

Bootcamp: Linux

Peter Green

<http://agathongroup.com/talks/bootcamp07/linux/>



Stuff We'll Cover

- CentOS 5 Installation and Configuration
- Virtualization
- Packet Mangling with iptables
- Configuring LAMP
- Cluster Architecture
- “Stupid Sysadmin Tricks”

Stuff We Won't Cover

- Email
- Writing code (Perl, Bash, PHP, Python, Ruby)
- Optimizing/fixing SQL
- Compilers, libraries, dependencies and other adventures in source code
- Revision control systems
- Network architecture
- Hardware
- File-sharing via NFS
- Licensing issues
- Network and system intrusion detection
- X-Windows
- VNC/remote GUI access
- Linux distribution wars
- vim versus emacs
- Compiled versus interpreted
- mutt versus pine versus elm
- bash versus tcsh
- Printing
- Samba, NCP, and the heterogeneous network
- Configuring and building a kernel from source
- BOOTP and diskless clients
- Jabber
- Time management
- Regular expressions
- sed, awk and other invaluable command-line utilities

CentOS 5 Installation and Configuration

CentOS 5: Pre- installation work

- BIOS
- Disk partitioning
- SELinux
- Package selection

CentOS 5: Post- installation work

- Silly server: GUIs are for desktops!
- System lockdown
- RPMForge
- System update

Virtualization

Virtualization: Why?

- System partitioning
- Better resource usage
- Root access for all
- Management flexibility

Virtualization: Hardware

- Intel Virtualization Technology (IVT)
- AMD Virtualization (AMD-V)
- Depends on virtualization method

Virtualization: Methods

- Virtual machine (VMWare)
- Paravirtualization (Xen)
- Virtualization on the OS level (OpenVZ)
- Our winner: OpenVZ

Virtualization: OpenVZ Installation

- Use OpenVZ yum repository
- Install and configure ovzkernel
- Configure sysctl
- Install tools

Virtualization: OpenVZ Templates

- Fastest way: use precreated template cache

<http://openvz.org/download/template/cache>

- Best way: create a template cache from template metadata

- No CentOS 5 template metadata; let's cheat!

<http://www.agathongroup.com/files/bootcamp07/linux/centos-5-metadata.tar>

Packet Mangling with iptables

Packet Mangling with iptables: Usage

- iptables [-AD] chain rule
- iptables -I chain [rulenum] rule
- iptables -D chain rulenum
- iptables -P chain target
- iptables -L chain

Packet Mangling with iptables: Common args

- [-sd] address[/mask]
- -p protocol
- [--sport/--dport] port[:port]
- -nv
- -j target
- -t table
- [-io] interface

Packet Mangling with iptables: “filter” table

- Used to filter traffic touched by the machine
- INPUT: packets destined for the local machine
- FORWARD: packets routing through the machine
- OUTPUT: packets generated on the local machine

Packet Mangling with iptables: “nat” table

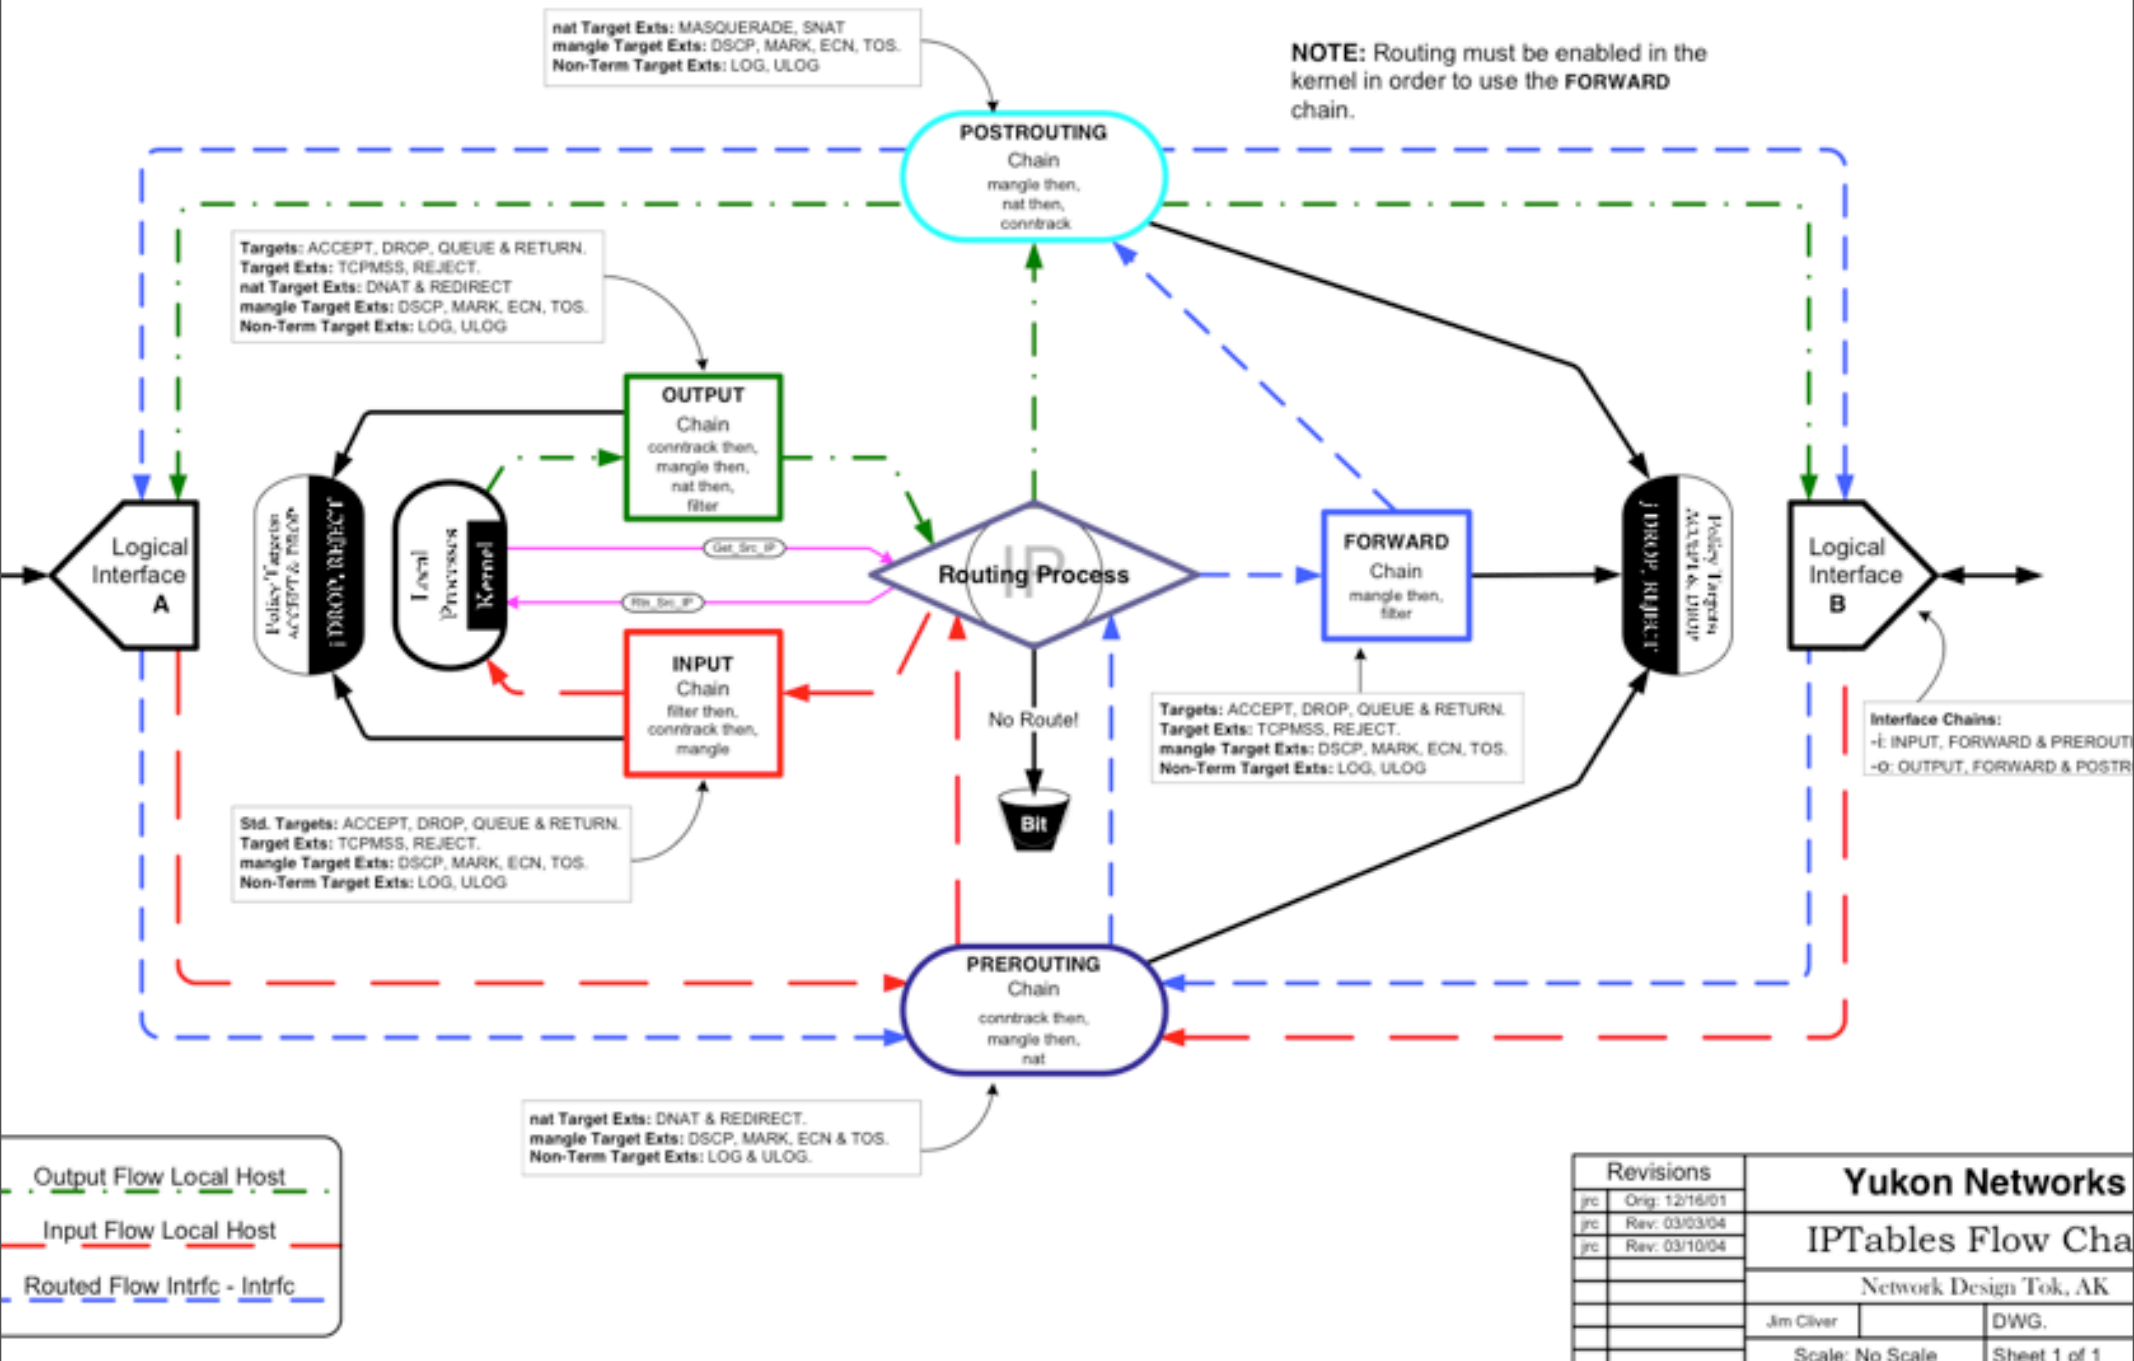
- Used when a packet that creates a new connection is encountered
- PREROUTING: “as soon as packets come in”, before routing decision
- OUTPUT: locally-generated packets, before routing decision
- POSTROUTING: “as packets are about to go out”, after routing

Packet Mangling with iptables: “mangle” table

- Used for specialized packet alteration
- INPUT: packets destined for the local machine
- FORWARD: packets routing through the local machine
- OUTPUT: locally-generated packets, before routing decision

Packet Mangling with iptables: “mangle” table

- PREROUTING: “as soon as packets come in”, before routing decision
- POSTROUTING: “as packets are about to go out”, after routing



Revisions	
jc	Orig: 12/16/01
jc	Rev: 03/03/04
jc	Rev: 03/10/04

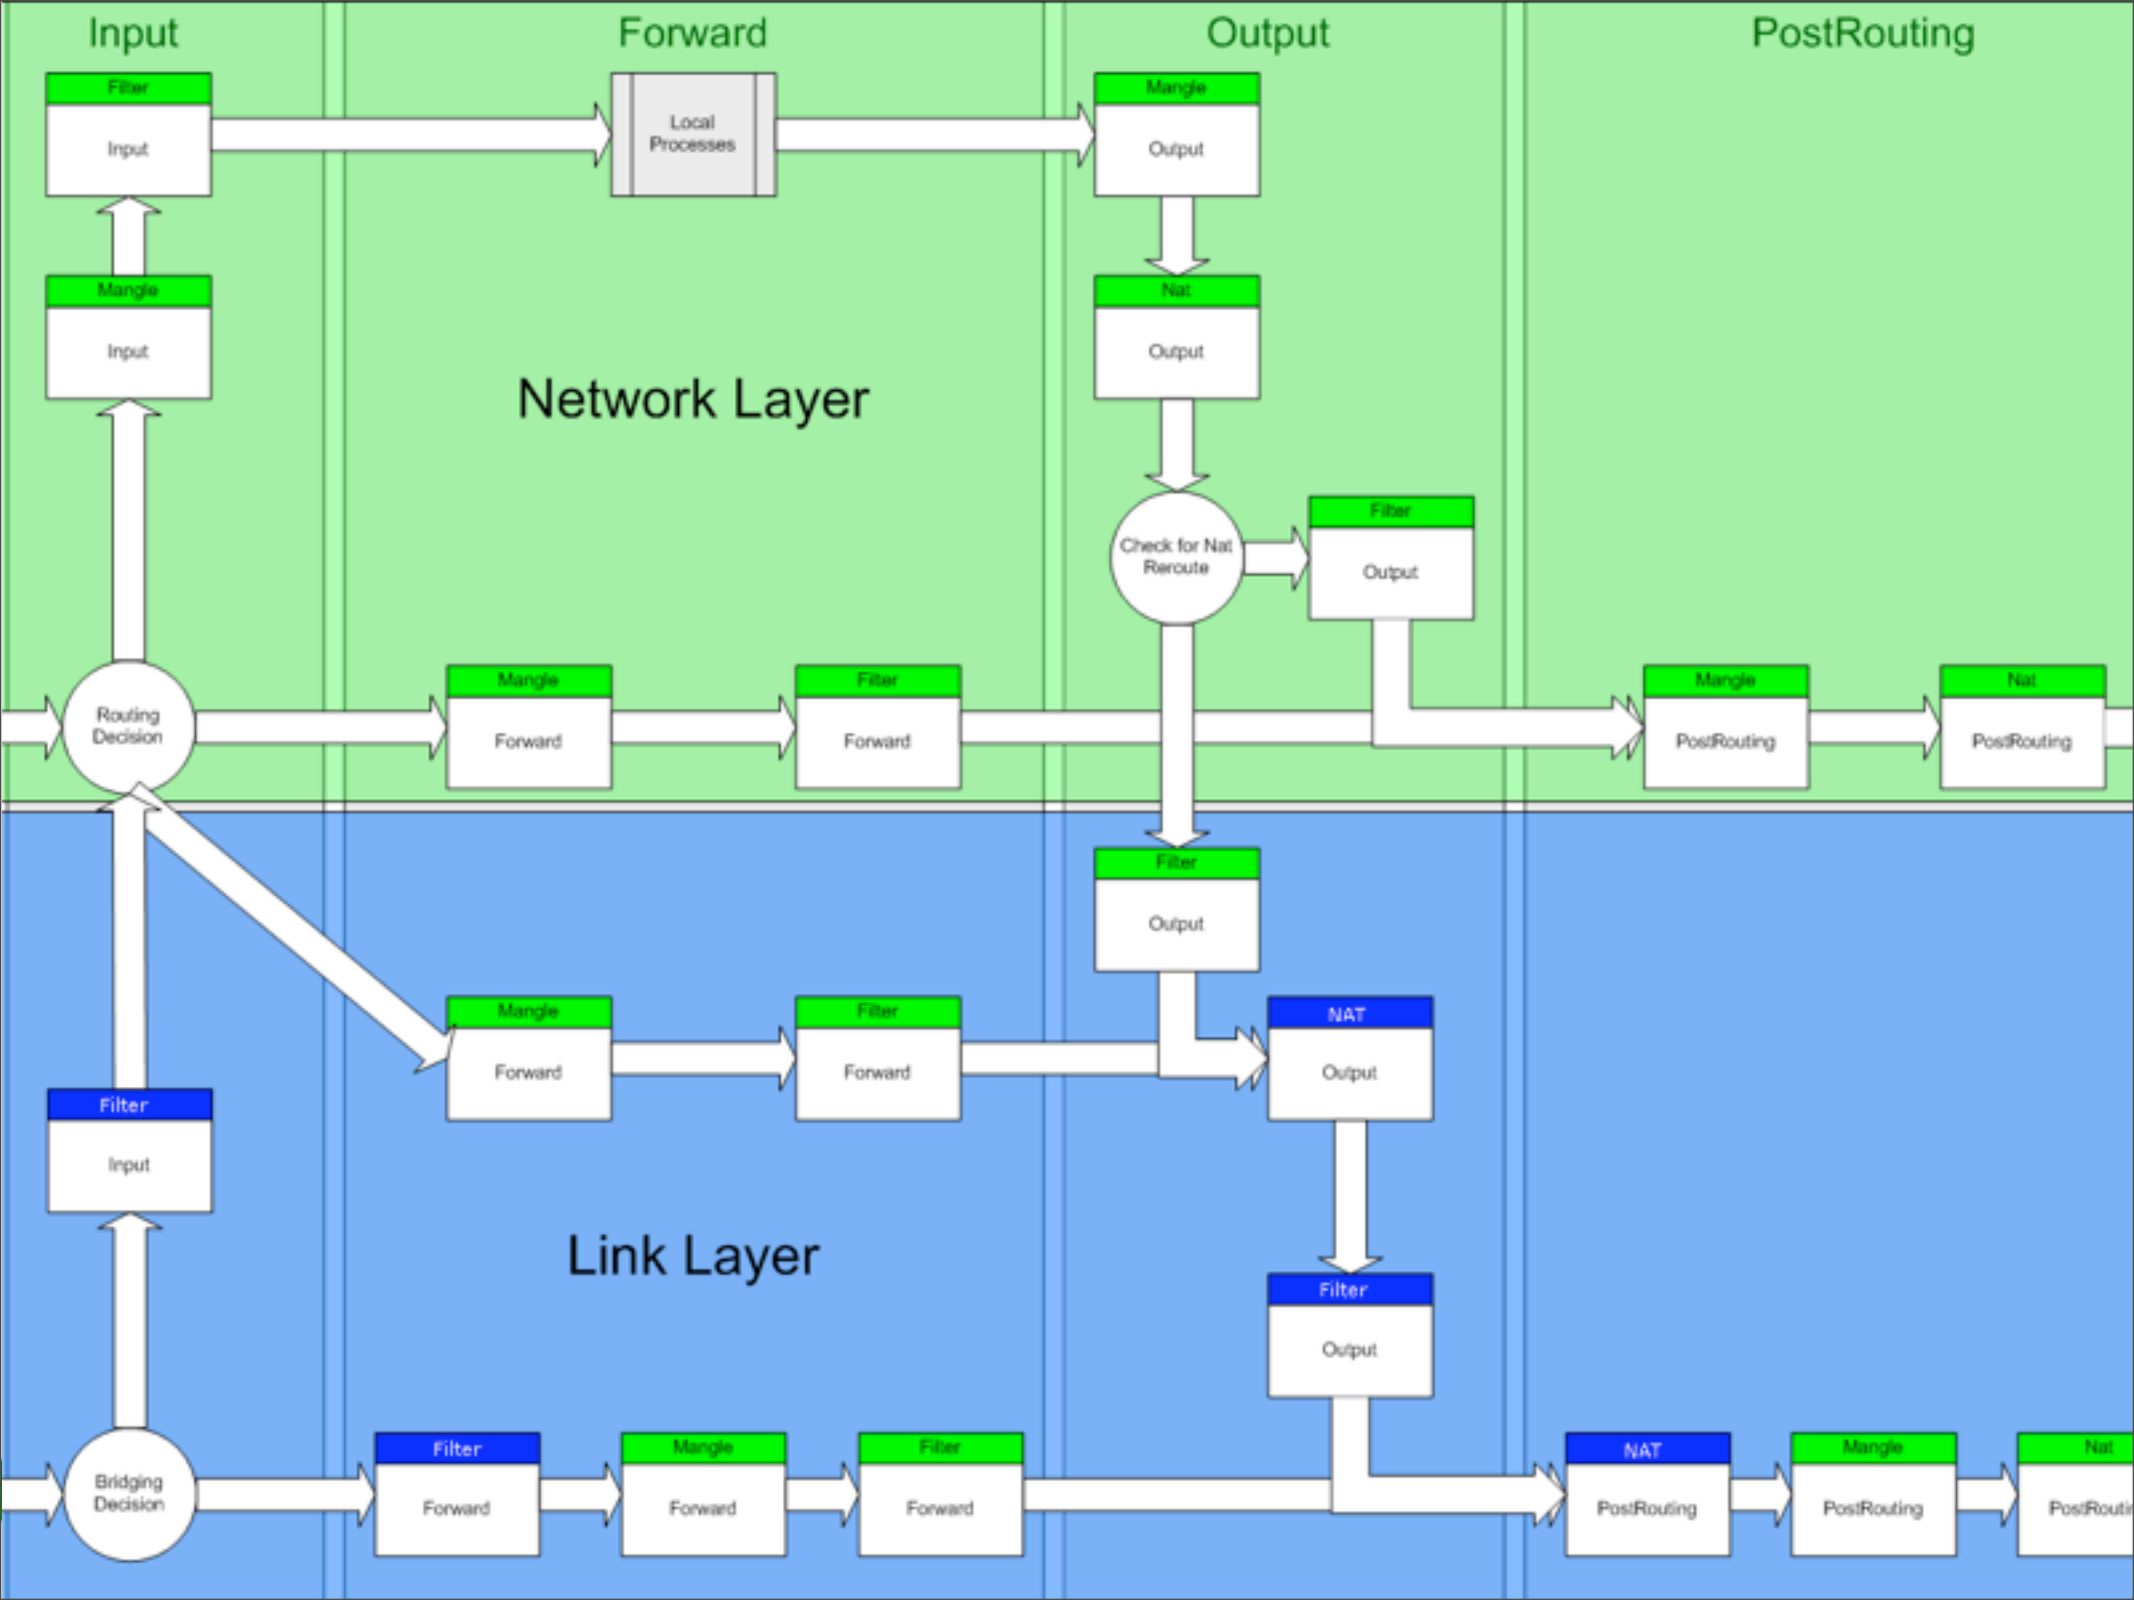
Yukon Networks

IPTables Flow Cha

Network Design Tok, AK

Jim Cliver DWG.

Scale: No Scale Sheet 1 of 1



Packet Mangling with iptables: Built-in targets

- ACCEPT: let the packet through
- DROP: drop the packet
- QUEUE: pass the packet to a userspace program
- RETURN: stop in this chain and return the to “calling” chain

Packet Mangling with iptables: Extension targets

- DNAT: alter the destination address
`nat:PREROUTING` and `nat:OUTPUT` only
- (What about MASQUERADE?)
- SNAT: alter the source address
`nat:POSTROUTING` only
- LOG: log matching packets via the kernel log

Packet Mangling with iptables: Extension targets

- **FILTER**: like DROP, but send back an error packet in response
`filter:* only`
- **MARK**: set the netfilter mark
`mangle:* only`
- **REDIRECT**: like DNAT to the primary address of incoming interface
`nat:PREROUTING and nat:OUTPUT only`

Packet Mangling with iptables: Examples

- Simple router (DHCP)

```
IPT=/sbin/iptables  
EXT_IF=eth0  
INT_IF=eth1  
INT_NET="192.168.1.0/24"  
$IPT -t nat -A POSTROUTING -o $EXT_IF \  
-s $INT_NET -j MASQUERADE
```

Packet Mangling with iptables: Examples

- Simple router (static, spoof protection)

```
IPT=/sbin/iptables
EXT_IF=eth0
INT_IF=eth1
INT_NET="192.168.1.0/24"
EXT_IP="64.78.150.1"
$IPT -A FORWARD -i $EXT_IF -s $INT_NET -j DROP
$IPT -A FORWARD -i $INT_IF '!' -s $INT_NET -j DROP
$IPT -t nat -A POSTROUTING -o $EXT_IF \
-s $INT_NET -j DNAT --to $EXT_IP
```

Packet Mangling with iptables: Examples

- Add “allow HTTP/S, deny all others”

```
IPT=/sbin/iptables
EXT_IF=eth0
INT_IF=eth1
INT_NET="192.168.1.0/24"
EXT_IP="64.78.150.1"
$IPT -A FORWARD -i $EXT_IF -s $INT_NET -j DROP
$IPT -A FORWARD -i $INT_IF '!' -s $INT_NET -j DROP
$IPT -A FORWARD -i $INT_IF -p tcp -m multiport \
--dports 80,443 -j ACCEPT
$IPT -A FORWARD -i $INT_IF -j DROP
$IPT -t nat -A POSTROUTING -o $EXT_IF \
-s $INT_NET -j DNAT --to $EXT_IP
```

Packet Mangling with iptables: Examples

- Log all outbound HTTP packets

```
IPT=/sbin/iptables
EXT_IF=eth0
INT_IF=eth1
$IPT -A FORWARD -i $INT_IF -o $EXT_IF \
-p tcp --dport 80 -j LOG -m limit --limit 20/min \
--log-prefix "FILTER "
```

Packet Mangling with iptables: Examples

- Drop XMAS/NULL packets

```
IPT=/sbin/iptables
```

```
EXT_IF=eth0
```

```
INT_IF=eth1
```

```
$IPT -A FORWARD -i $EXT_IF -o $INT_IF \
```

```
-p tcp --tcp-flags ALL ALL -j DROP
```

```
$IPT -A FORWARD -i $EXT_IF -o $INT_IF \
```

```
-p tcp --tcp-flags ALL NONE -j DROP
```

Configuring LAMP

Configuring LAMP: httpd

- httpd is Apache, “the most popular web server ... since April 1996”
- Managing config files
- Addressing virtual hosts
- Abusing error handlers
- mod_rewrite: The Swiss Army Knife of URL Manipulation

Configuring LAMP: httpd

Managing config files

- Include /etc/httpd/conf.d/*.conf

```
-rw-r--r-- 1 root root 279 Jun 8 12:28 alive.conf
-rw-r--r-- 1 root root 311 Aug 22 11:20 awstats.conf
-rw-r--r-- 1 root root 321 Jun 8 13:01 cac.conf
-rw-r--r-- 1 root root 349 Jul 23 09:48 cac.include
-rw-r--r-- 1 root root 303 Jun 8 12:34 creation.conf
-rw-r--r-- 1 root root 392 Jun 26 16:28 README
-rw-r--r-- 1 root root 9677 Jun 26 16:28 ssl.conf
drwxr-xr-x 2 root root 4096 Aug 10 15:30 ssl.crt
-rw-r--r-- 1 root root 286 May 7 15:21 ssl.include
drwxr-xr-x 2 root root 4096 Aug 10 15:29 ssl.key
-rw-r--r-- 1 root root 258 May 1 16:31 TEMPLATE
```

Configuring LAMP: httpd

Managing config files

- alive.conf: the typical config file

```
<VirtualHost *:80>  
  ServerName aliveupdates.com  
  ServerAlias www.aliveupdates.com  
  ServerAdmin webmaster@aliveupdates.com  
  DocumentRoot /home/alive/public_html  
  TransferLog logs/vhosts/alive/access-log  
  ScriptAlias /cgi-bin/ /home/alive/public_html/cgi-bin/  
</VirtualHost>
```

Configuring LAMP: httpd

Managing config files

- cac.conf: dealing with SSL

```
<VirtualHost *:80>
  Include conf.d/cac.include
</VirtualHost>
<VirtualHost 10.0.3.30:443>
  Include conf.d/cac.include
  Include conf.d/ssl.include
  SSLCertificateFile      conf.d/ssl.crt/www.comealivecruises.com.crt
  SSLCertificateKeyFile  conf.d/ssl.key/www.comealivecruises.com.key
</VirtualHost>
```

Configuring LAMP: httpd

Managing config files

- cac.include: dealing with SSL (cont.)

```
ServerName www.comealivecruises.com
ServerAlias comealivecruises.com
ServerAdmin webmaster@comealivecruises.com
DocumentRoot /home/cac/public_html
TransferLog logs/vhosts/cac/access-log
ScriptAlias /cgi-bin/ /home/cac/public_html/cgi-bin/
```

Configuring LAMP: httpd

Managing config files

- ssl.include: dealing with SSL (cont.)

```
SSLEngine on
SSLCipherSuite \
ALL:!ADH:!EXPORT56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP:+eNULL
<Files ~ "\.(cgi|shtml|phtml|php3?)$" >
    SSLOptions +StdEnvVars
</Files>
SetEnvIf User-Agent ".*MSIE.*" \
    nokeepalive ssl-unclean-shutdown \
    downgrade-1.0 force-response-1.0
```

Configuring LAMP: httpd

Managing config files

- Use templates

```
<VirtualHost *:80>  
  ServerName #DOMAIN#  
  ServerAlias www.#DOMAIN#  
  ServerAdmin webmaster@#DOMAIN#  
  DocumentRoot /home/#USER#/public_html  
  TransferLog logs/vhosts/#USER#/access-log  
  ScriptAlias /cgi-bin/ /home/#USER#/public_html/cgi-bin/  
</VirtualHost>
```

```
shell> sed 's/#DOMAIN#/ais.cx/g; s/#USER#/ais/g' < Tmpl > ais.conf
```

Configuring LAMP: httpd

Addressing virtual hosts

- IP address or wildcard?

- Non-SSL: wildcard

```
NameVirtualHost *:80
```

```
<VirtualHost *:80>
```

- SSL: IP address

```
<VirtualHost 10.0.3.30:443>
```

- Don't mix IPs and wildcards!
- Don't hardcode IPs unless necessary!

Configuring LAMP: httpd

Abusing error handlers

- ErrorDocument code document
- Serve static content (images, CSS) from the filesystem
- Set a 404 ErrorDocument to serve other content “dynamically”
- But this is an ugly solution...

Configuring LAMP: httpd

Abusing error handlers

- Defeats stats
- Allows for very simplistic pages only
- Creates additional load
- Inelegant
- There has to be a better way...

Configuring LAMP: httpd mod_rewrite

- mod_rewrite: The Swiss Army Knife of URL Manipulation
- Rewrite URLs transparently
- Proxy requests to backend servers
- Lots, lots more

http://httpd.apache.org/docs/2.2/mod/mod_rewrite.html

Configuring LAMP: MySQL

- Flexible, extensible, scalable, high-performance, open-source “Standard Query Language” (SQL) database server
- Useful my.cnf variables
- Replication
- Up and coming

Configuring LAMP: MySQL

Useful my.cnf variables

- `old_passwords = 1`
 - ▶ Useful only when upgrading from MySQL 3.x
- `skip-innodb`
- `max_connections = N`
- `log-bin`
 - ▶ Necessary for replication, but useful even without
- `port = N`
- `socket = [path]`

Configuring LAMP: MySQL

Useful my.cnf variables

- Server tuning – black art?
- Sample .cnf files

<http://dev.mysql.com/doc/refman/5.0/en/server-parameters.html>

<http://dev.mysql.com/doc/refman/5.0/en/server-system-variables.html>

<http://dev.mysql.com/doc/refman/5.0/en/server-status-variables.html>

Configuring LAMP: MySQL Replication

- Take data and copy it elsewhere
- Data backup/security
- Database backups
- Analytics
- Geographic diversity

Configuring LAMP: MySQL Replication

- The Good
 - ▶ Asynchronous
 - ▶ Fast
 - ▶ Simple to switch to a new master
- The Bad
 - ▶ Only one master (usually!)
 - ▶ Asynchronous

Configuring LAMP: MySQL

Up and coming

- Circular (multi-master) replication
- MySQL Cluster
- MySQL Proxy

Configuring LAMP: PHP

- PHP – “PHP Hypertext Processor”
- “Widely-used general-purpose scripting language that is especially suited for Web development”
- Security
- Performance
- Contextual configuration

Configuring LAMP: PHP

Security

- “It’s getting better, a little better all the time.” — Paul
- “(It can’t get no worse!)” — John
- Security “versus” usability
- Fundamental problems lead to poorly-designed “fixes”

Configuring LAMP: PHP Security

- `safe_mode = ?`
- `register_globals = off`
 - ▶ Both going away in PHP6 — YAY!
- `mail()` is obsolete
 - ▶ PHPMailer
 - ▶ PEAR's Mail class
- Session handling

Configuring LAMP: PHP Security

<http://php-security.org/>

<http://hardened-php.net/>

<http://phpsec.org/>

<http://phpsecurity.org/>

“8 chapters. 30 exploits.”

[http://www.google.com/search?
q=PHP+security](http://www.google.com/search?q=PHP+security)

Configuring LAMP: PHP Performance

- `memory_limit = N`
- `register_long_arrays = false`
- Extensions
- PEAR

Configuring LAMP: PHP

Contextual configuration

- `php_value/php_admin_value`
- `php_flag/php_admin_flag`
- No PHP constants (`E_ALL`)!

<http://us.php.net/manual/en/ini.php#ini.list>

Cluster Architecture

Cluster Architecture: Concepts and Options

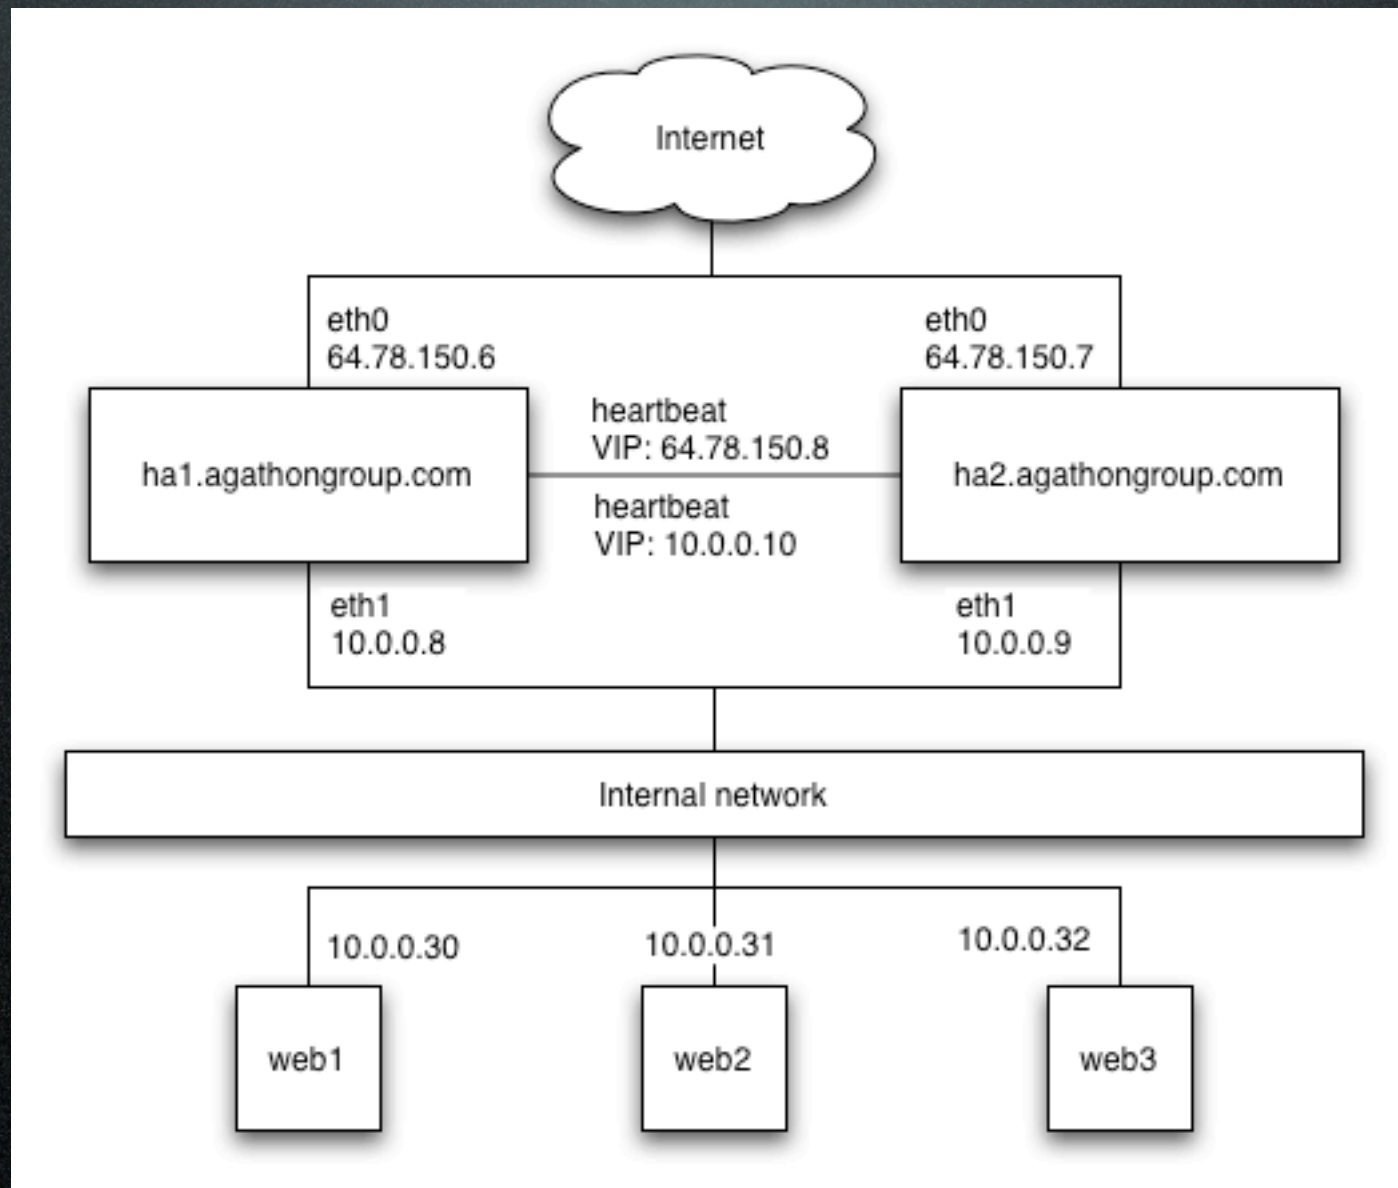
- High availability
- Load balancing
- Payware (e.g., F5 BIG-IP)
- Open Source (e.g., Ultramonkey)

Cluster Architecture:

Why cluster?

- Single server
 - ▶ Single point of failure
 - ▶ Simple (on-disk storage, programming)
- Cluster
 - ▶ Redundant systems
 - ▶ Complex (replicated/shared storage)

Cluster Architecture: Typical setup



Cluster Architecture: High availability (HA)

- heartbeatd
- “Are you there?”
- One primary, one backup
- Simple: live web server, hot failover web server
- Modular: HA load balancers

Cluster Architecture: High availability (HA)

- Sample configuration (incomplete!)

```
ha.agathongroup.com \  
  IPaddr::10.0.0.10/8/eth1 \  
  IPaddr::64.78.150.8/26/eth0 ldirector::agathongroup.cf \  
  service::iptables \  
  service::ldap::restart
```

```
#baud 19200  
#serial /dev/ttyS0  
#bcast eth1  
#mcast eth0 225.0.0.1 694 1 0  
ucast eth1 10.0.0.8  
auto_failback on  
node ha1.agathongroup.com  
node ha2.agathongroup.com
```

Cluster Architecture: Load balancing

- Linux Virtual Service (LVS)
- 1+ servers behind a load balancer
- Identical content
- Potentially disparate servers

Cluster Architecture: Load balancing

- Simple web service

```
checktimeout=2
checkinterval=1
autoreload=no
logfile="/var/log/ldirectord/agathongroup.log"
quiescent=yes
virtual=64.78.150.8:80
    fallback=127.0.0.1:80
    real=10.0.0.30:80 masq
    real=10.0.0.31:80 masq
    checktype=connect
    scheduler=wrr
    protocol=tcp
```

Cluster Architecture: Load balancing

- Silly iptables+LVS trick: Use MARK to group packets for routing by LVS

```
$IPT -t mangle -A PREROUTING -p tcp -m multiport \  
--dports 25,110,143,993,995 -j MARK --set-mark 0x1
```

```
virtual=1  
real=10.0.0.4 masq  
service=smtp  
checktype=connect  
checkport=25  
scheduler=rr  
protocol=fwm
```

Cluster Architecture: Centralization

- Storage
 - ▶ Data synchronization
 - ▶ NFS
- Authentication, authorization and identification
 - ▶ Old school: NIS/YP
 - ▶ LDAP

Cluster Architecture: LDAP

- LDAP service
- Authentication and authorization:
pam_ldap
- Identification: Name Service
Switch (NSS) with nss_ldap

“Stupid Sysadmin Tricks”

“Stupid Sysadmin Tricks”

Exploring /proc

- /proc: a real-time, memory-resident pseudo-filesystem that tracks the state of the system
- /proc/\$PID
- /proc/sys/net/ipv4

“Stupid Sysadmin Tricks”

Exploring /proc

- `/proc/scsi/*` and `/proc/ide/*`
- `/proc/cpuinfo`
- `/proc/mdstat`
- `/proc/net`
- `/proc/config.gz`

“Stupid Sysadmin Tricks”

Password-less ssh

- Goal: allow logins using keys, rather than passwords
- Local side: Build `~/.ssh/id_rsa.pub`
`ssh-keygen -N '' -t rsa`
- Remote side: copy `id_rsa.pub` into
`~/.ssh/authorized_keys`

“Stupid Sysadmin Tricks”

Password-less ssh

- RSA vs. DSA, SSHv1 vs. SSHv2
- sshd_config:
PubkeyAuthentication yes
- Permissions: go-wx on:
~/.ssh/authorized_keys/
~/.ssh/
~

“Stupid Sysadmin Tricks”

Booting utilities

- Knoppix live CD
- linux rescue
- `init=/bin/sh`

“Stupid Sysadmin Tricks”

System stats with lsof

- lsof: “list open files”
- What is a file in *nix?
 - ▶ Regular file or directory
 - ▶ Block or character special file
 - ▶ Executing text reference
 - ▶ Library
 - ▶ Stream
 - ▶ Network file (Internet socket, NFS file or UNIX domain socket)

“Stupid Sysadmin Tricks”

System stats with lsof

- See what's listening to TCP port 80
`lsof -n -i tcp:80`
- See what has /sbin/init open
`lsof /sbin/init`

“Stupid Sysadmin Tricks”

Useless Use of cat

- `cat file | less`
`less file`
- `cat file | wc -l`
`wc -l file` OR `wc -l < file`
- `for i in `cat file`; do echo $i; done`
`while read i; do echo $i; done < file`
- `for i in `cat servers`; do \
ssh $i uptime; done`
`xargs -i ssh {} uptime < servers`

“Stupid Sysadmin Tricks”

UseFUL Use of cat

- `(foo ; bar ; cat file ; baz) | quux`
- `cat file1 file2 | wc -l`
- `cat<<EOF`
`data`
`EOF`
- `cat` == “concatenate”! In general, don’t use with `<2` files.

“Stupid Sysadmin Tricks”

sudo

- Why sudo rules
 - Dangers of logging in as root
 - Audit trail
 - Granular permissions
- Incorrect uses of sudo
 - NOPASSWD
 - sudo su -

“Stupid Sysadmin Tricks”

sudo

- Where sudo falls short
 - `foo > /root/file`
 - `foo | sudo tee /root/file`
 - `foo < /etc/shadow`
 - `sudo cat /etc/shadow | foo`
 - UUOC!
 - Not really
- Fast typing can lead to mistakes similar to those of being logged in as root

“Stupid Sysadmin Tricks”

One-liners

- Remount a running filesystem
`mount -o remount,[new options] /`
- Remove a file whose name starts with a hyphen:
`rm -- -file`
`rm ./-file`
- Ignore STDERR
`grep foo * 2>/dev/null`

“Stupid Sysadmin Tricks”

One-liners

- Find files older than (e.g.) 10 days

```
find . -type f -mtime -10
```

- Find mail() calls in all PHP files

```
find . -name "*.php" | xargs grep 'mail('
```

- Find mail() calls in all PHP files altered on August 22 (ghetto!)

```
find . -name "*.php" -print0 | \  
xargs -0 ls -ld | grep "Aug 22" | \  
awk '{print $9}' | xargs grep 'mail('
```