

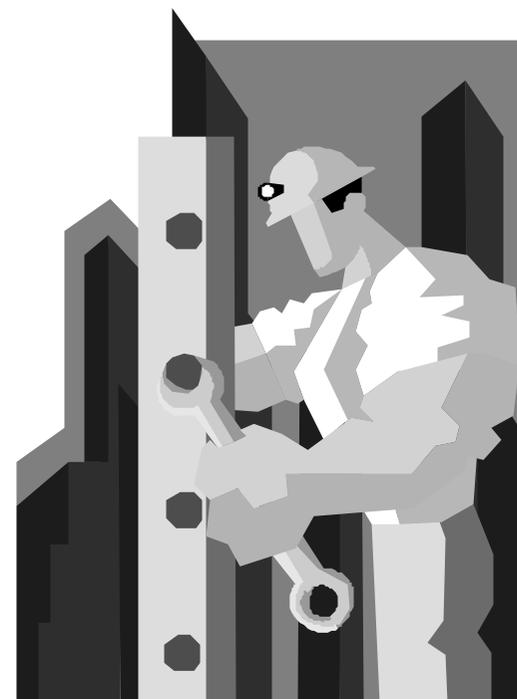
Linux

가

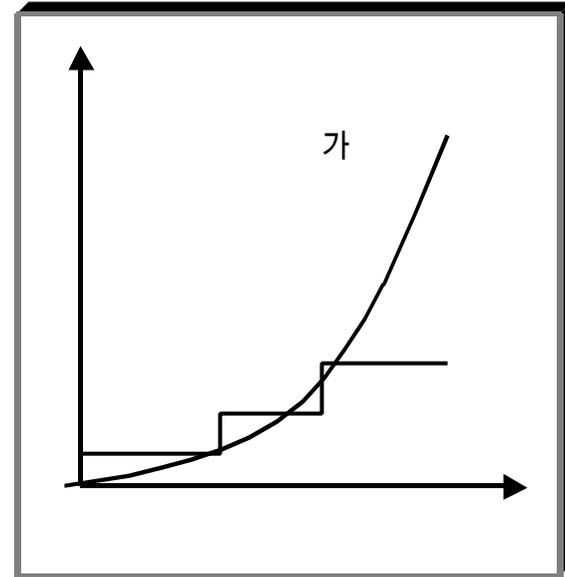
- Linux Revolution 2000
- 2000. 12. 7
-
- loveme@wowlinux.com

Contents

-
-
- Availability & Scalability
- Linux Virtual Server
- Technology
 - Heartbeat
 - IP Address takeover
 - ARP
 - Gratuitous ARP
 - Ldirectord
 - ipvsadm
- ULTRA MONKEY LVS
 - Topologies
 - Preparing Cluster nodes
 - Real Server
 - Director
 - Heartbeat & LVS
 - (1) - real server /
 - (2) - master Director /
-
-



■ (Traffic)
100% 가
■ 가 (Traffic)
가 가
■ 가 가





■ CPU , Memory , HDD



node 가



가



2

2



가



가

가

,

Availability & Scalability

■ Availability

■ 가

가

가

가

(Availability)가

.

■ Scalability

■

가

100

가

1000

가

scalability가

,

.

가

가		
90%	1	1
99%	4	2
99.9%	9	3
99.99%	1	4
99.999%	5	5
99.9999%	30	6
99.99999%	3	7

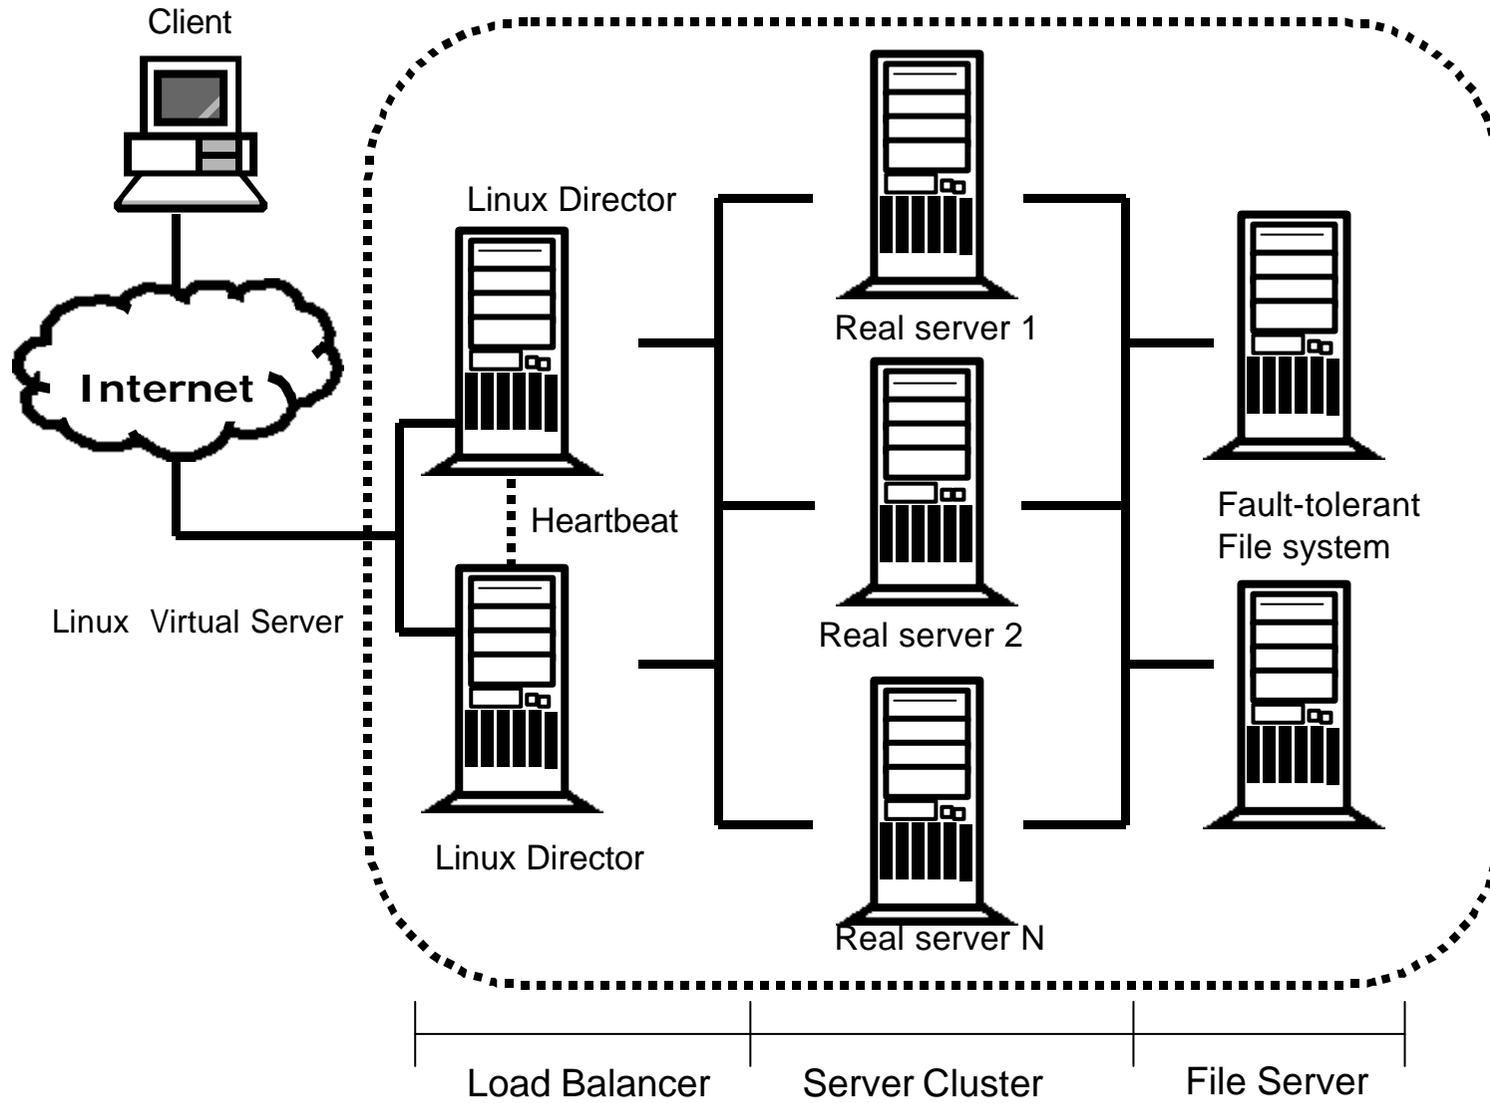
가

■ 가 5 가

LVS – Linux Virtual Server

- : , , 가
- 가
가
- 가
- 가
- LVS (front-end)
■
■ IP 가
■ 가
■ 가
가

LVS -



LVS Load Balancing method

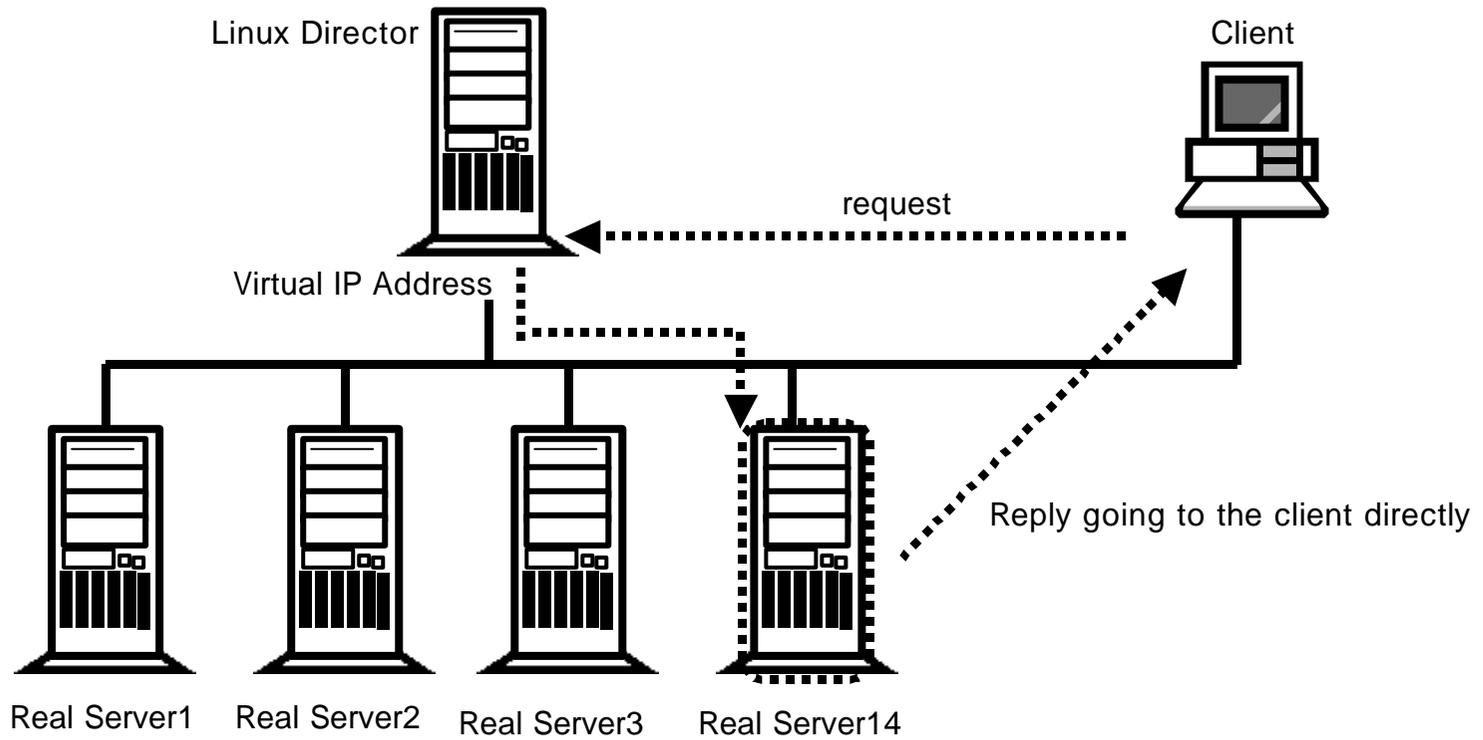
- NAT 가
- IP 가
- 가
 - 가 가
 -
 - 가
 - IP 가

LVS-DR

-
-
-
-
-

가 IP 가

가



Direct Routing

- Director 가
 - Director 가
 - Director 가
 - Director 가
 - Director 가
 - 가 -> DRLB 가
- 가

가

■ Round-Robin Scheduling

-
-

■ Weighted Round-Robin Scheduling

■ 가

- (가) 가 . 가 1 .
▪) 가 A,B,C 가 가 4,3,2
ABCABCABA

■ 가 가 가

- - 가 가
 - 가 가
 - 가 가

가

■ Least - Connection Scheduling

-
- 가
-
- 가 .
- 가
- TCP TIME_WAIT
-

■ Weighted Least - Connection Scheduling

- 가
-
- 가 가
- 가 . 가
- 1 .
- 가 가 가

Technology

- Heartbeat
- IP Address takeover
- ARP
- Gratuitous ARP
- Ldirectord
- ipvsadm

Heartbeat

- heartbeat

-

-

- Heartbeat

IP

가

IP

- Master 가

- Master

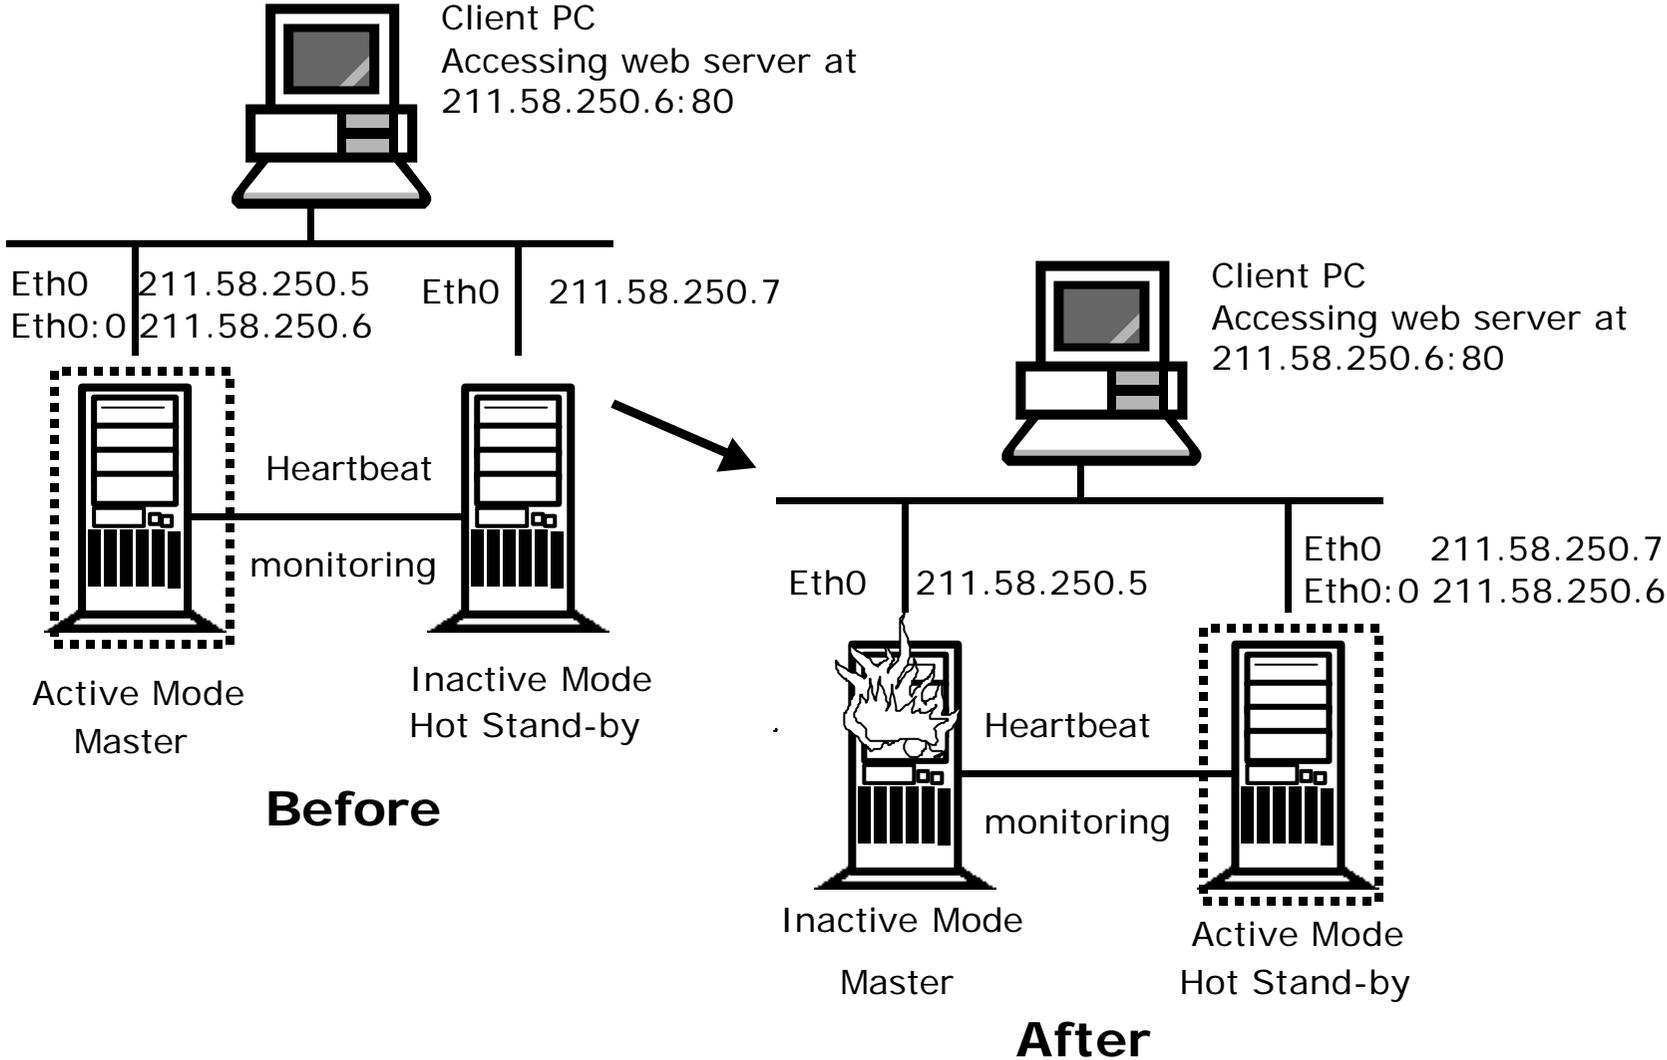
Standby

가 가

IP

IP Address Takeover

IP Address takeover

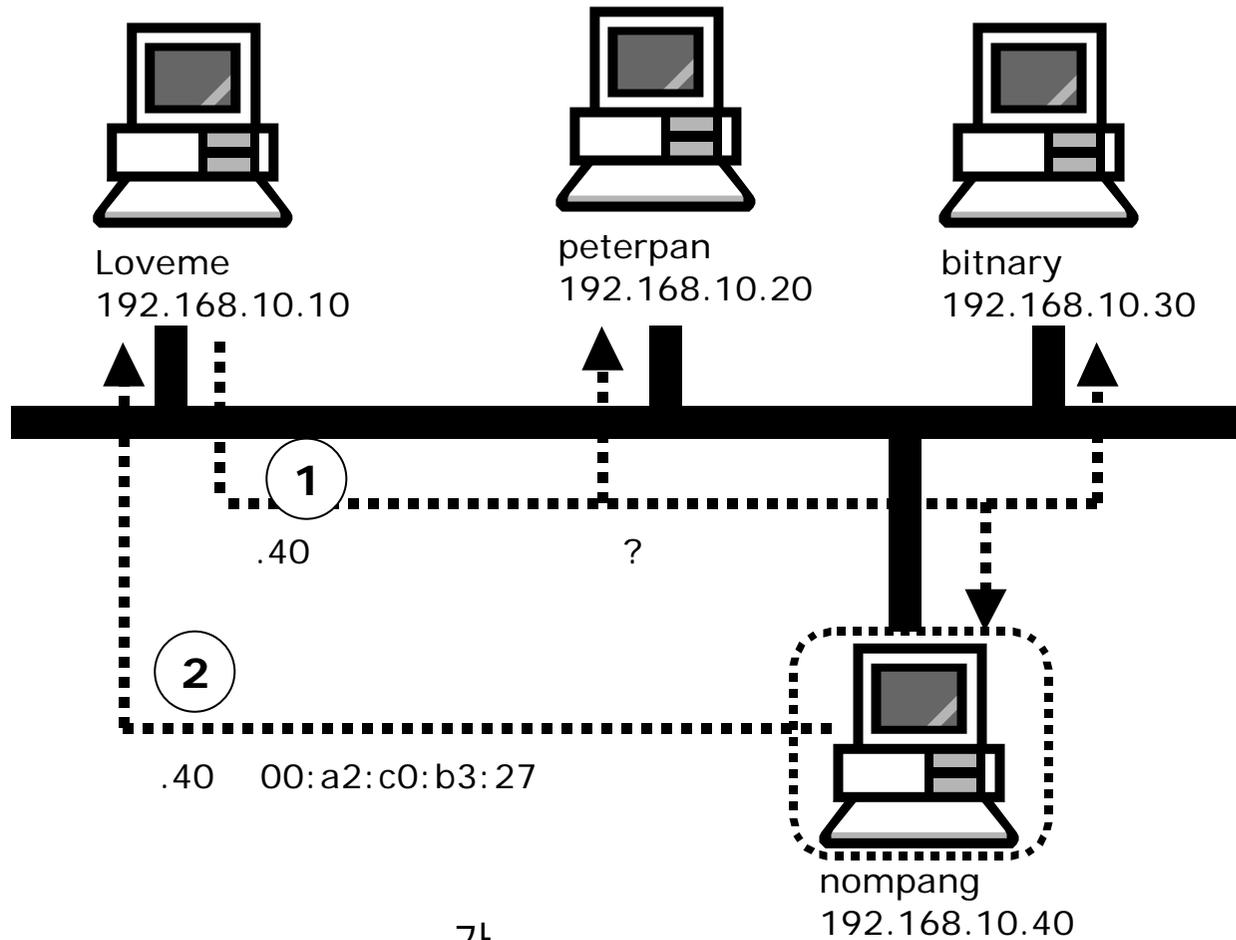


ARP

- IP (MAC)
- IP
- ARP (Address Resolution Protocol:)
- Data Link Layer
 - OSI Reference Model *
- 가 IP (MAC)
 - IP
 - ARP

ARP

■ ARP



- ARP
- IP

가

- IP

ARP

Gratuitous ARP

- Gratuitous ARP

- ARP (MAC) IP

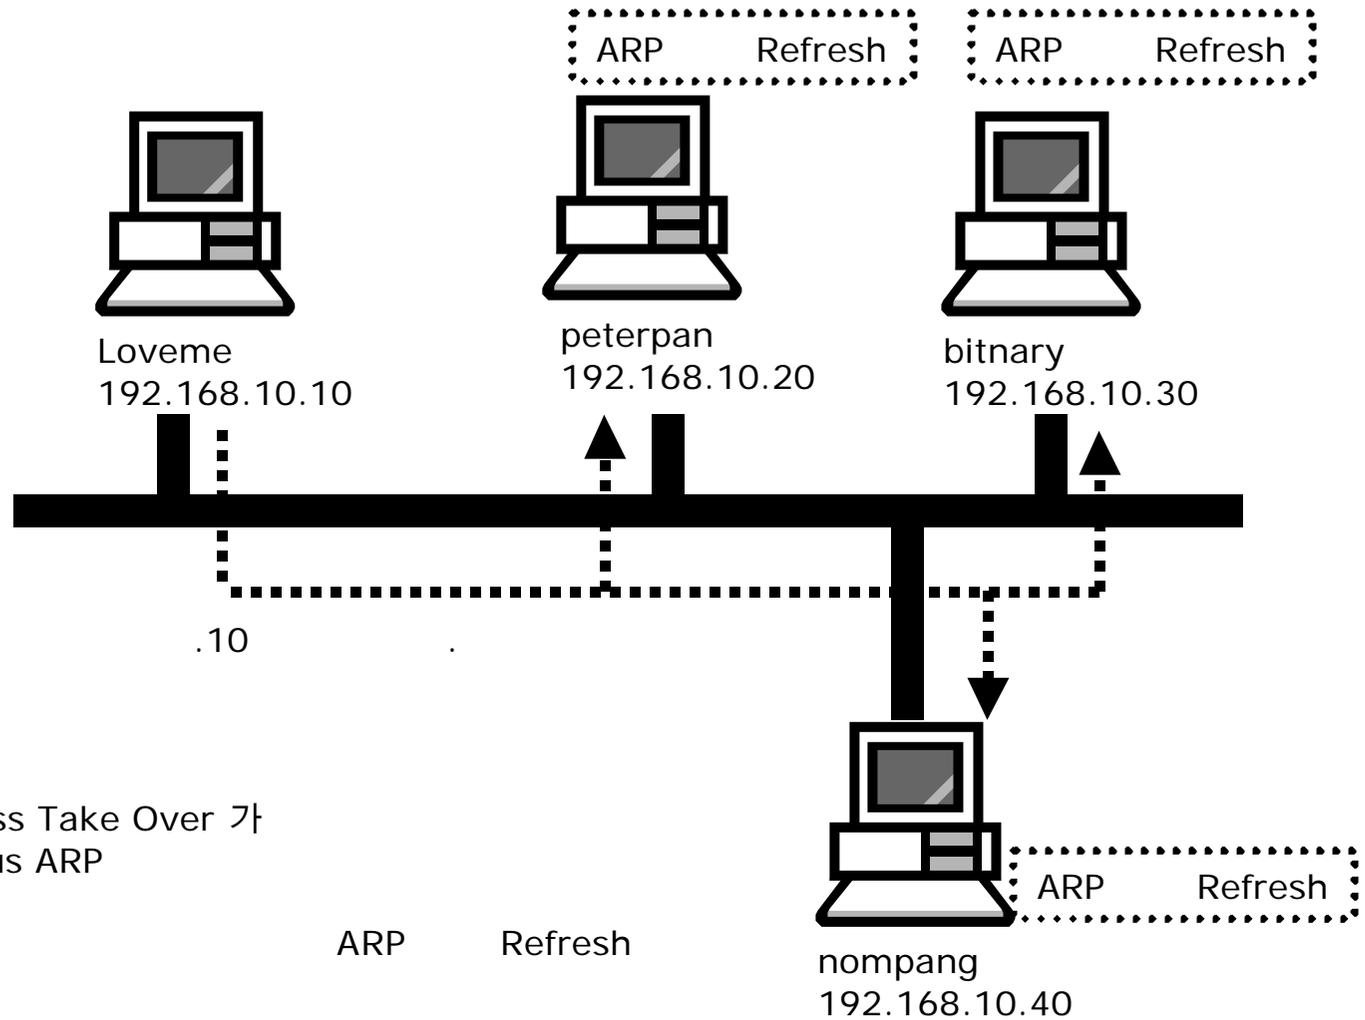
- Gratuitous ARP

ARP Refresh ()

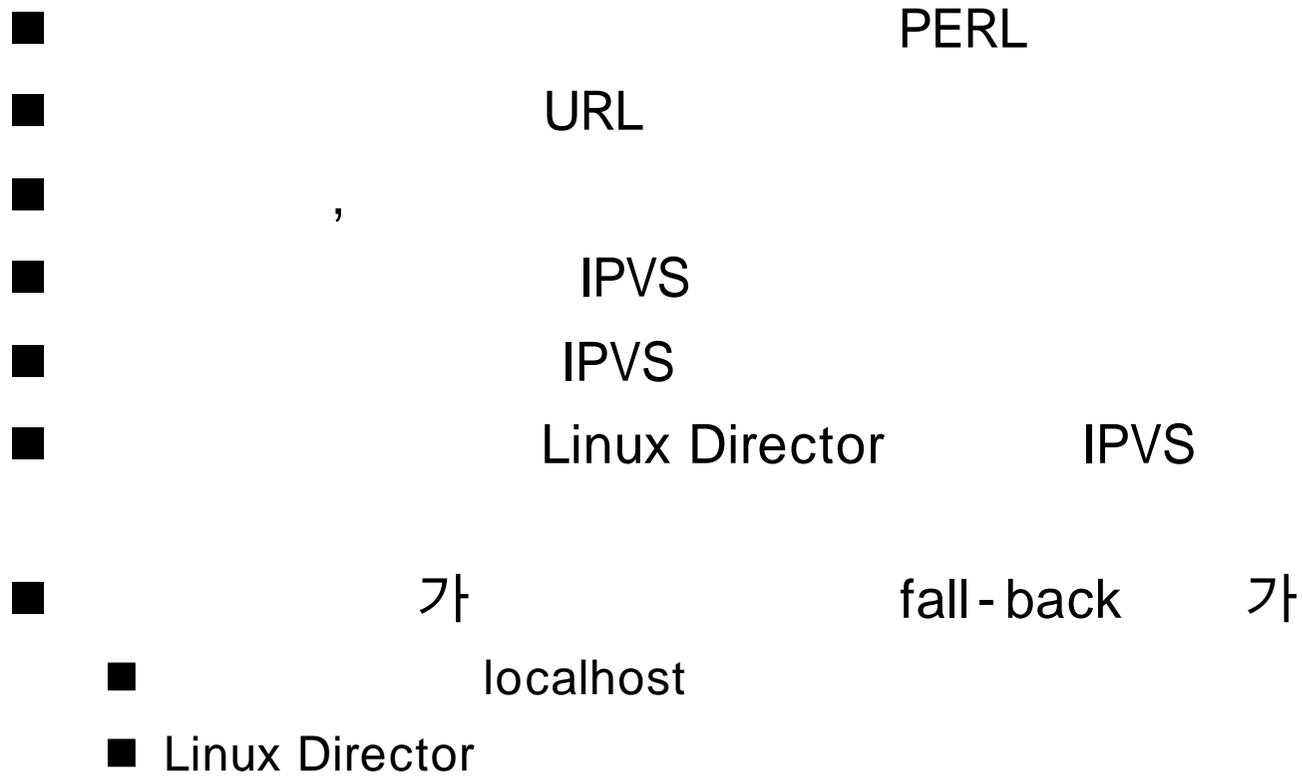
- host
120

- - IP Address Take Over ARP
NIC 가 IP 가 NIC (MAC) 가 Refresh

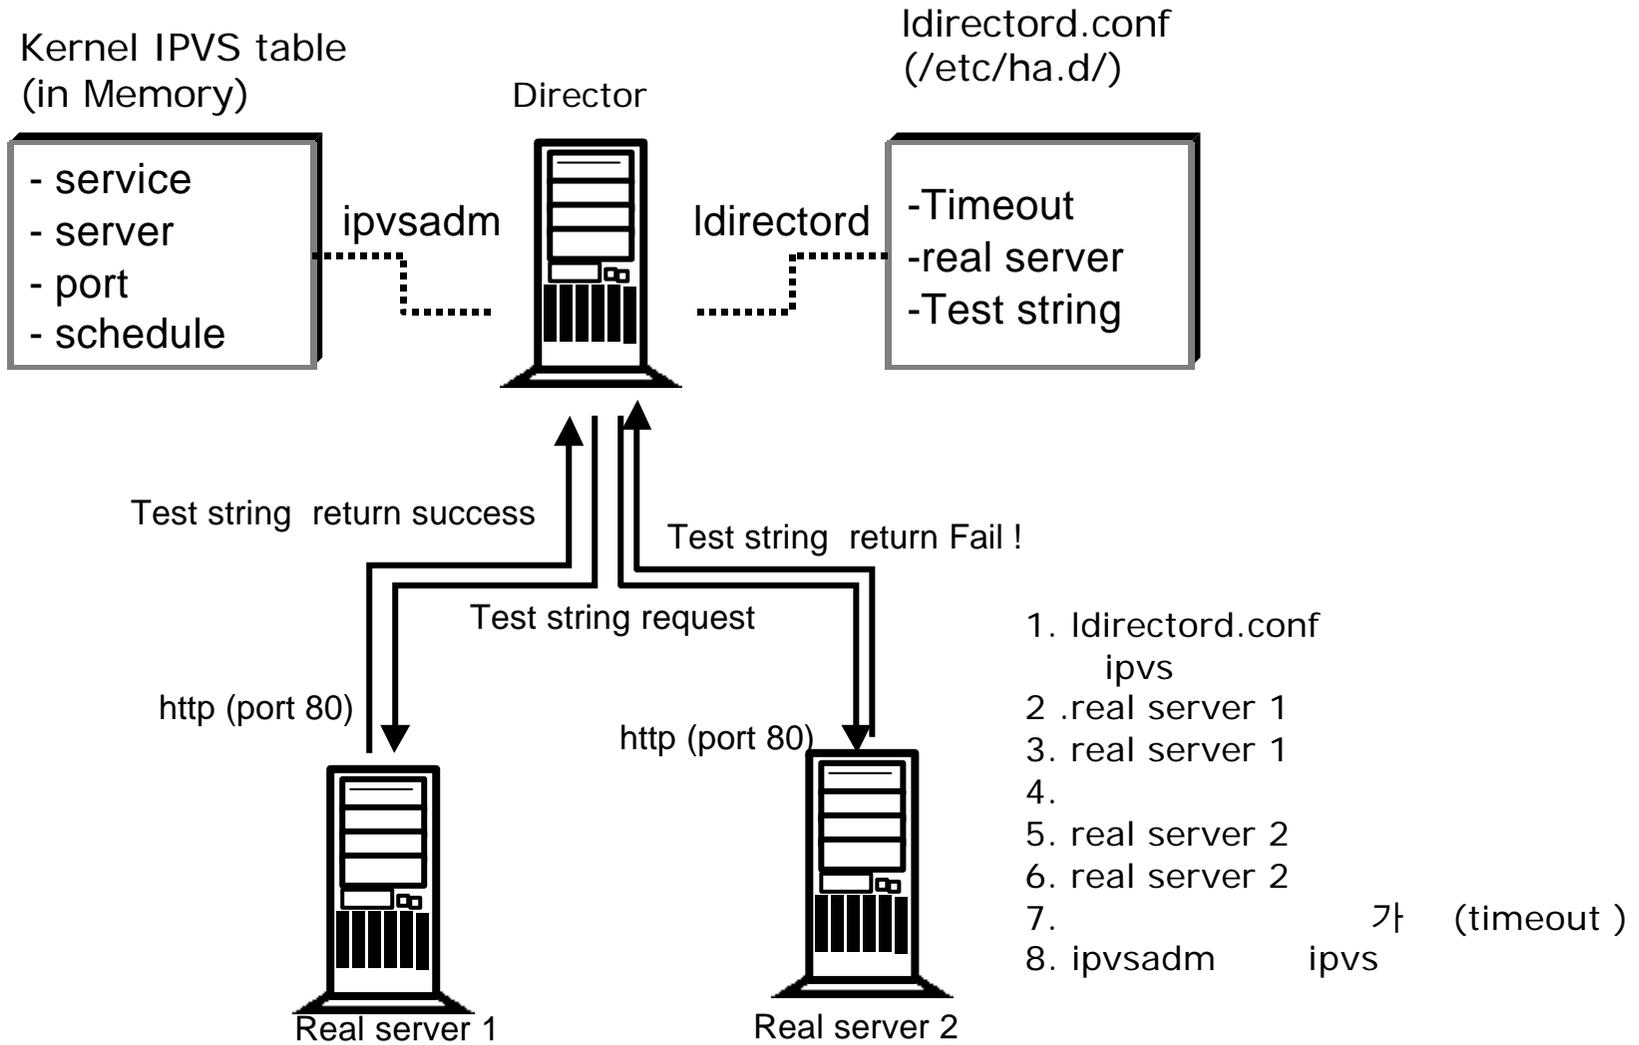
Gratuitous ARP



Ldirectord



Ldirectord



ipvsadm

- 2 가 IPVS
- IPVS (Director)
- 2
 - TCP UDP
- 3가
 - NAT, tunneling, direct routing
- 4가
 - round robin, weighted round robin,
 - least-connection weighted least-connection

ULTRAMONKEY

LVS

- Topologies
- Preparing Cluster nodes
- Real Server
- Director
- Heartbeat & LVS
- (1) - real server /
- (2) - master Director /

ULTRA MONKEY

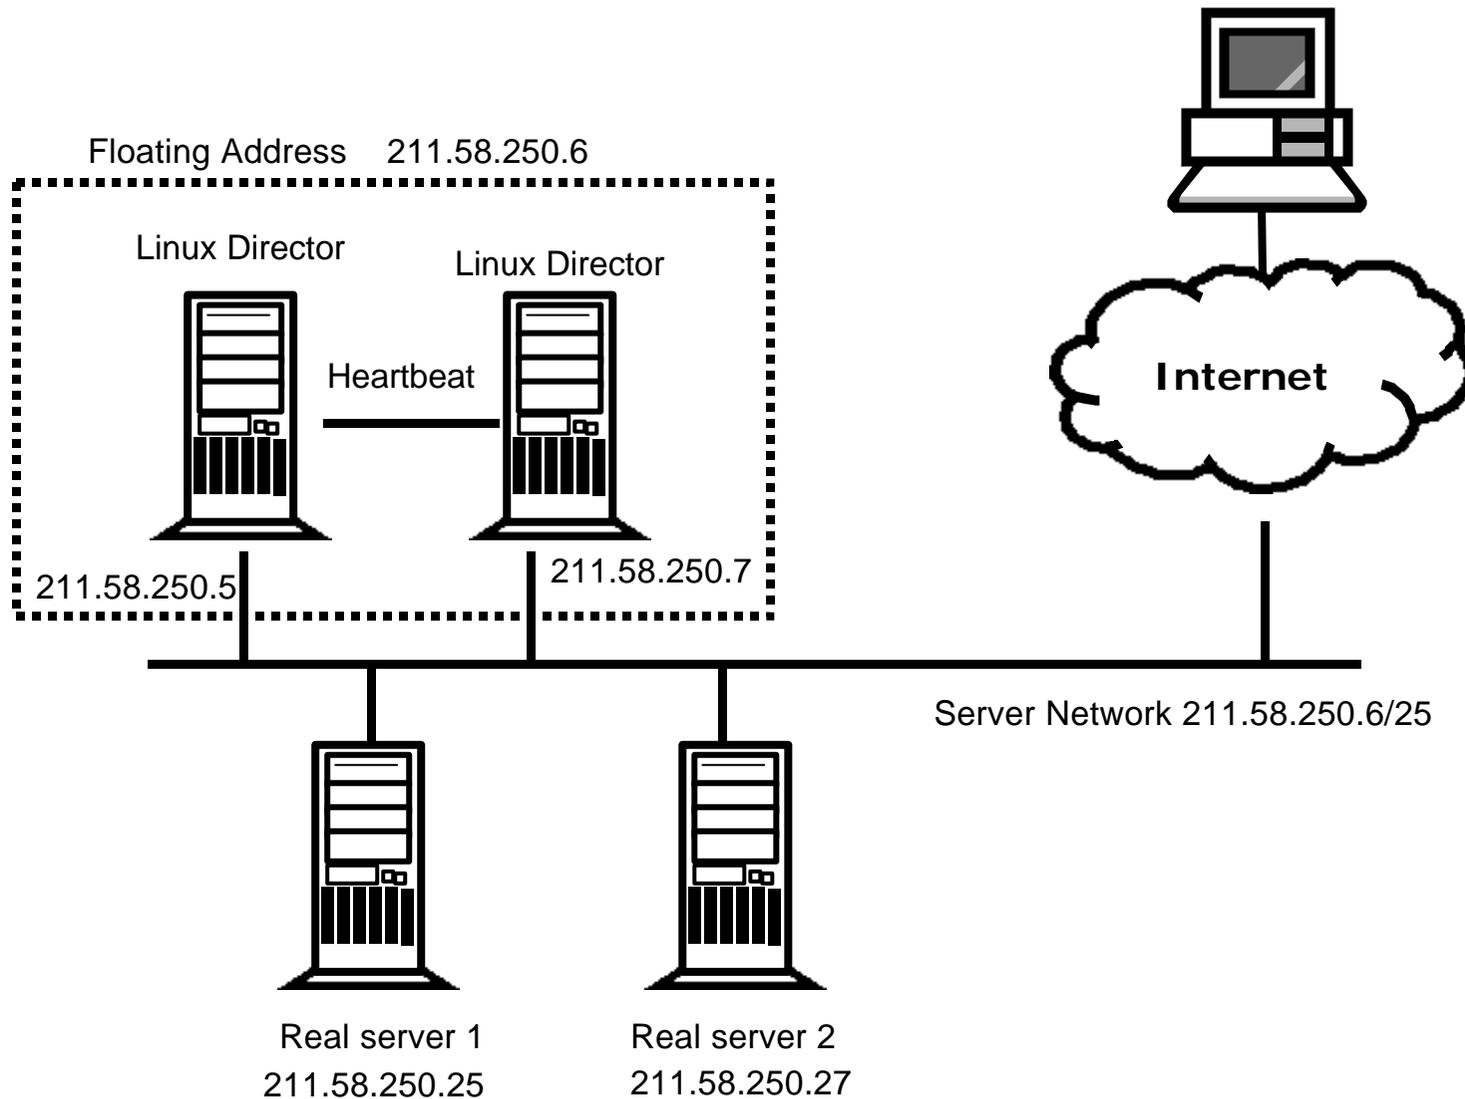
- High Availability and Load Balancing Solution



- Key Features

- Layer 4 Switching using The Linux Virtual Server
 - Easily expandable to a large number of IP based virtual services
 - High Availability provided by Heartbeat protocol
 - Service level monitoring using Ldirectord
 - All Code is Open Source
- Easily expandable to other services such as email and FTP.

Topologies — High Availability Load Balancing Service



Topologies – overview

- 2 Director Heartbeat
 - Directors (Master)
 - Director (hot stand - by)

- DNS 가 IP가
- director 가 IP (211.58.250.6)
- 80 port Apache HTTP
- LVS/DR
-

Preparing Cluster nodes

■ HW – Director & Web node

■ Compaq SB

- Intel Celeron 500MHz, Intel 810 Chipset ,64MB RAM, 8.4GB HDD , RealTeck 8139c NIC
- 48x CD-ROM, Integrated Graphic Controller(4MB)

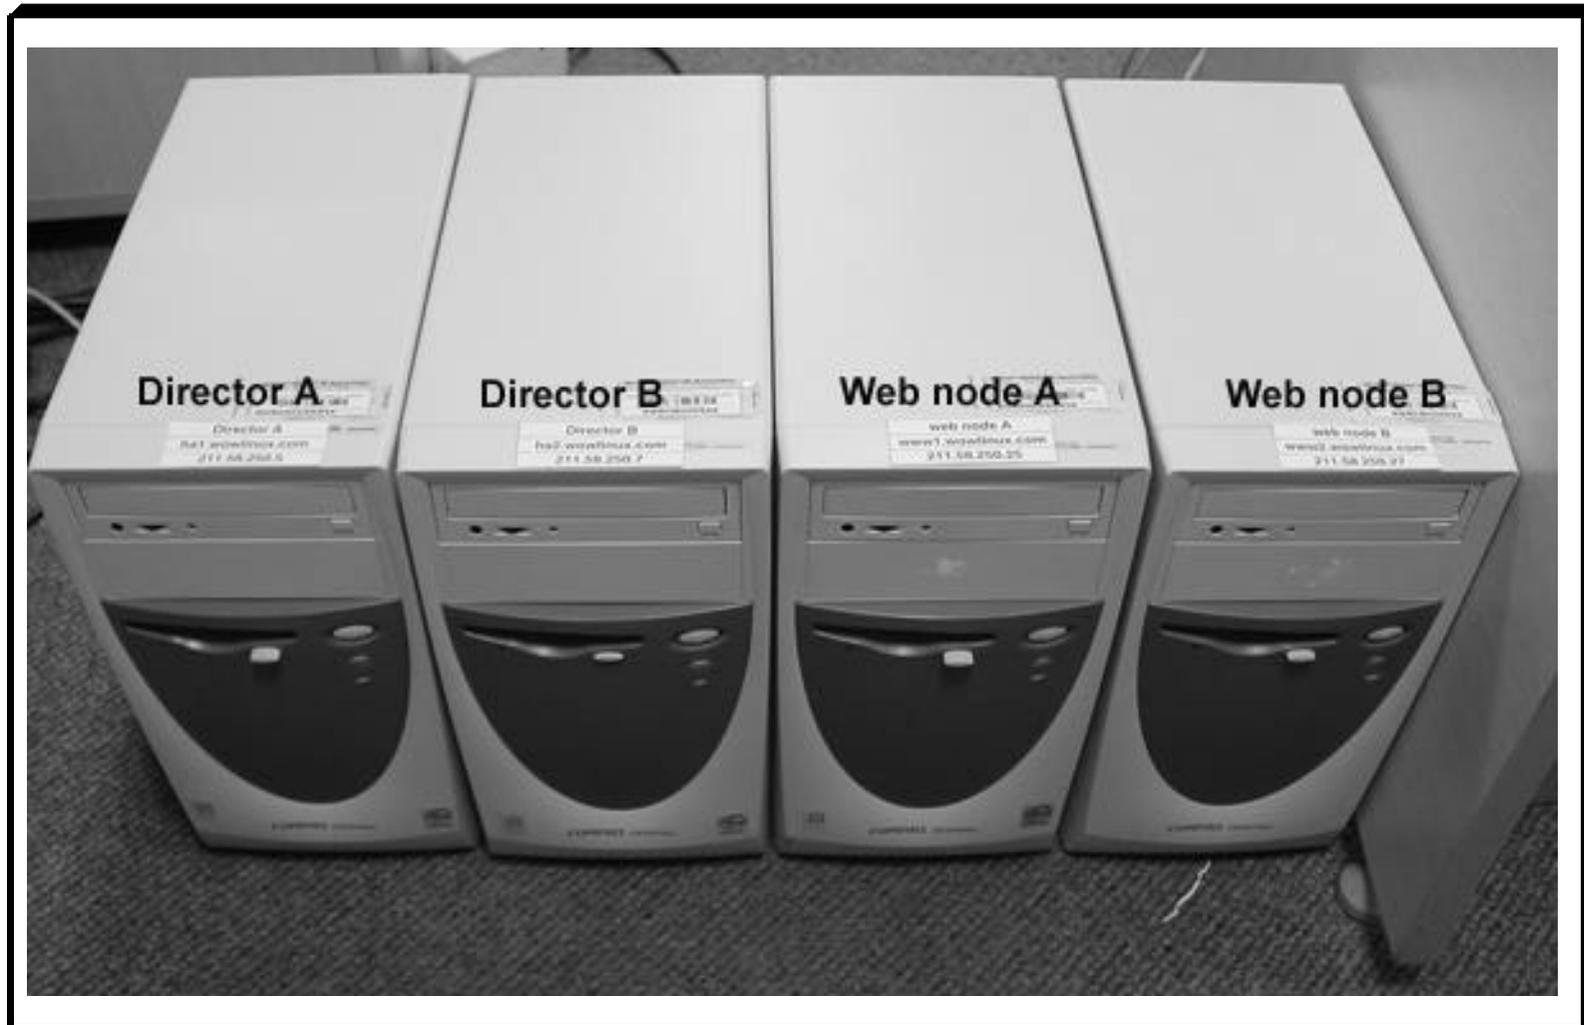
■ OS

- Wowlinux 6.2 update (kernel 2.2.16) Full Install

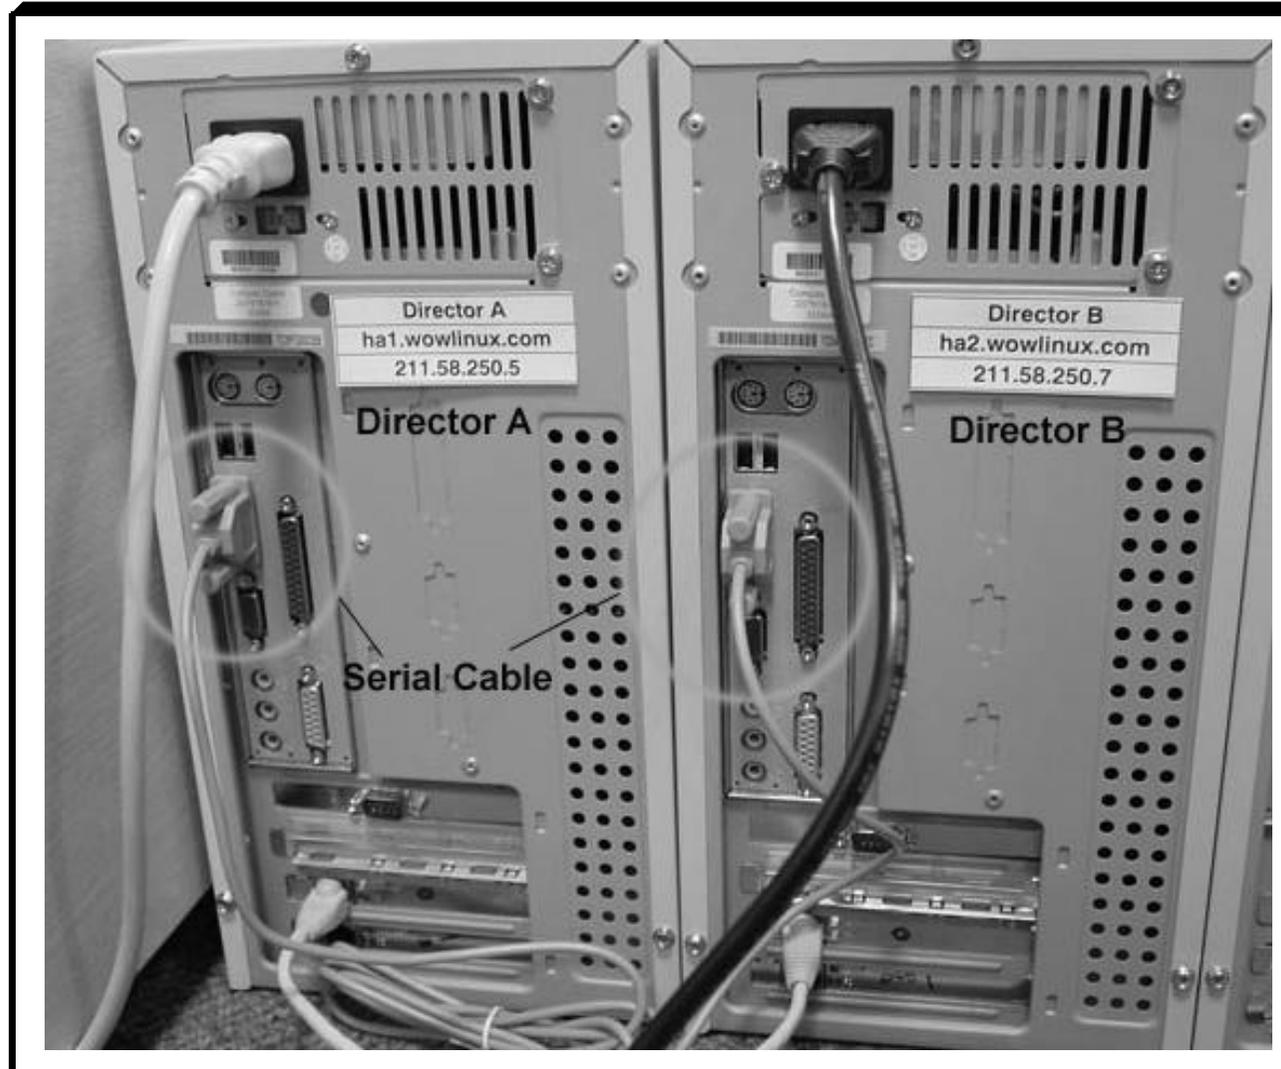
■ Network

- Ha.wowlinux.com : 211.58.250.6 (Virtual IP)
- Ha1.wowlinux.com : 211.58.250.5 (Director 1)
- Ha2.wowlinux.com : 211.58.250.7 (Director 2)
- www1.wowlinux.com : 211.58.250.25 (web node 1)
- www2.wowlinux.com : 211.58.250.27 (web node 2)
- Ha1 & ha2

Preparing Cluster nodes



Preparing Cluster nodes



Network -

■ /etc/hosts

127.0.0.1	localhost.localdomain	localhost
211.58.250.6	ha.wowlinux.com	ha
211.58.250.5	ha1.wowlinux.com	ha1
211.58.250.7	ha2.wowlinux.com	ha2
211.58.250.25	www1.wowlinux.com	www1
211.58.250.27	www2.wowlinux.com	www2

Real Server (1)

- 가 IP (211.58.250.6)
 - loopback device IP Alias
 - /etc/sysconfig/network-scripts/ifcfg-lo:0

```
DEVICE=lo:0
IPADDR=211.58.250.6
NETMASK=255.255.255.255
NETWORK=211.58.250.0
BROADCAST=211.58.250.127
ONBOOT=yes
NAME=loopback
```

- - /etc/sysconfig/network
 - GATEWAYDEV=eth0

eth0

Real Server (2)

- IP alias

- #/sbin/ifup lo

- /sbin/ifconfig 가 .

```
lo:0 Link encap:Local Loopback
      inet addr:211.58.250.6 Mask:255.255.255.255
      UP LOOPBACK RUNNING MTU:3924 Metric:1
```

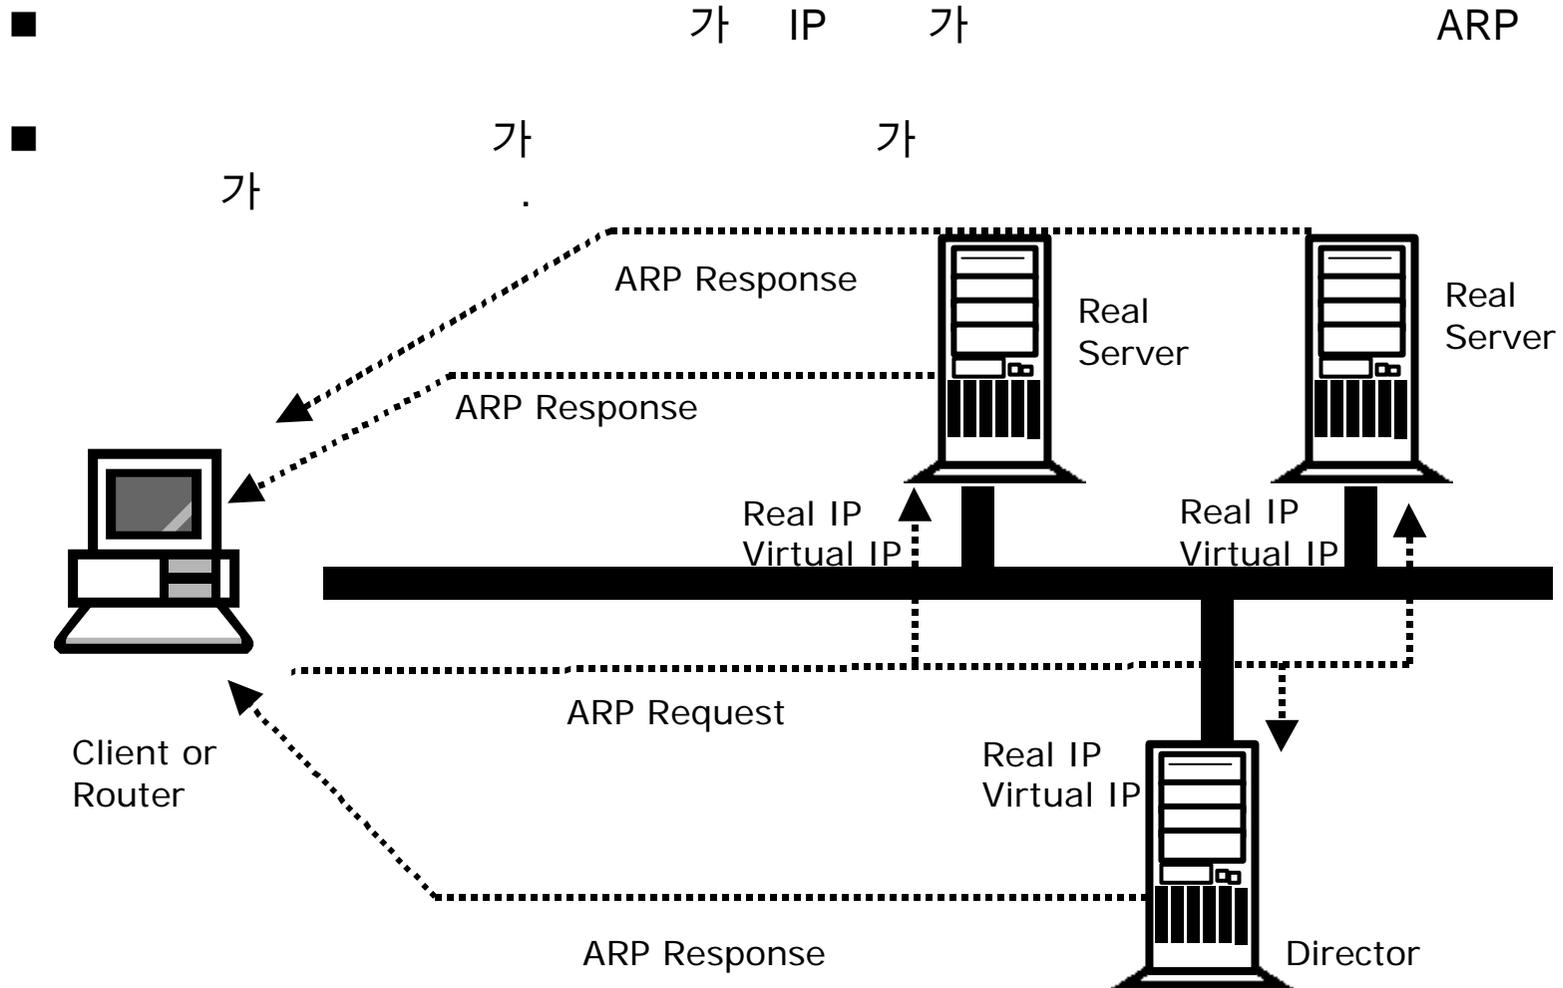
- : Loopback 가
211.58.250.6
255.255.255.255

- web node ntsysv

- /etc/ha.d/conf/ldirectord.cf Index.html Test

- Test Page of www1.wowlinux.com
 - Test Page of www2.wowlinux.com

Real Server (3) - ARP



Real Server (3) ARP

- 2.2.14 Hidden Device (ARP Request
)
가
■ <http://www.linuxvirtualserver.org/arp.html>

- Real server loopback 가 ARP
hidden
■ /etc/sysctl.conf 가.

```
# Enable configuration of hidden devices  
net.ipv4.conf.all.hidden = 1  
# Make the loopback device hidden  
net.ipv4.conf.lo.hidden = 1
```

- ■ /sbin/sysctl -p

Serial cable / (1)

- 2 LVS

- ha2

```
[root@ha2 /root]# cat < /dev/ttyS0
```

- ha1

```
[root@ha1 /]# echo HELLO > /dev/ttyS0  
[root@ha1 /]# echo 'THIS IS HEARTBEAT TEST from  
HA1' > /dev/ttyS0
```

- ha2

가

```
[root@ha2 /root]# cat < /dev/ttyS0  
HELLO  
  
THIS IS HEARTBEAT TEST from HA1
```

Serial cable / (2)



```
#cat /proc/tty/driver/serial
serinfo:1.0 driver:4.27
0: uart:16550A port:3F8 irq:4 baud:9600 tx:18 rx:0 CTS|DSR|CD
1: uart:16550A port:2F8 irq:3 baud:9600 tx:11 rx:0
2: uart:unknown port:3E8 irq:4
...
```



CD Carrier Detect

.

Director (1) - Install

- - 2.2.14 .
 - kernel 2.2.16 IPVS 가
 - 6.2up

- Ultramonkey RPM Package
 - <ftp://ultramonkey.sourceforge.net/pub/ultramonkey/ultramonkey-1.0.1/RPMS/>

- Piranha
 - Redhat piranha Ultramonkey

 - Piranha Clustering .
 - /mnt/cdrom/RedHat/base/comps
 - piranha , piranha-docs , pvm , lam, piranha-gui, pvm-gui

Director (2) - RPM Package install

- HA1 11 가
 - perl-libnet - 1.0607 - 1.i386.rpm
 - perl-HTML-Parser - 3.07 - 1.i386.rpm
 - MIME-Base64 - 2.11 - 2.i386.rpm
 - URI - 1.06 - 1.i386.rpm
 - openssl - 0.9.5a - 2.i386.rpm
 - perl-Net-SSLeay - 1.05 - 4.i386.rpm
 - LWP.pm - 5.48 - 2.i386.rpm
 - heartbeat - 0.4.7bpre1 - 1.um.1.i386.rpm
 - initscripts - 5.00 - 1.1.um.1.i386.rpm
 - ipvsadm - 1.9 - 1.um.1.i386.rpm
 - ultramonkey-doc - 1.0.1 - 1.noarch.rpm

- 가

- Rebooting !

Director (3) -

■ Director

■ IPVS4

IPVS
가

■ /etc/sysctl.conf

가

```
# Disables packet forwarding
net.ipv4.ip_forward = 1
# Enables source route verification
net.ipv4.conf.all.rp_filter = 1
# Disables automatic defragmentation
net.ipv4.ip_always_defrag = 0
# Disables the magic-sysrq key
kernel.sysrq = 0
```

Director (3) -

- net.ipv4.ip_forward = 1

-

```
[root@ha1 conf]# /sbin/sysctl -p
net.ipv4.ip_forward = 1
net.ipv4.conf.all.rp_filter = 1
net.ipv4.ip_always_defrag = 0
kernel.sysrq = 0
```

Director (4) - ldirectord

- /etc/ha.d/conf/ldirectord.cf
- /usr/doc/ultramonkey-doc-1.0.1/topologies/config/
- /etc/ha.d/conf/ldirectord.cf

```
timeout=3  
checkinterval=1  
fallback=127.0.0.1:80  
virtual=211.58.250.6:80  
    real=211.58.250.25:80 gate  
    real=211.58.250.27:80 gate  
service=http  
request="index.html"  
receive="Test Page"  
scheduler=rr  
protocol=tcp
```

Director (4) - Idirectord

- checkinterval :
- fallback :
- virtual : 가
- real : IP ,
 - gate = Direct Routing
 - masq = Network address translation
 - ipip = IP Tunneling
- service :
- request :
- receive :
- scheduler :
 - rr = Round-Robin
 - wrr = Weighted Round-Robin
 - lc = Least-Connection
 - wlc = Weighted Least-Connection

Director (5) - Heartbeat

- LVS /etc/ha.d/ 3가
 - authkeys : LVS
 - ha.cf : heartbeat .
 - haresources : heartbeat
 - <http://www.linux-ha.org/download/GettingStarted.html>

■ authkeys

```
auth 1
1 crc
#2 sha1 ultramonkey
#3 md5 Hello!
```

- crc .
-
- 600

Director (5) - Heartbeat

■ ha.cf

```
debugfile /var/log/ha-debug
logfile /var/log/ha-log
keepalive 3
deadtime 9
serial /dev/ttyS0
baud 19200
node ha1.wowlinux.com
node ha2.wowlinux.com
```

- keepalive : Heartbeat
- deadtime :
- serial :
- baud :
- node : Heartbeat
 - #uname -n

Director (5) - Heartbeat

■ Haresource

```
...  
ha1.wowlinux.com      211.58.250.6  ldirectrod  
...
```

■ ha1.wowlinux.com : master LVS node

■ 211.58.250.6 : LVS 가 IP

- heartbeat IP IP Alias
- master node 가 hot standby 가 IP
eth0:0 가 (IP Take over)

■ ldirectrod : heartbeat

- /etc/rc.d/init.d/

LVS , - HDD

- run level 1 HDD
- #init 1
- #cat /dev/hda > /dev/hdc
- 가 HDD

Heartbeat

- www1, www2

```
[root@www1 /root]# /etc/rc.d/init.d/httpd start  
Starting httpd: [ OK ]
```

- ha1 , ha2 heartbeat

```
[root@ha1 ha.d]# /etc/rc.d/init.d/heartbeat start  
Starting High-Availability services: [ OK ]
```



- ldirectord 가
- ntsysv ldirevtord
- heartbeat .

Heartbeat & LVS

■ Heartbeat

```
[root@ha1 /etc]# cat /var/log/ha-log  
  
heartbeat:11/17_13:01:55 INFO: Running /etc/ha.d/rc.d/status status  
heartbeat:11/17_13:02:03 INFO: Running /etc/ha.d/resource.d/IPaddr  
211.58.250.6 status  
heartbeat:11/17_13:02:03 INFO: Running /etc/ha.d/rc.d/ip-request ip-  
request  
heartbeat:11/17_13:02:13 Acquiring resource group: ha1.wowlinux.com  
211.58.250.6 ldirectord  
heartbeat:11/17_13:02:13 INFO: Running  
/etc/ha.d/resource.d/ldirectord start  
heartbeat:11/17_13:02:14 INFO: Running /etc/ha.d/resource.d/IPaddr  
211.58.250.6 start  
heartbeat:11/17_13:02:14 INFO: ifconfig eth0:0 211.58.250.6 netmask  
255.255.255.128 broadcast 211.58.250.127  
heartbeat:11/17_13:02:14 Sending Gratuitous Arp for 211.58.250.6 on  
eth0:0 [eth0]
```

Heartbeat & LVS

■ ldirectord

```
[root@ha1 /etc]# cat /var/log/ldirectord.log

[17 13:02:14 2000|ldirectord] Starting Linux Director Daemon
[17 13:02:14 2000|ldirectord] Adding virtual server:
211.58.250.6:80
[17 13:02:14 2000|ldirectord] Starting fallback server for:
211.58.250.6:80
[17 13:02:14 2000|ldirectord] Adding real server: 211.58.250.25:80
(1*211.58.250.6:80)
[17 13:02:14 2000|ldirectord] Turning off fallback server for:
211.58.250.6:80
[17 13:02:14 2000|ldirectord] Adding real server: 211.58.250.27:80
(2*211.58.250.6:80)
```

Heartbeat & LVS

■ ipvsadm

PIVS

```
[root@ha1 /etc]# ipvsadm
```

```
IP Virtual Server version 0.9.12 (size=4096)
```

```
Prot LocalAddress:Port Scheduler Flags
```

```
-> RemoteAddress:Port      Forward Weight ActiveConn
```

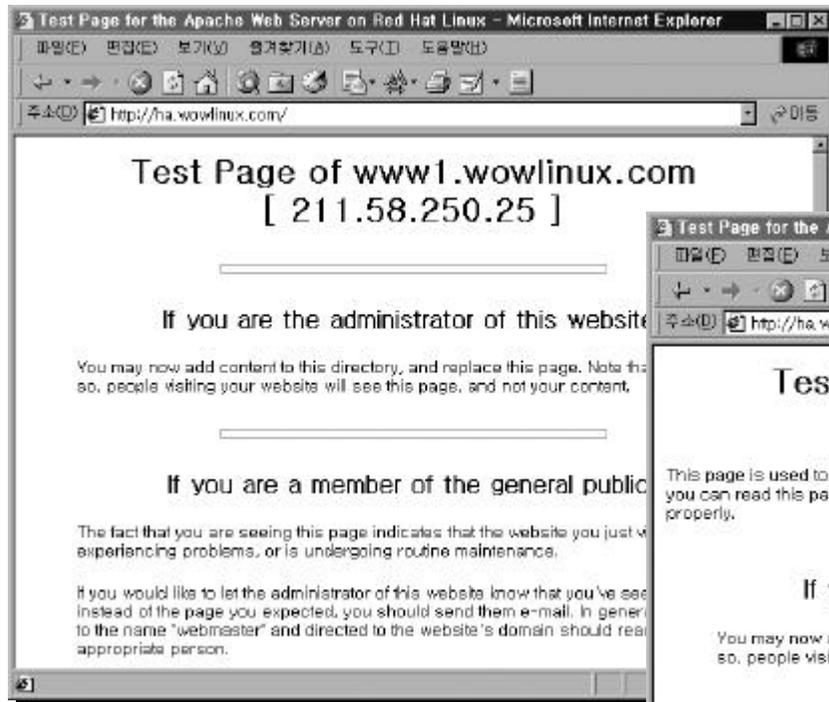
```
InActConn
```

```
TCP  ha.wowlinux.com:www rr
```

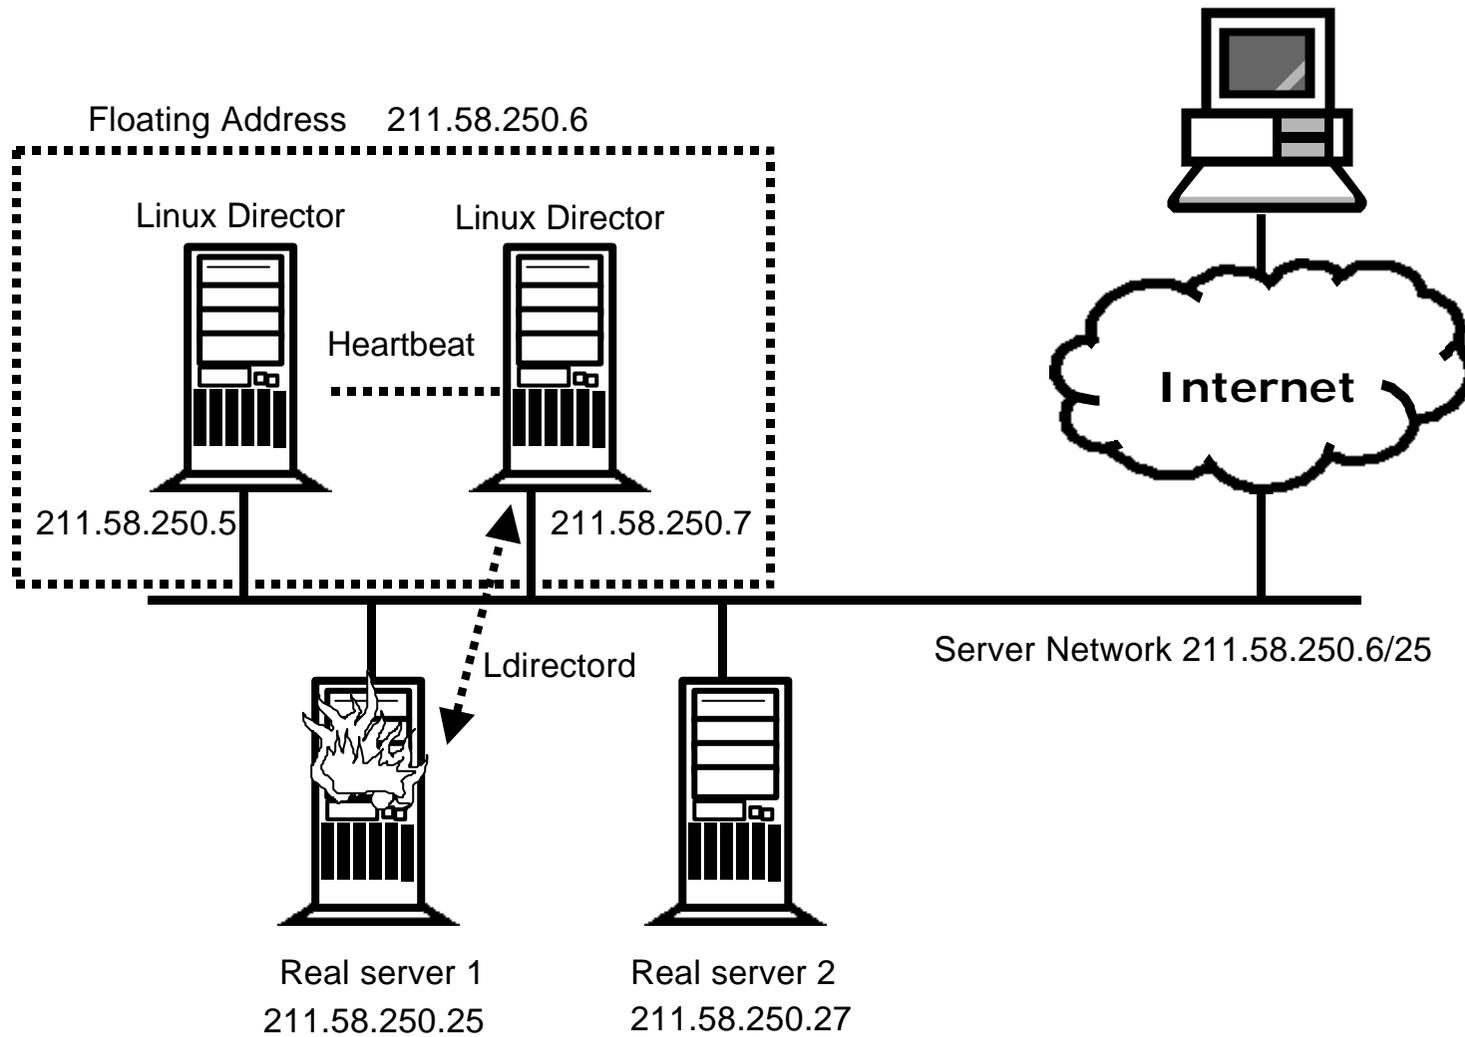
```
-> www2.wowlinux.com:www    Route  1    0    0
```

```
-> www1.wowlinux.com:www    Route  1    0    0
```

■ <http://ha.wowlinux.com>



(1) – real server /



(1) – real server

■ www1

```
[root@www1 /root]# /etc/rc.d/init.d/httpd stop
Stopping httpd: [ OK ]
```

■ Ha1 /var/log/ldirectord.log

```
...
[Fri Nov 17 13:24:28 2000|ldirectord.cf] Removing real
server: 211.58.250.25:80 (1*211.58.250.6:80)
```

■ Ha1 ipvsadm

```
[root@ha1 etc]# ipvsadm
IP Virtual Server version 0.9.12 (size=4096)
Prot LocalAddress:Port Scheduler Flags
  -> RemoteAddress:Port          Forward Weight ActiveConn
InActConn
TCP  ha.wowlinux.com:www rr
  -> www2.wowlinux.com:www      Route  1    0    0
```

(1) – real server

■ www1

```
[root@www1 /root]# /etc/rc.d/init.d/httpd start
Starting httpd: [ OK ]
```

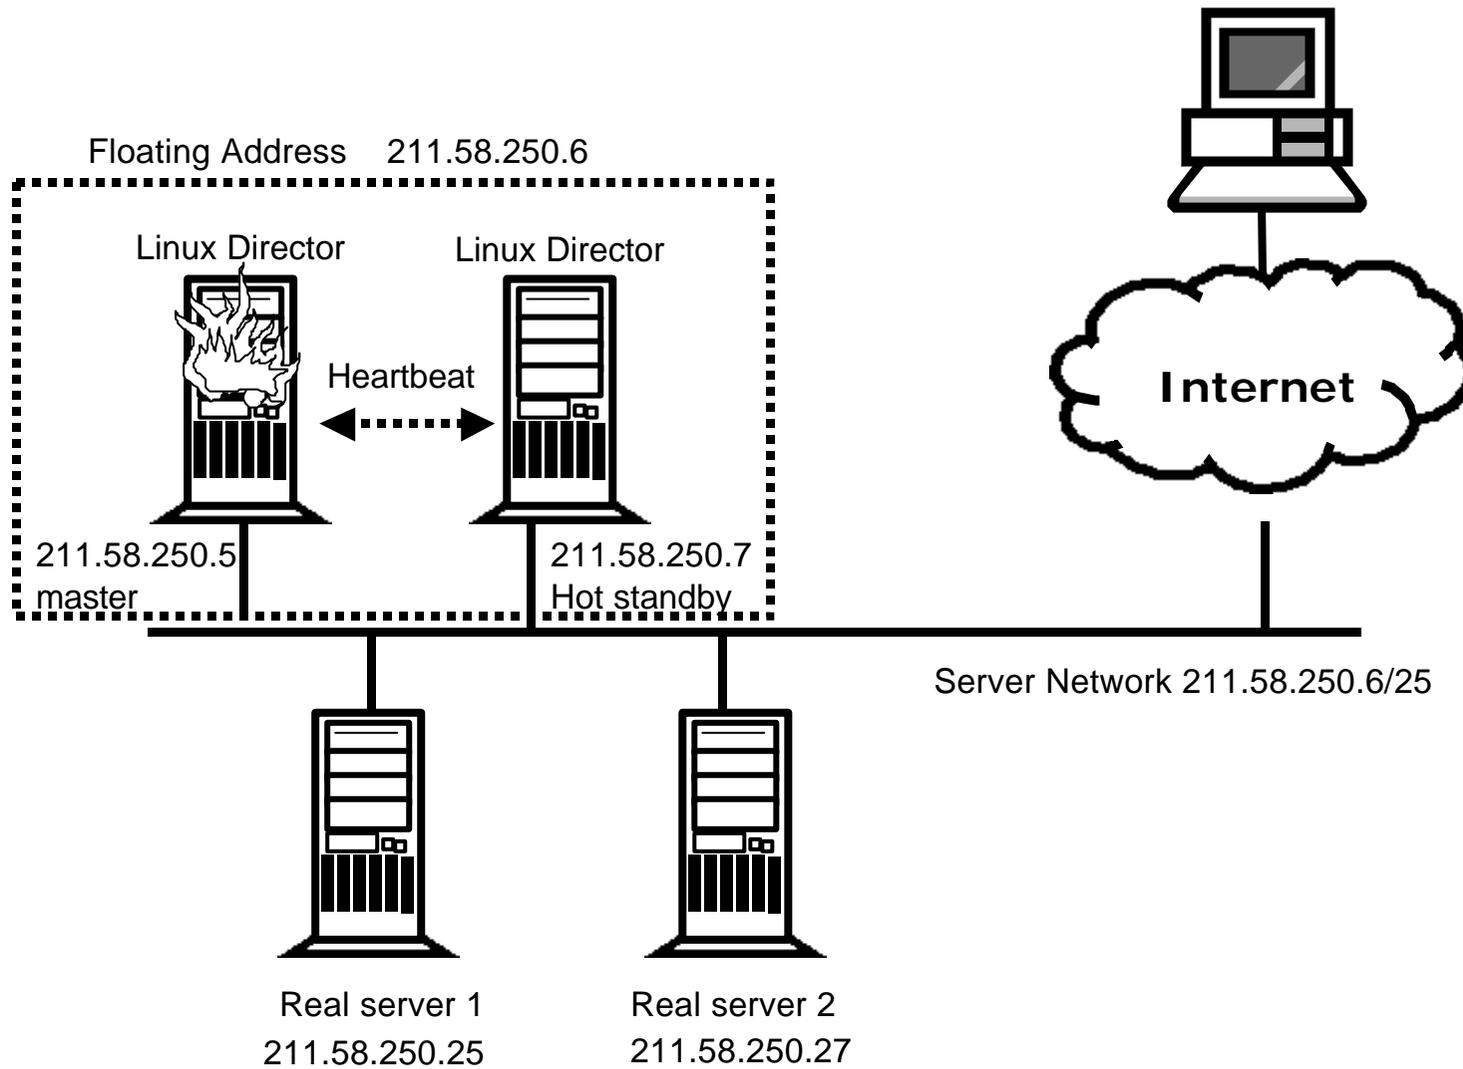
■ Ha1 /var/log/ldirectord.log

```
...
[Fri Nov 17 13:29:38 2000|ldirectord.cf] Adding real server:
211.58.250.25:80 (2*211.58.250.6:80)
```

■ Ha1 ipvsadm

```
# ipvsadm
IP Virtual Server version 0.9.12 (size=4096)
Prot LocalAddress:Port Scheduler Flags
  -> RemoteAddress:Port          Forward Weight ActiveConn
InActConn
TCP  ha.wowlinux.com:www rr
  -> www1.wowlinux.com:www      Route 1    0    0
  -> www2.wowlinux.com:www      Route 1    0    0
```

(2) – master Director /



(2) – master Director

■ Ha1 heartbeat

```
[root@ha1 /root]# /etc/rc.d/init.d/heartbeat stop
Stopping High-Availability services: [ OK ]
```

■ Ha2 /var/log/ha-log

```
...
heartbeat: 2000/11/17_13:10:55 Taking over resource group
211.58.250.6
heartbeat: 2000/11/17_13:10:55 Acquiring resource group:
ha1.wowlinux.com 211.58.250.6 ldirectord
heartbeat: 2000/11/17_13:10:55 INFO: Running
/etc/ha.d/resource.d/ldirectord start
heartbeat: 2000/11/17_13:10:55 INFO: Running
/etc/ha.d/resource.d/IPaddr 211.58.250.6 start
heartbeat: 2000/11/17_13:10:56 INFO: ifconfig eth0:0 211.58.250.6
netmask 255.255.255.128 broadcast 211.58.250.127
heartbeat: 2000/11/17_13:10:56 Sending Gratuitous Arp for
211.58.250.6 on eth0:0 [eth0]
```

(2) – master Director

■ Ha1 heartbeat

```
[root@ha1 /root]# /etc/rc.d/init.d/heartbeat start  
Starting High-Availability services: [ OK ]
```

■ Ha2 /var/log/ha-log

```
...  
heartbeat: 2000/11/17_21:13:37 Releasing resource group:  
ha1.wowlinux.com 211.58.250.6 ldirectord  
heartbeat: 2000/11/17_21:13:37 INFO: Running  
/etc/ha.d/resource.d/IPaddr 211.58.250.6 stop  
heartbeat: 2000/11/17_21:13:37 IP Address 211.58.250.6 released  
heartbeat: 2000/11/17_21:13:37 INFO: Running  
/etc/ha.d/resource.d/ldirectord stop
```

(2) – master Director

■ Ha1 /var/log/ha-log

```
heartbeat: 2000/11/17_21:18:01 received ip-request-resp
211.58.250.6 OK
heartbeat: 2000/11/17_21:18:01 Acquiring resource group:
ha1.wowlinux.com 211.58.250.6 ldirectord
heartbeat: 2000/11/17_21:18:01 INFO: Running
/etc/ha.d/resource.d/ldirectord start
heartbeat: 2000/11/17_21:18:01 INFO: Running
/etc/ha.d/resource.d/IPaddr 211.58.250.6 start
heartbeat: 2000/11/17_21:18:01 INFO: ifconfig eth0:0 211.58.250.6
netmask 255.255.255.128 broadcast 211.58.250.127
heartbeat: 2000/11/17_21:18:01 Sending Gratuitous Arp for
211.58.250.6 on eth0:0 [eth0]
```



■ GUI

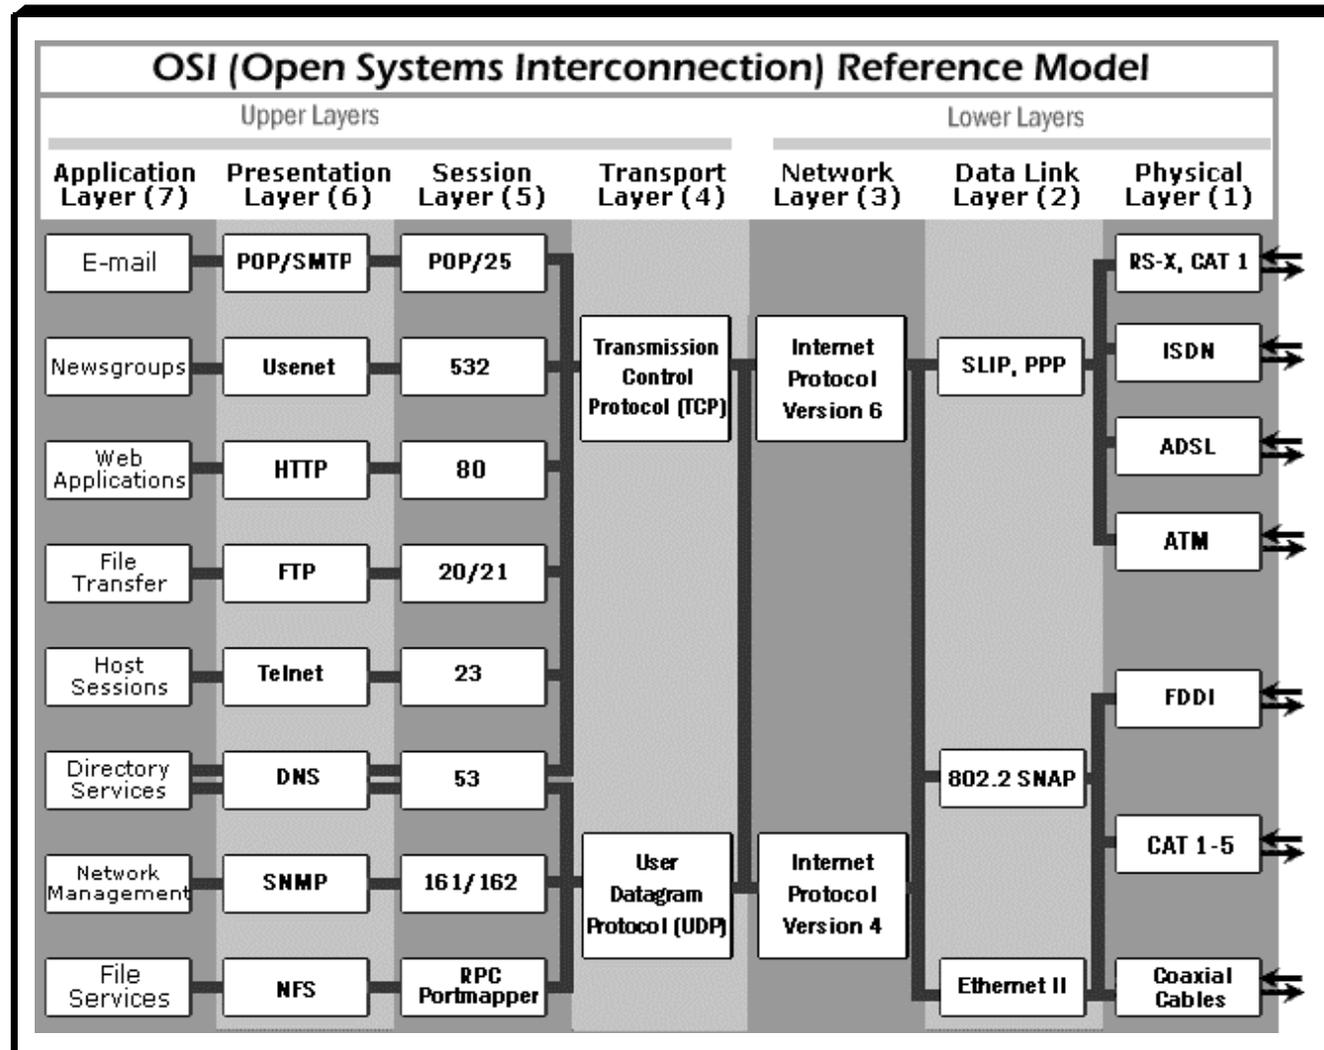
■ WEB

■ Shared File system

- GFS
- Rsync

- 가
 - <http://www.linuxvirtualserver.org>
 - <http://tunelinux.pe.kr/virtual/>
- 가
 - http://kldp.org/KoreanDoc/html/Virtual_Server - KLDP/Virtual_Server - KLDP.html
- ULTRA MONKEY
 - <http://www.au.ultramonkey.org/>
- High Availability systems under Linux
 - <http://www.linuxfocus.org/English/November2000/article179.shtml>
- Internet Core Protocols: The Definitive Guide
 - O'reilly , Erick A. hall
- 2000.7
 - 가 ,
- IP Aliasing
 - <http://kldp.org/HOWTO/mini/html/IP - Alias/IP - Alias.html>

: Open Systems Interconnection Reference Model



: OSI Reference Model

- 가 가
 - 7 : ... , , ,
 - 6 : ...
 - 5 :
 - 4 :
 - 3 : ... ()
 - 2 : ... , 5
 - 1 : ...
- : Internetworking Basics
 - <http://www.cisco.com/cpress/cc/td/cpress/fund/ith/ith01gb.htm>

!



<http://wowlinux.com>