU



## Physical Hardware / OS

# ***What is Virtualization ??***

Application Virtualization **Ex. VMWare ThinApp**

Desktop Virtualization  
Client Virtualization **Ex. XenDesktop**

Presentation Virtualization **Ex. VNC, M\$ RDP**


OS-level Virtualization **Ex. Xen, KVM**

Network Virtualization **Ex. OpenFlow**

Storage Virtualization **Ex. NetApp**

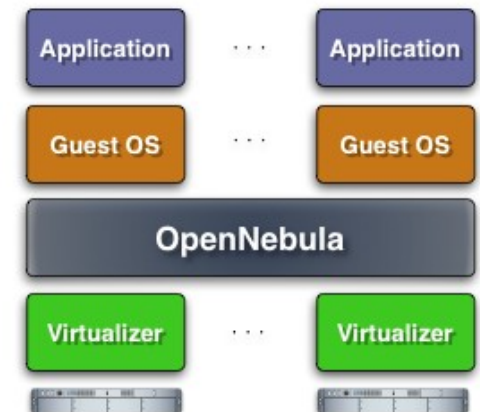

Database Virtualization

Data Virtualization



- <http://open.eucalyptus.com/>
- It was a research project of UCSB, USA
- Now Eucalyptus System provide technical supports.
- It designed to help user **to build their own Amazon EC2**
- Its feature is compatible with existing EC2 client.
- [Ubuntu Enterprise Cloud powered by Eucalyptus in 9.04](#)
- You can register trail account at <http://open.eucalyptus.com/>
- Cons : you might need to type commands in some case

- <http://www.opennebula.org>
- Sponsor by European Union FP7
- Turn Physical Cluster into Virtual Cluster
- manage **status, scheduling and migration of virtual cluster**
- [Ubuntu 9.04 provide package of opennebula](#)
- Cons : You need to type commands to check or migration



# Building IaaS using DRBL-Xen

- DRBL-Xen is still need more work to intergrate into DRBL
- Manual procedure could be found at
  - [http://trac.nchc.org.tw/grid/wiki/jazz/DRBL\\_Xen](http://trac.nchc.org.tw/grid/wiki/jazz/DRBL_Xen)

The screenshot displays the Trac web interface. At the top left is the Trac logo with the text "trac Integrated SCM & Project Management". To the right is a search bar and the text "logged in as jazz". Below this is a navigation menu with links: Wiki, Timeline, Roadmap, Browse Source, View Tickets, New Ticket, Search, Admin. Further right are links: Start Page, Index, History, Last Change.

The main content area shows a ticket titled "Massive Deployment of Xen enabled Virtual Cluster using Diskless Remote Boot Linux". Under the title, it says "Test Environment". Below that is a list of bullet points:

- VMWare Server 1.0.3
- VMWare Image debian 4.0r1 as DRBL Server

To the right of the main content is a yellow box containing a list of links for the ticket:

- Massive Deployment of Xen enabled Virtual Cluster using Diskless Remote ...
- Test Environment
- Pre-installation
- Install and Configure DRBL
- Configure PXELinux for Xen Network Booting
- Enable Xend in default rcX.d
- Know Issue of DRBL () with Xen
- Create Xen DomU configure files
- Xen DomU PXE Booting
  - Case 1: HVM supported
  - Case 2: only Para-Virtualization supported
- Reference

At the bottom of the screenshot is a window titled "DRBL\_Xen" showing details for a virtual machine. The state is "Powered off". The guest OS is "Other Linux". The configuration file is "C:\Virtual Machines\DRBL\_Xen\drbl\_xen.vmx". The version is "Current virtual machine for VMware Server 1.0.3". Below this are sections for "Commands" and "Devices". The "Commands" section has "Start this virtual machine" and "Edit virtual machine settings". The "Devices" section shows "Memory" at 256 MB and "Hard Disk (SCSI 0:0)".

# Building PaaS using DRBL-Hadoop

## Application

Social Computing, Enterprise, ISV, ...

eyeOS, Nutch, ICAS,  
X-RIME, ...

## Programming

Web 2.0, Mashups, Workflows, ...

Hadoop (MapReduce),  
Sector/Sphere, AppScale

## Management

Qos Negotiation, Admission Control,  
Pricing, SLA Management, Metering...

OpenNebula, Enomaly,  
Eucalyptus, OpenQRM, ...

## Virtualization

VM, VM management and Deployment

Xen, KVM, VirtualBox,  
QEMU, OpenVZ, ...

## Physical Hardware

Infrastructure: Computer, Storage, Network

# *Hadoop*

- <http://hadoop.apache.org>
- Hadoop is Apache Top Level Project
- Major sponsor is Yahoo!
- Developed by Doug Cutting
- Written by Java, it provides HDFS and MapReduce API
- Used in Yahoo since year 2006
- It had been deploy to 4000+ nodes in Yahoo
- Design to process dataset in Petabyte
- Facebook 、 Last.fm 、 Joost are also powered by Hadoop





- <http://sector.sourceforge.net/>
- Developed by National Center for Data Mining, USA
- Written by C/C++, so performance is better than Hadoop
- Provide file system **similar to** Google File System and MapReduce API
- Based on [UDT](#) which enhance the network performance
- [Open Cloud Consortium](#) provide [Open Cloud Testbed](#) and develop [MalStone toolkit for benchmark](#)



# Building PaaS using DRBL-Hadoop

- Used in <http://hadoop.nchc.org.tw>
- drbl-hadoop – mount local disk for HDFS and MapReduce  
svn co <http://trac.nchc.org.tw/pub/grid/drbl-hadoop>
- hadoop-register – web interface with ssh applet  
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 ▲

↑ ../

▶ 📁 etc

📄 adduser.php

Size

1.3 kB

Rev

103

85

Age

4 weeks

6 weeks

Last

wa

wa

# Demo : [hadoop.nchc.org.tw](http://hadoop.nchc.org.tw) for multi-users

- **DRBL Server x 1 (hadoop)**
- **DRBL Client x 19 (hadoop101~hadoop119)**
- **Based on Cloudera Debian package and enhance security setting and permission for multi-users.**

The image shows two overlapping windows from a Mozilla Firefox browser. The left window displays a terminal session with the following text:

```
Linux hadoop 2.6.26-2-amd64 #1 SMP Fri Mar 27 04:02:59 UTC 2009 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: Thu Jul 16 09:05:28 2009 from wr189-050.nchc.org.tw
hadoop004@hadoop:~$
```

The right window displays the 'hadoop Hadoop Map/Reduce Administration' page. The page title is 'hadoop Hadoop Map/Reduce Administration'. The state is 'RUNNING'. The started time is 'Sun Jul 19 22:48:19 EDT 2009'. The version is '0.18.3-4cloudera0.3.0, r'. The compiled time is 'Fri May 29 23:29:49 UTC 2009 by root'. The identifier is '200907192248'.

Cluster Summary

Maps	Reduces	Total Submissions	Nodes	Map Task Capacity	Reduce Task
0	0	711	19	38	38

Running Jobs

Running Jobs
none

# Building SaaS using DRBL-biocluster

## 應用軟體 Application

Social Computing, Enterprise, ISV, ...

eyeOS, Nutch, ICAS,  
X-RIME, ...

## 程式語言 Programming

Web 2.0 介面, Mashups, Workflows, ...

Hadoop (MapReduce),  
Sector/Sphere, AppScale

## 控制管理 Control

Qos Negotiation, Admission Control,  
Pricing, SLA Management, Metering...

OpenNebula, Enomaly,  
Eucalyptus, OpenQRM, ...

## 虛擬化 Virtualization

VM, VM management and Deployment

Xen, KVM, VirtualBox,  
QEMU, OpenVZ, ...

## 硬體設施 Hardware

Infrastructure: Computer, Storage, Network

# Building SaaS using DRBL-biocluster

- **Need more time to package related software.**
- **drbl-biocluster – batch script of Debian to install bioinformatics related softwares**

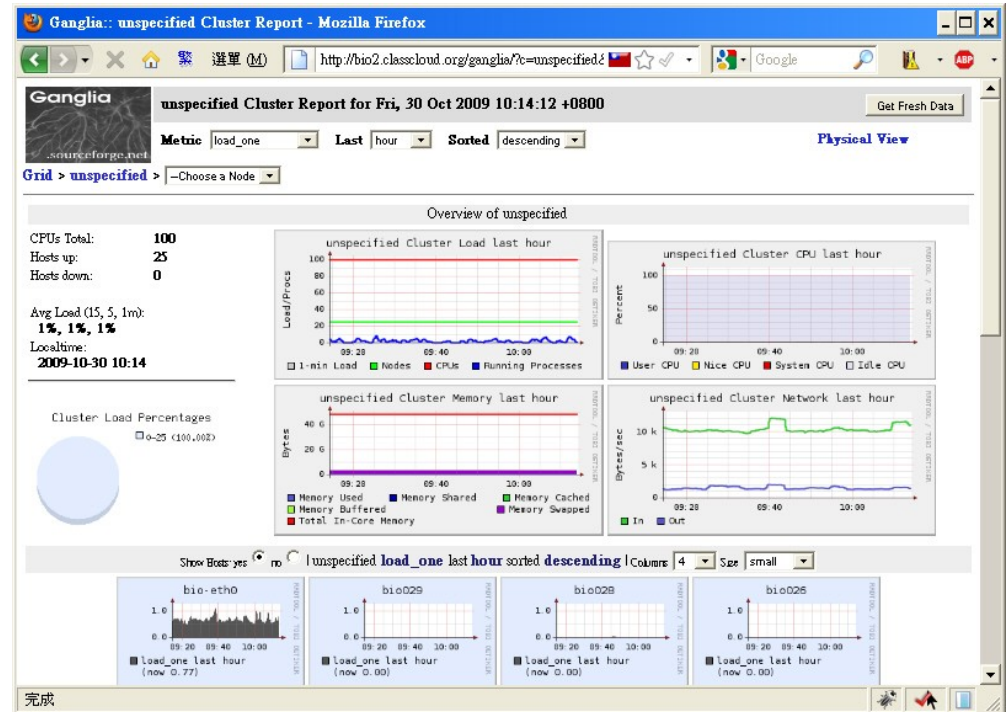
**svn co <http://trac.nchc.org.tw/pub/grid/drbl-biocluster>**

- **Including DRBL、MPICH2、R、Rmpi、BioConductor、Ganglia、Nagios、AutoFACT、BLAST、SIM4、Clustal、PipMaker、Phylip、Eland、Velvet、Bowtie、SOAP**

```
Linux hadoop 2.6.26-2-amd64 #1 SMP Fri Mar 27 04:02:59 UTC 2009 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: Thu Jul 16 09:05:28 2009 from wr189-050.nchc.org.tw
hadoop004@hadoop:~$
```



# 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**