

# jail(8)

**Hacking UNIX with FreeBSD jail(8), Secure Virtual Servers**  
Presentation for DefCon 14, by Isaac Levy, (.ike)



# .ike Context

- I have used jails extensively for web application servers and software development purposes
- the methodology I'm presenting here is attempting to be 'stock' UNIX (no 'ike-specific' magic formulas)
- I am not a jail author, no commit bit...

# Warranty / Announcement

- I'll be out and about later if anyone has more complex questions or strategies they want to discuss
- I'm \*trying\* to stick to classic UNIX process and ideas, and 'stock' methodology (no ike-specific magic)
- I'm assuming you all know your way around various \*NIX Operating Systems

scale, patterns, complexity  
(a big picture exercise)






Film: Powers of Ten, 1977, Charles and Ray Eames

<http://www.powersof10.com/>  
<http://www.youtube.com/watch?v=4i6B7HzijSo>















Film: Powers of Ten, 1977, Charles and Ray Eames

<http://www.powersof10.com/>  
<http://www.youtube.com/watch?v=4i6B7HzijSo>












Film: Powers of Ten, 1977, Charles and Ray Eames

<http://www.powersof10.com/>  
<http://www.youtube.com/watch?v=4i6B7HzijSo>





Internet universe, (according to ike, today.)



# America's CryptoKids™

Future Codemakers & Codebreakers



**MEET THE GANG**

CHARACTER BIOGRAPHIES

**HOW CAN I WORK FOR NSA?**

STUDENT RESOURCES

**OPERATION DIT DAH**

GAMES & ACTIVITIES

**MAKE YOUR OWN SECRET CODES!**

CODES & CIPHERS

**BATTLE OF THE BADGES**

GAMES & ACTIVITIES

*Codes & Ciphers*

*Games & Activities*

*Student Resources*

*Character Biographies*

[NSA.gov](#)

[Links](#)

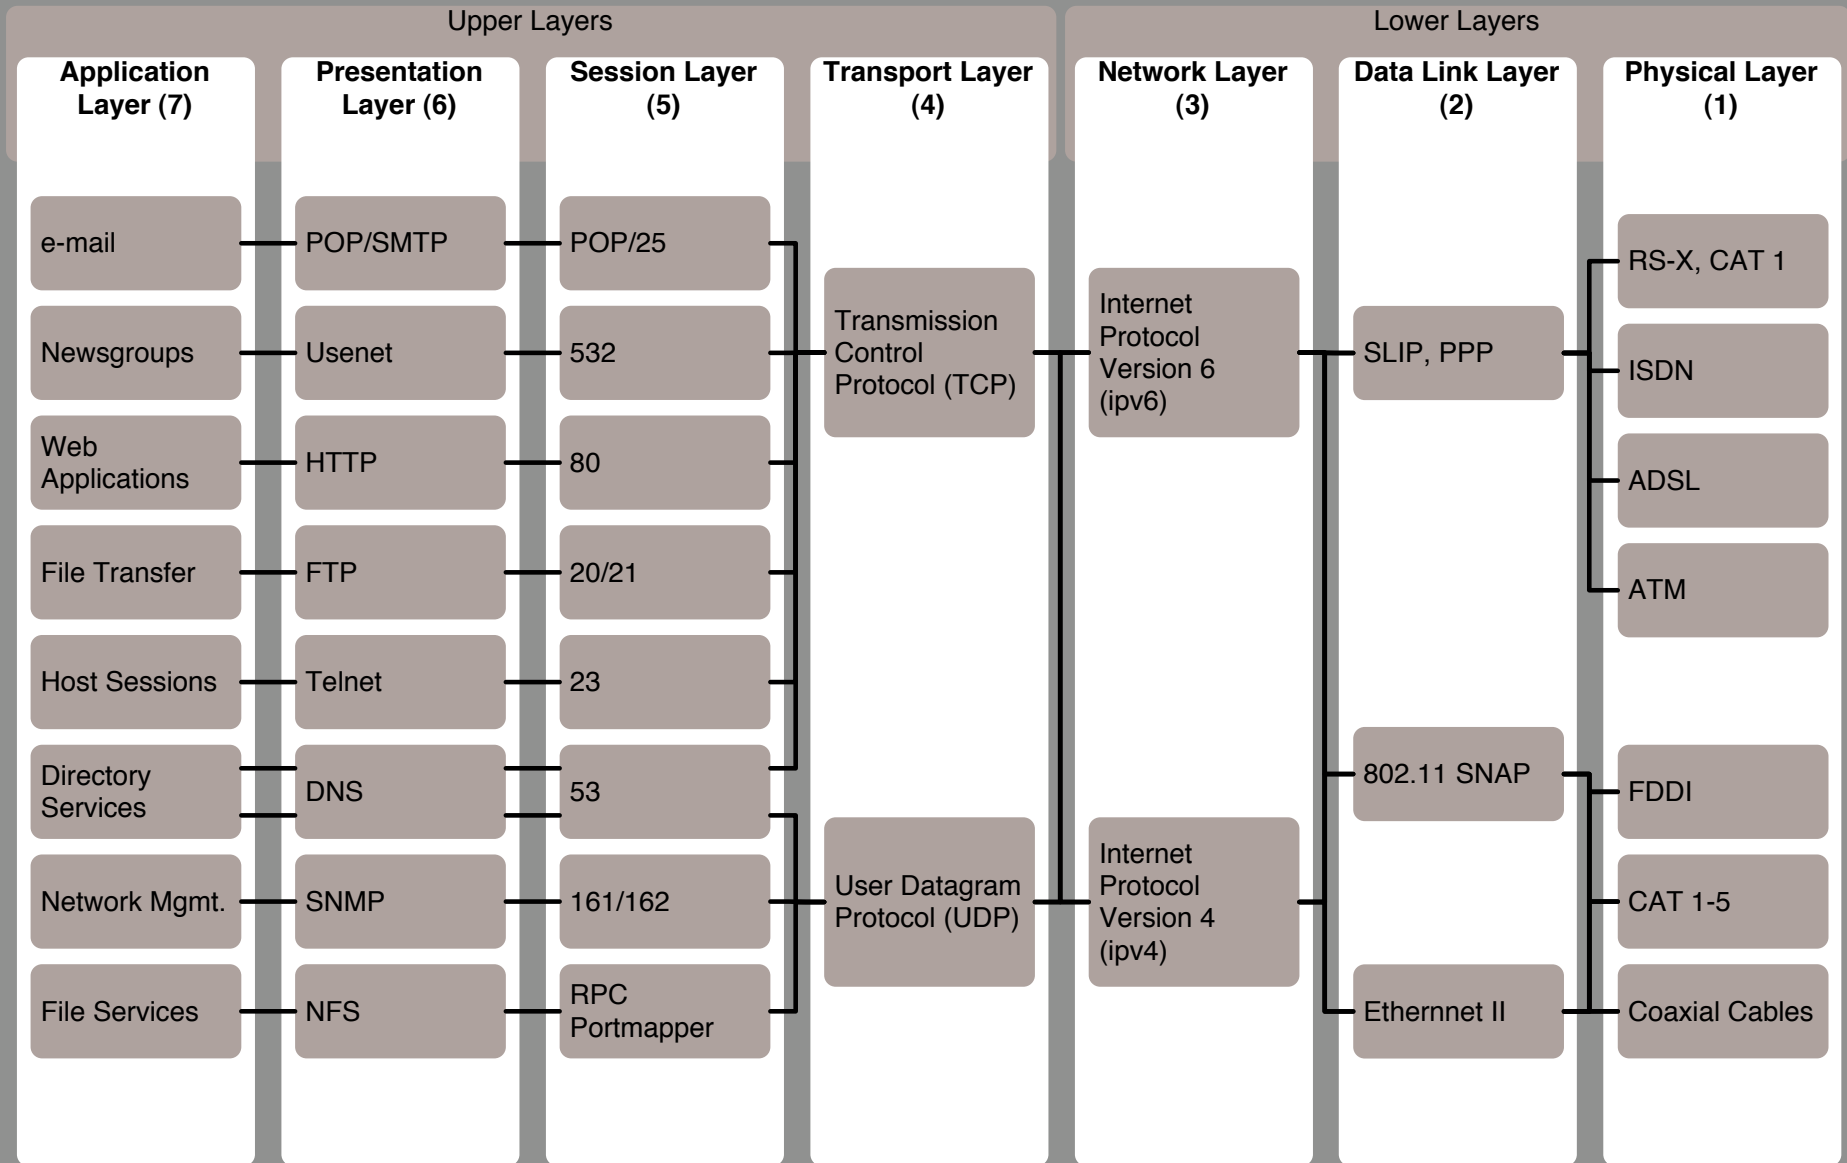
[Parents & Teachers](#)

[Trademark](#)








[Accessibility](#)

[Privacy & Security](#)

# Open Systems Interconnection (OSI) Reference Model




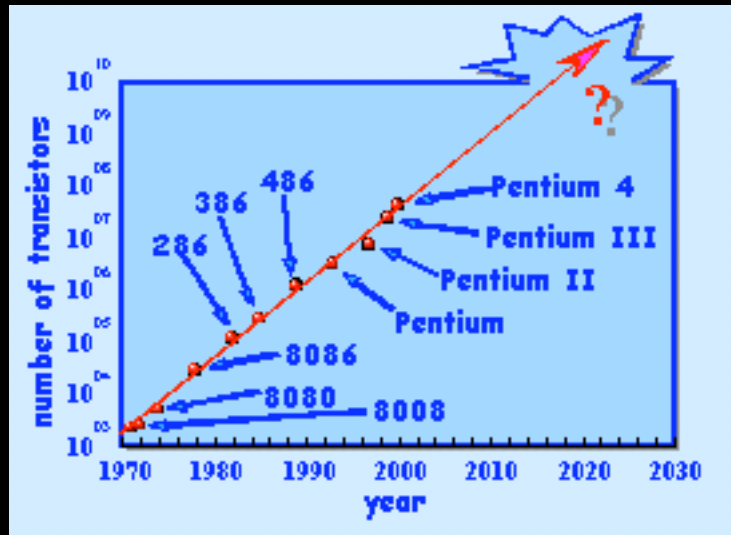













Internet universe, (according to ike, today.)

# yadda yadda






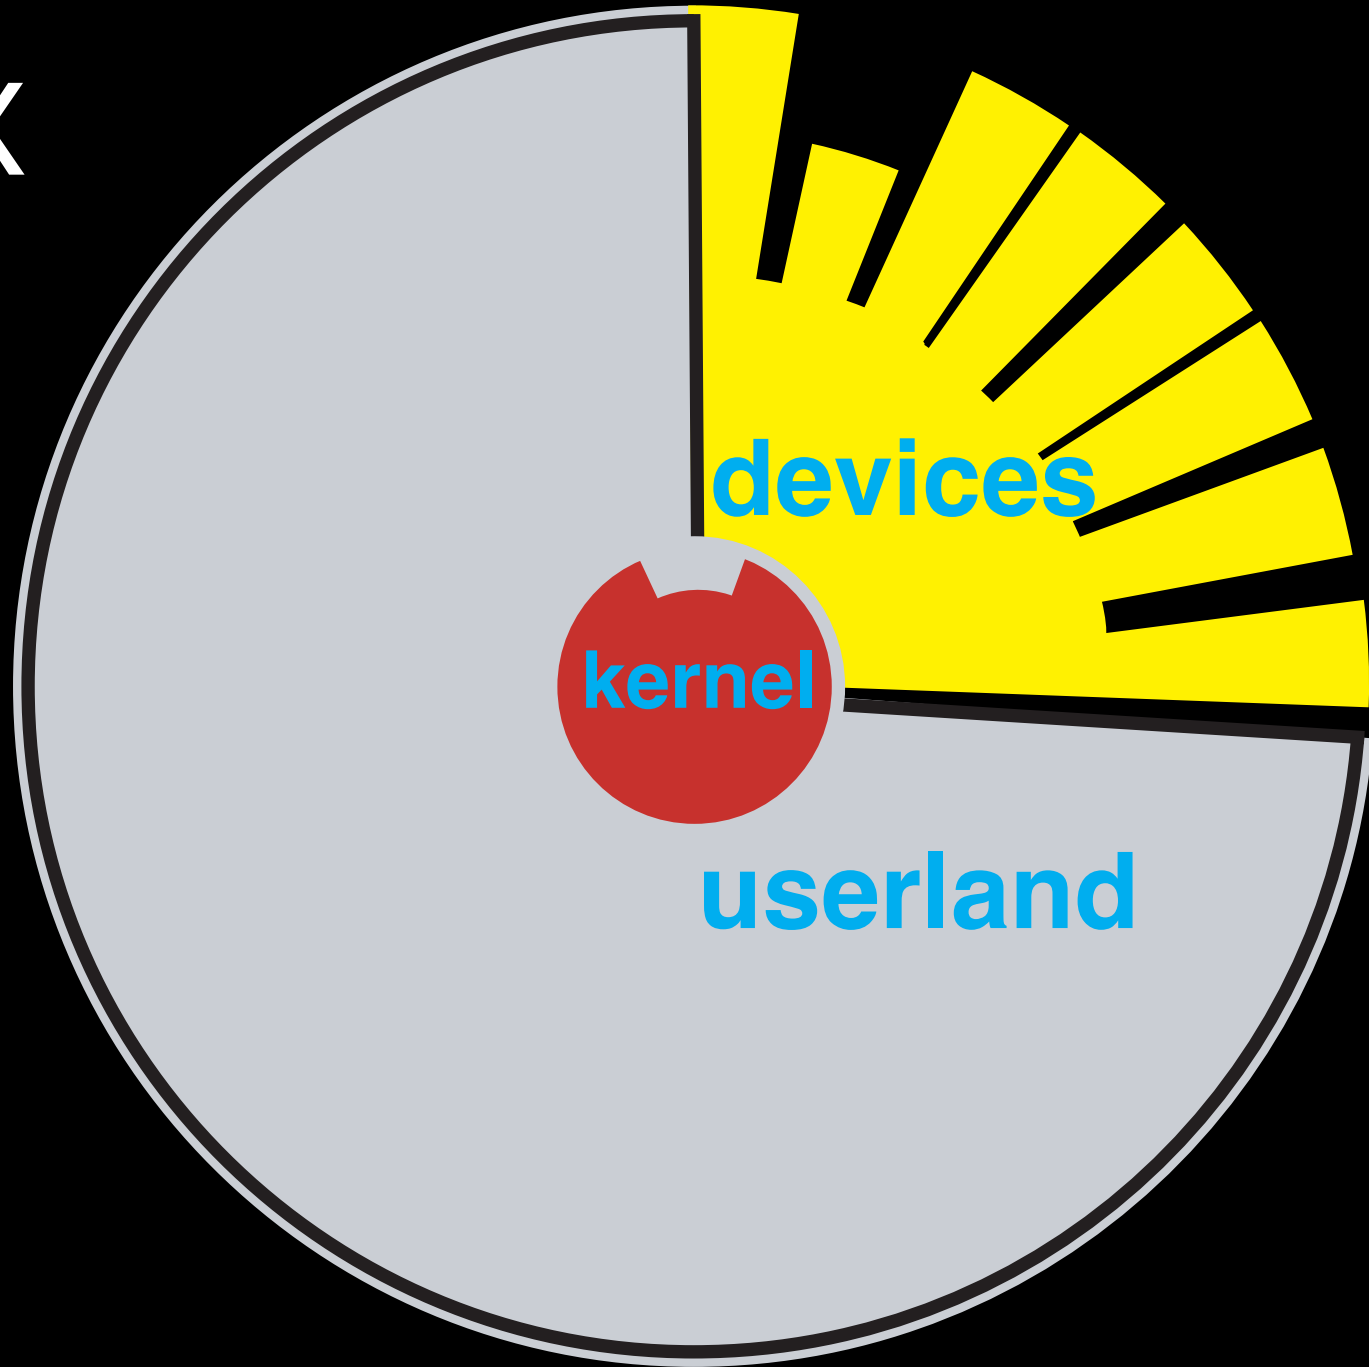


Internet universe, (according to ike, today.)


# UNIX




# UNIX



# UNIX



# UNIX




# Spiral Galaxy




*Spiral Galaxy NGC 1232*

Our world is complex

# UNIX




# Helium Atom




Our world is simple too...

# UNIX










Mandelbrot Fractal - Julia set





virtual  
UNIX's







kernel

root

usr

home

user



user

local

src

ports



jails





# User


# ports








ser



S



virtual  
UNIX's



# The UNIX Time-Sharing System\*

*D. M. Ritchie and K. Thompson*


## ABSTRACT

Unix is a general-purpose, multi-user, interactive operating system for the larger Digital Equipment Corporation PDP-11 and the Interdata 8/32 computers. It offers a number of features seldom found even in larger operating systems, including

- i A hierarchical file system incorporating demountable volumes,
- ii Compatible file, device, and inter-process I/O,
- iii The ability to initiate asynchronous processes,
- iv System command language selectable on a per-user basis,
- v Over 100 subsystems including a dozen languages,
- vi High degree of portability.

This paper discusses the nature and implementation of the file system and of the user command interface.

**NOTE:** \* Copyright 1974, Association for Computing Machinery, Inc., reprinted by permission. This electronic edition of this paper is a reprint of the version appearing in *The Bell System Technical Journal* 57 no. 6, part 2 (July-August 1978). In turn, that was a revised version of an article that appeared in *Communications of the ACM* 17, No. 7 (July 1974), pp. 365-375. That article was a revised version of a



*You get the idea-*

So what real-world contexts  
warrant virtualizing the  
**ENTIRE** operating system?



external  
security  
threats



development  
messes

# Mutually Untrusted Users





# Mutually Untrusted Users



# Mutually Untrusted Users





# Mutually Untrusted Users



telnet forever!

# Mutually Untrusted Users

login:admin  
pass:love

su  
24/7 ?



**Moron: Using an electric tool, in a pool, with a metal ladder.  
Stupid Moron: Standing barefoot on the ladder while doing it.**

# Mutually Untrusted Users



# Mutually Untrusted Users



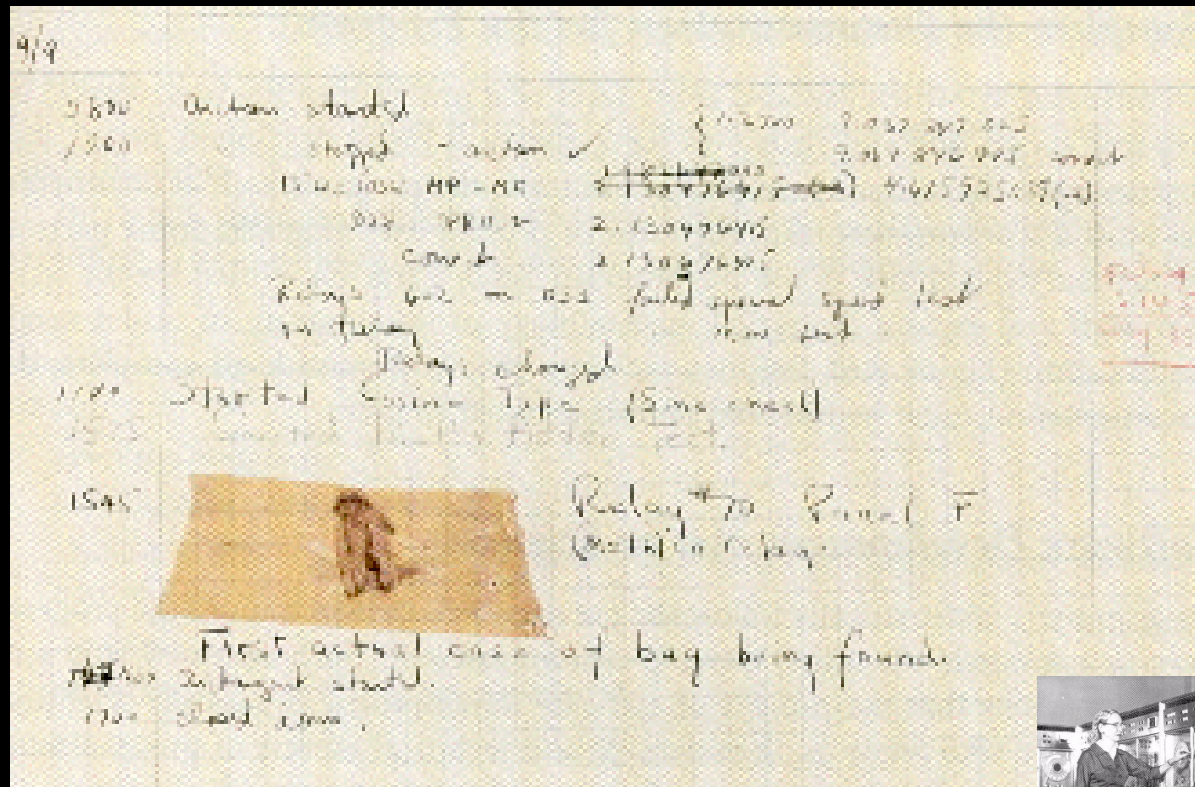


# Mutually Untrusted Users



You run  
\*WHAT\* as  
CGI?

# Mutually Untrusted Users



programs are users too...

# Mutually Untrusted Users



**muscle memory kills!**

SOME ASSHOLE TALKING  
ON HIS CELL PHONE  
GOT CREAMED

TRUCKS

9304






# Harmony.






# UNIX



Once upon a time,  
wasn't UNIX \*fun\*?

# maintaining old junk?



## Rack full of stuff Example:

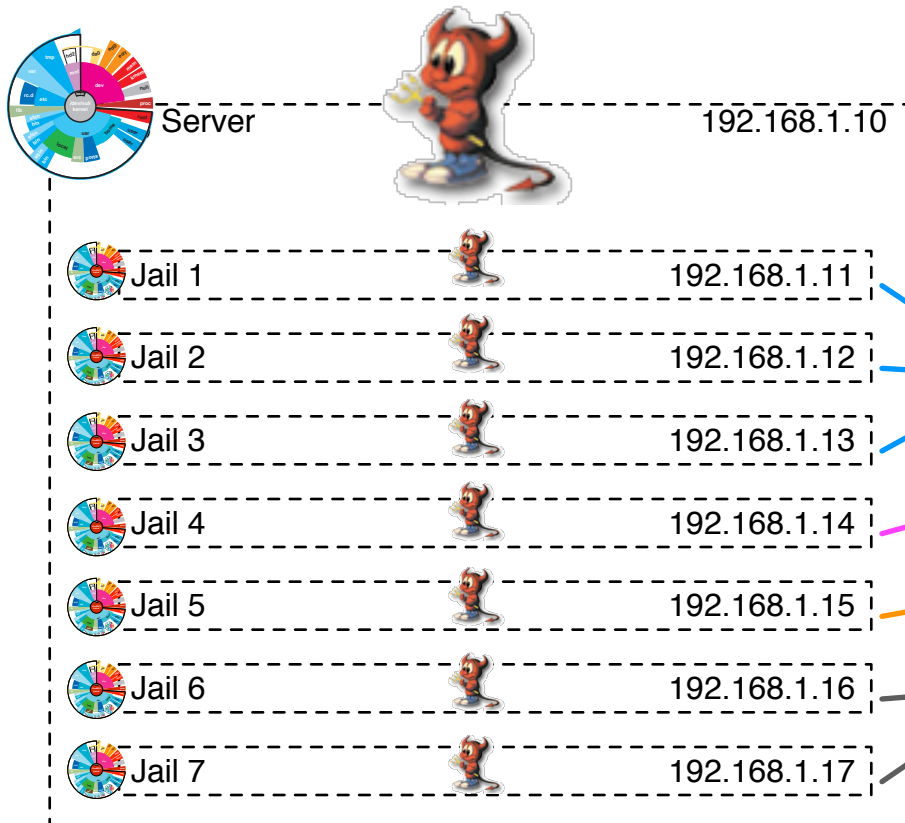
- 3 webservers
- 1 local-use dns cache
- fileserver (for 2 people)
- 2 dev servers





# jail(8)!


**Rack full of stuff ,  
becomes 1u server!**



**3 webservers**

**1 local-use dns cache  
fileserver (for 2 people)**

**2 dev servers**



**jail(8)**

# Definitions

- what is a jail(8):
  - a user space utility, like ifconfig(8)
  - produces a virtual system image
  - process tree based
- what is jail(2):
  - a system call to imprison a process
  - it calls chroot and attaches to IP
  - a very few lines of source code!

# Definitions

- what jail is *not*:
  - it is **not** a classical machine emulator
  - it is **not** chroot ('jail' vocabulary is commonly misused with other \*NIX cultures)

# Great Uses for jail(8)

- hardware resource sharing, an entire OS can be dedicated to a given service
- securely separate untrusted users/processes
- learning/development/testing/hacking
- insane high availability possibilities
- honeypots
- highly vulnerable network services




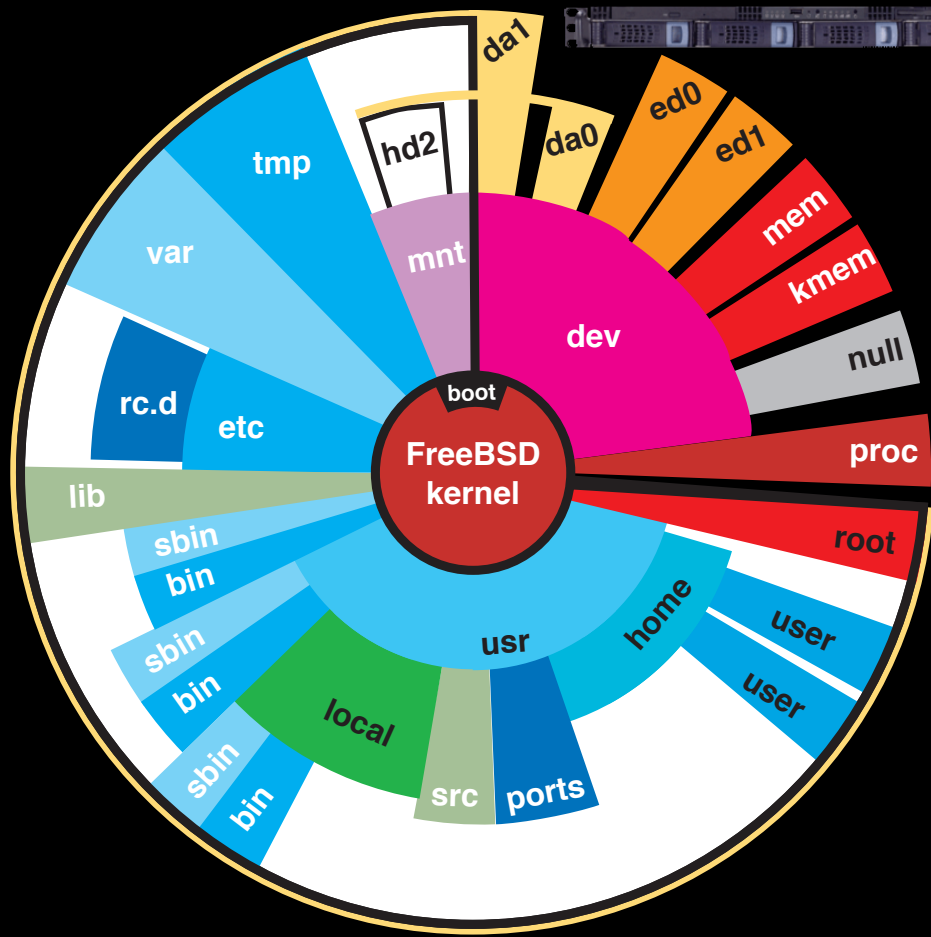
# Poor Uses for jail(8)

- kernel access (you don't get a kernel)
- limited network interface access
- limited device driver access
- when chroot(8) will simply do the job
- some applications require particular low-level system calls:
  - Notably, PostgreSQL doesn't run (securely) in jails based on SysV IPC

# How To jail(8)

- DEFINITIVE instructions in **jail man pages**,
  1. compile a FreeBSD userland from source somewhere on host machine, minor tweaks.
  2. create an IP alias on a network interface
  3. run the jail(8) call with the IP, and userland, to 'boot' the jail, (so to speak).

# Practical Comparison



making a jail

# Host Machine



```
Terminal — ssh — 80x24

o Security advisories and updated errata information for all releases
  at http://www.FreeBSD.org/releases/ - always consult the ERRATA
  for your release first as it's updated frequently.

o The Handbook and FAQ documents are at http://www.FreeBSD.org/ and,
  along with the mailing lists, can be searched by going to
  http://www.FreeBSD.org/search/. If the doc distribution has
  been installed, they're also available formatted in /usr/share/doc.

If you still have a question or problem, please take the output of
`uname -a`, along with any relevant error messages, and email it
as a question to the questions@FreeBSD.org mailing list. If you are
unfamiliar with FreeBSD's directory layout, please refer to the hier(7)
manual page. If you are not familiar with manual pages, type `man man`.

You may also use sysinstall(8) to re-enter the installation and
configuration utility. Edit /etc/motd to change this login announcement.

chicken:/home/ike ike$ uname -a
FreeBSD chicken.diversaform.net 6.1-RELEASE-p3 FreeBSD 6.1-RELEASE-p3 #3: Fri Jul 14 11:43:08 EDT 2006
ERIC i386
chicken:/home/ike ike$
```

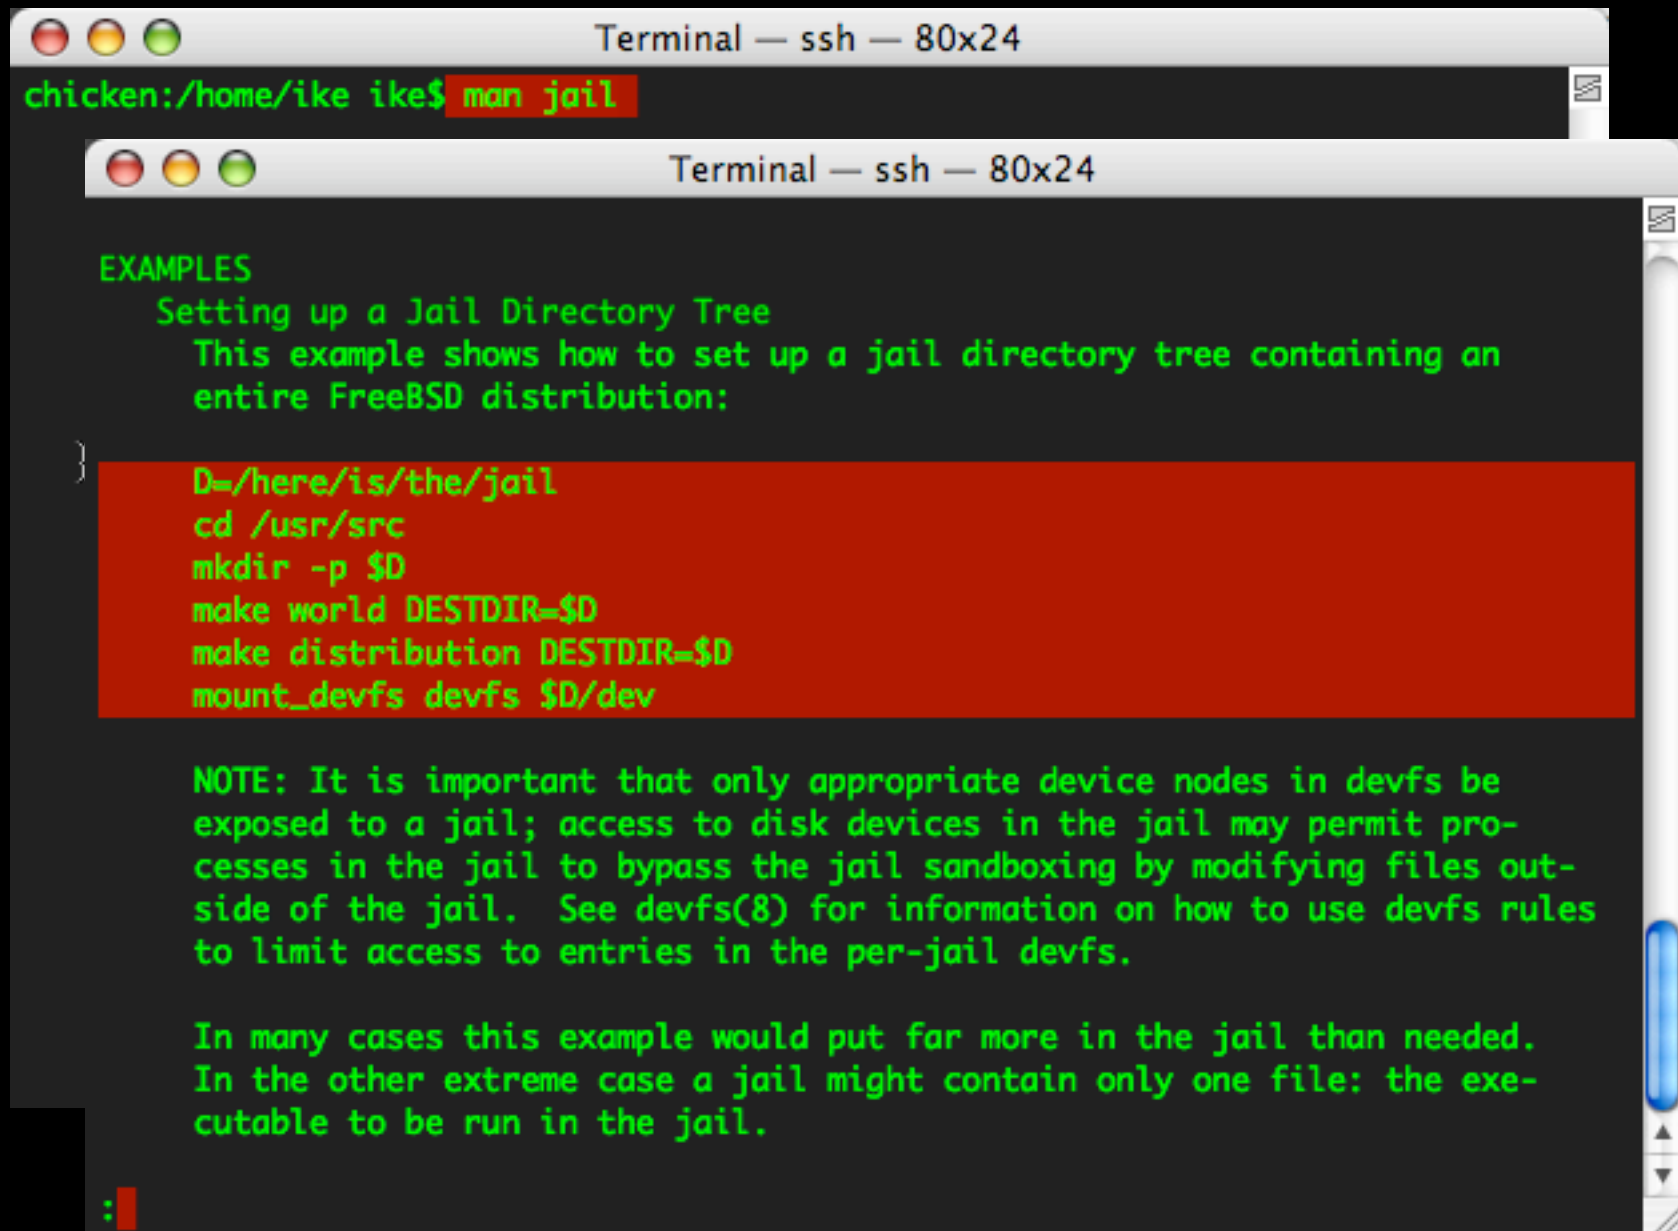


# preflight (simple)

1. get source to build with (cvsup is great)
2. make somewhere for the jails to live (partitions, disk mounts, etc...)
3. make somewhere for jail-related start/mgmt scripts to live

(starting jails from `/etc/rc.d/jail` can **thrash violently** in most contexts! Bad!)

# preflight- (man, definitive)



The image shows two overlapping terminal windows. The top window shows a user at a shell prompt typing 'man jail'. The bottom window displays the man page for 'jail', which includes an 'EXAMPLES' section with instructions on setting up a jail directory tree, a list of commands, a 'NOTE' about device nodes, and a concluding paragraph.

```
Terminal — ssh — 80x24
chicken:/home/ike ike$ man jail

Terminal — ssh — 80x24

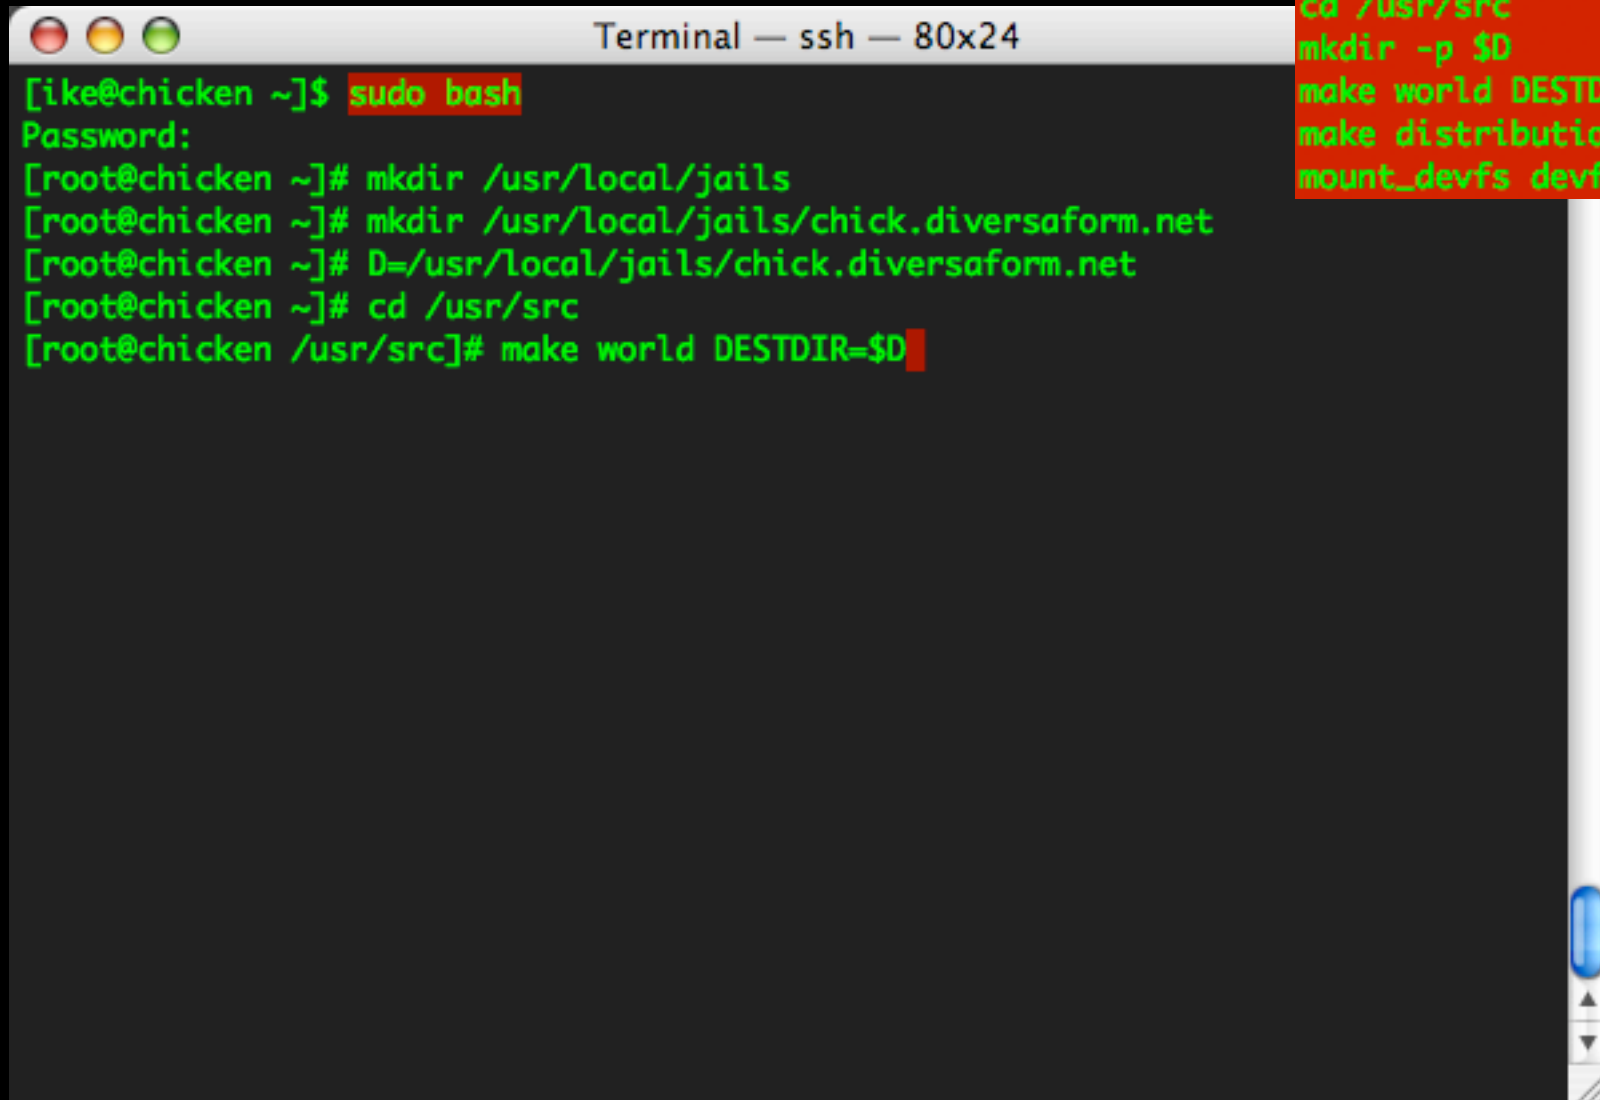
EXAMPLES
Setting up a Jail Directory Tree
This example shows how to set up a jail directory tree containing an
entire FreeBSD distribution:
}
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev

NOTE: It is important that only appropriate device nodes in devfs be
exposed to a jail; access to disk devices in the jail may permit pro-
cesses in the jail to bypass the jail sandboxing by modifying files out-
side of the jail. See devfs(8) for information on how to use devfs rules
to limit access to entries in the per-jail devfs.

In many cases this example would put far more in the jail than needed.
In the other extreme case a jail might contain only one file: the exe-
cutable to be run in the jail.
```



# preflight- (build from src)



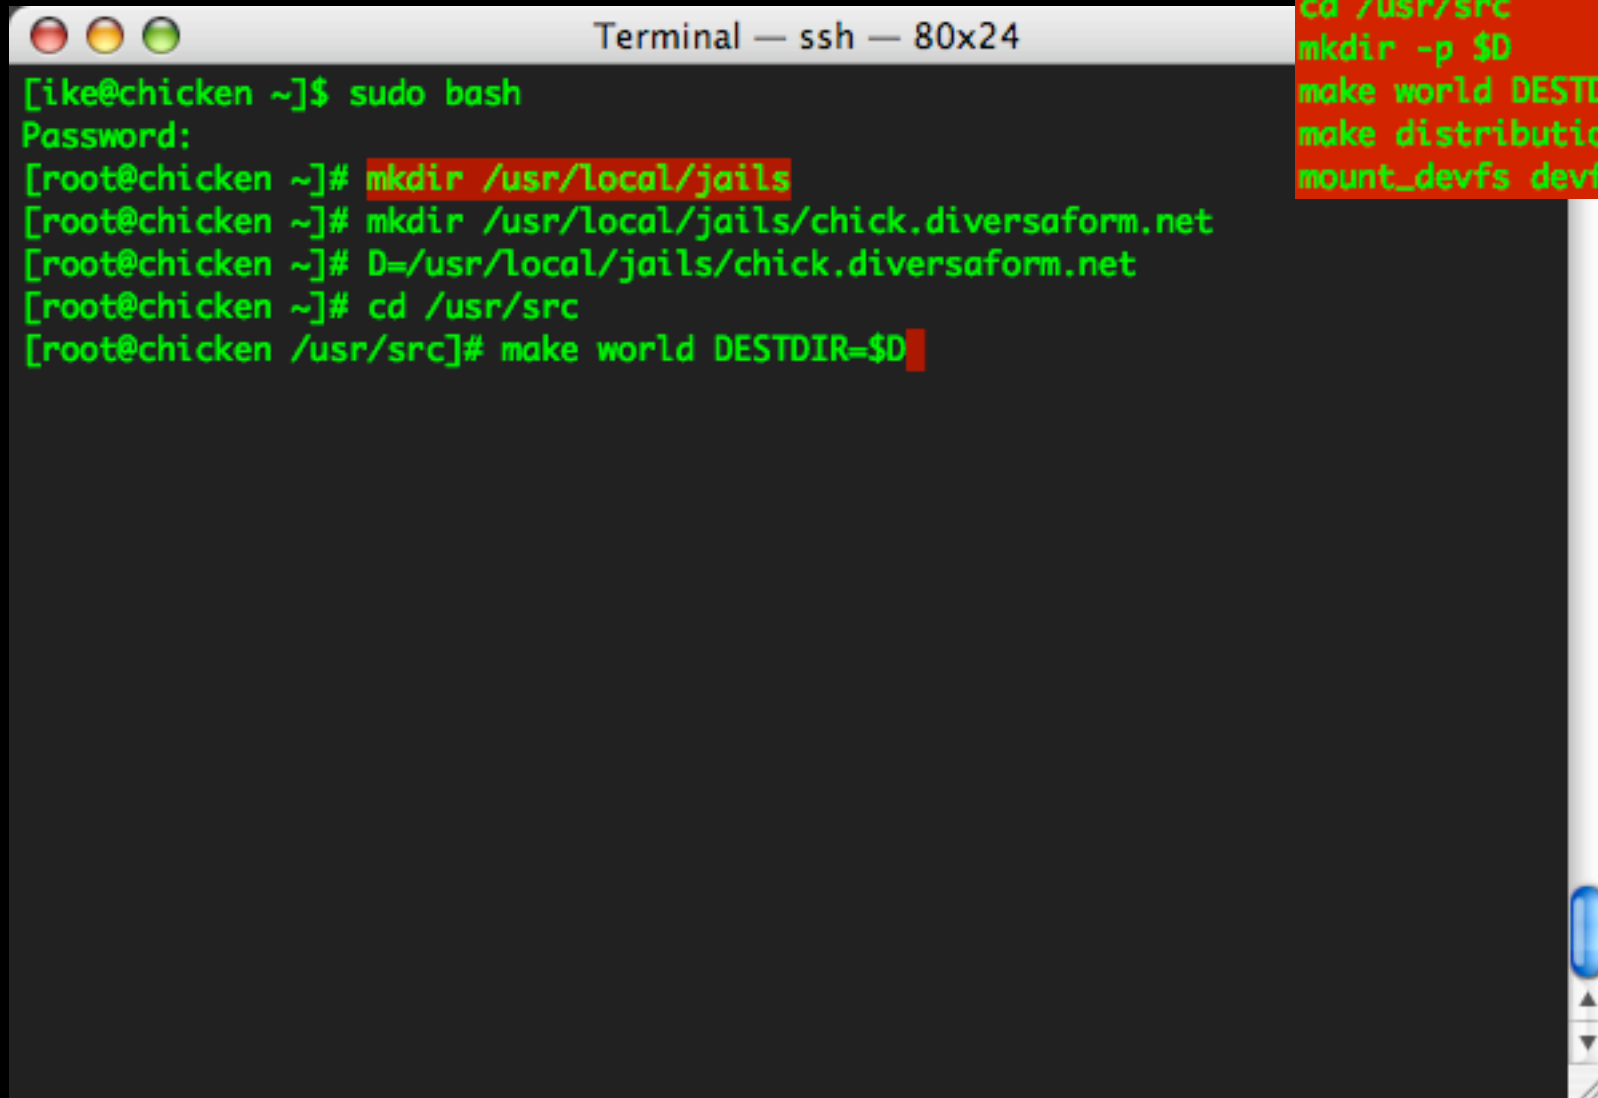
The image shows a terminal window titled "Terminal — ssh — 80x24". The user "ike" is logged in as "root" on a host named "chicken". The terminal shows the following commands and output:

```
[ike@chicken ~]$ sudo bash
Password:
[root@chicken ~]# mkdir /usr/local/jails
[root@chicken ~]# mkdir /usr/local/jails/chick.diversaform.net
[root@chicken ~]# D=/usr/local/jails/chick.diversaform.net
[root@chicken ~]# cd /usr/src
[root@chicken /usr/src]# make world DESTDIR=$D
```

To the right of the terminal window, there is a red box containing the following commands:

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

# preflight- (build from src)

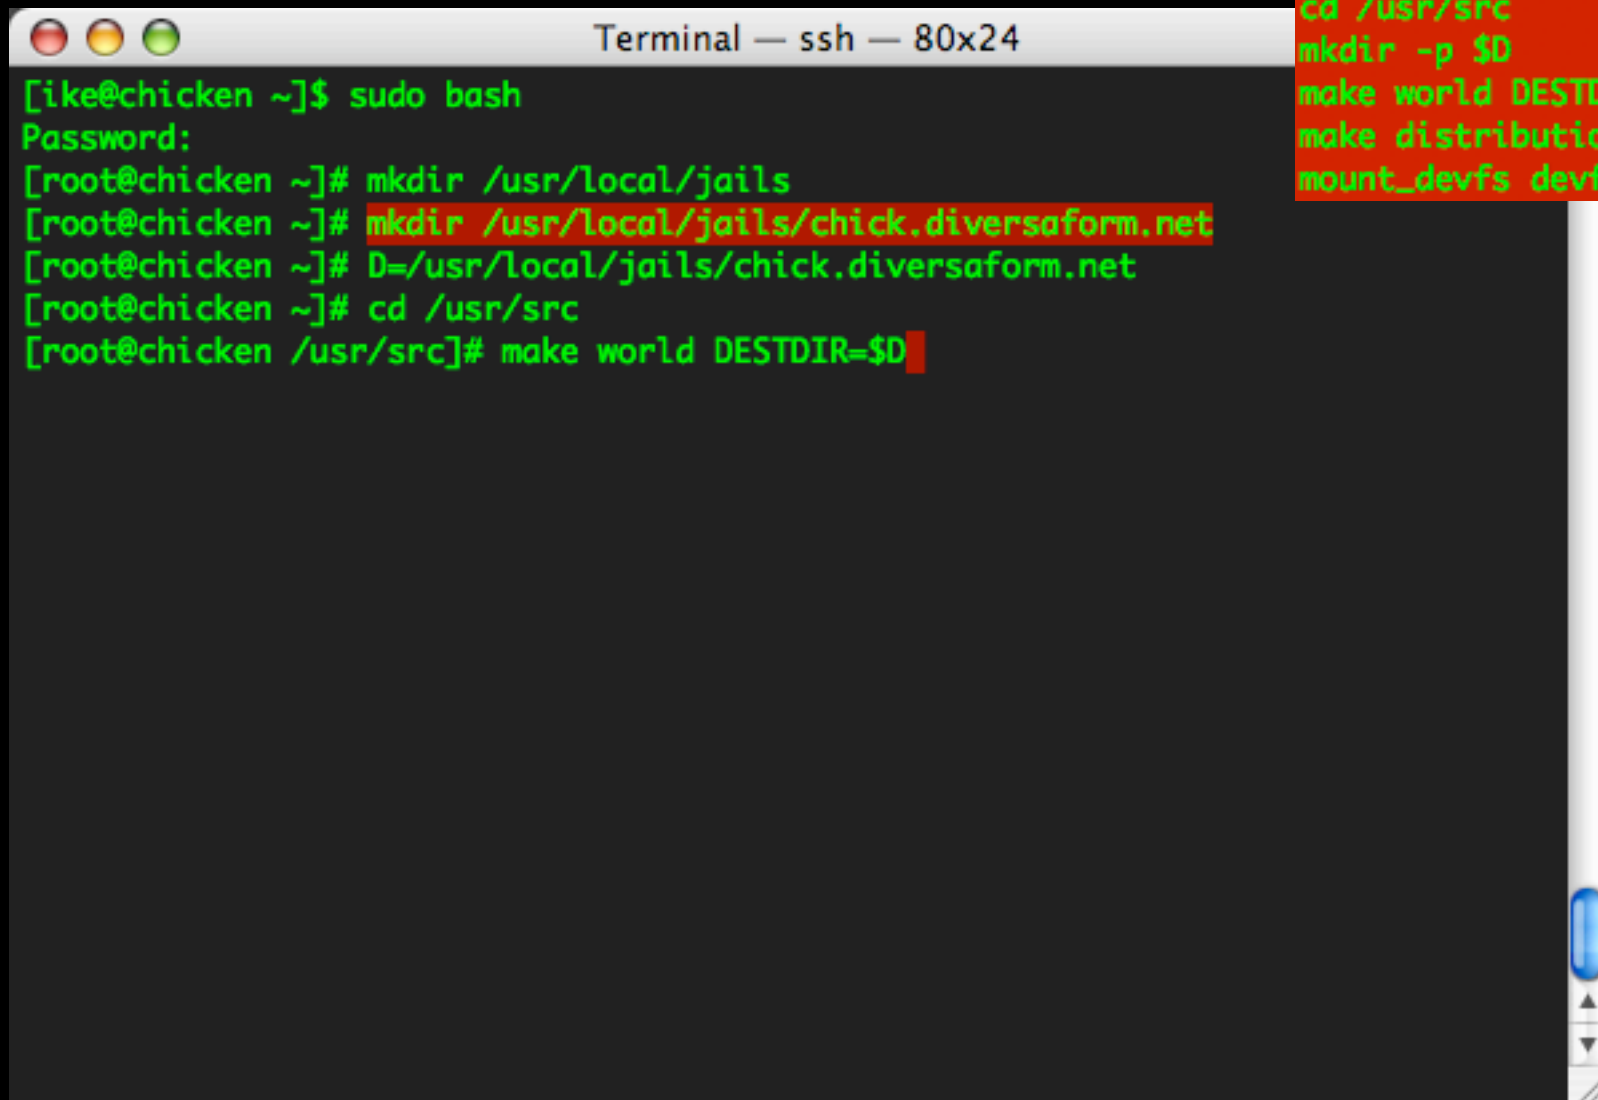


A terminal window titled "Terminal — ssh — 80x24" showing the execution of a preflight build process. The user starts as [ike@chicken ~] and runs 'sudo bash'. After entering the password, they become root and execute a series of commands: 'mkdir /usr/local/jails', 'mkdir /usr/local/jails/chick.diversaform.net', 'D=/usr/local/jails/chick.diversaform.net', 'cd /usr/src', and 'make world DESTDIR=\$D'. A red box on the right side of the terminal contains the same commands, with the first line 'D=/here/is/the/jail' being different from the terminal output.

```
[ike@chicken ~]$ sudo bash
Password:
[root@chicken ~]# mkdir /usr/local/jails
[root@chicken ~]# mkdir /usr/local/jails/chick.diversaform.net
[root@chicken ~]# D=/usr/local/jails/chick.diversaform.net
[root@chicken ~]# cd /usr/src
[root@chicken /usr/src]# make world DESTDIR=$D
```

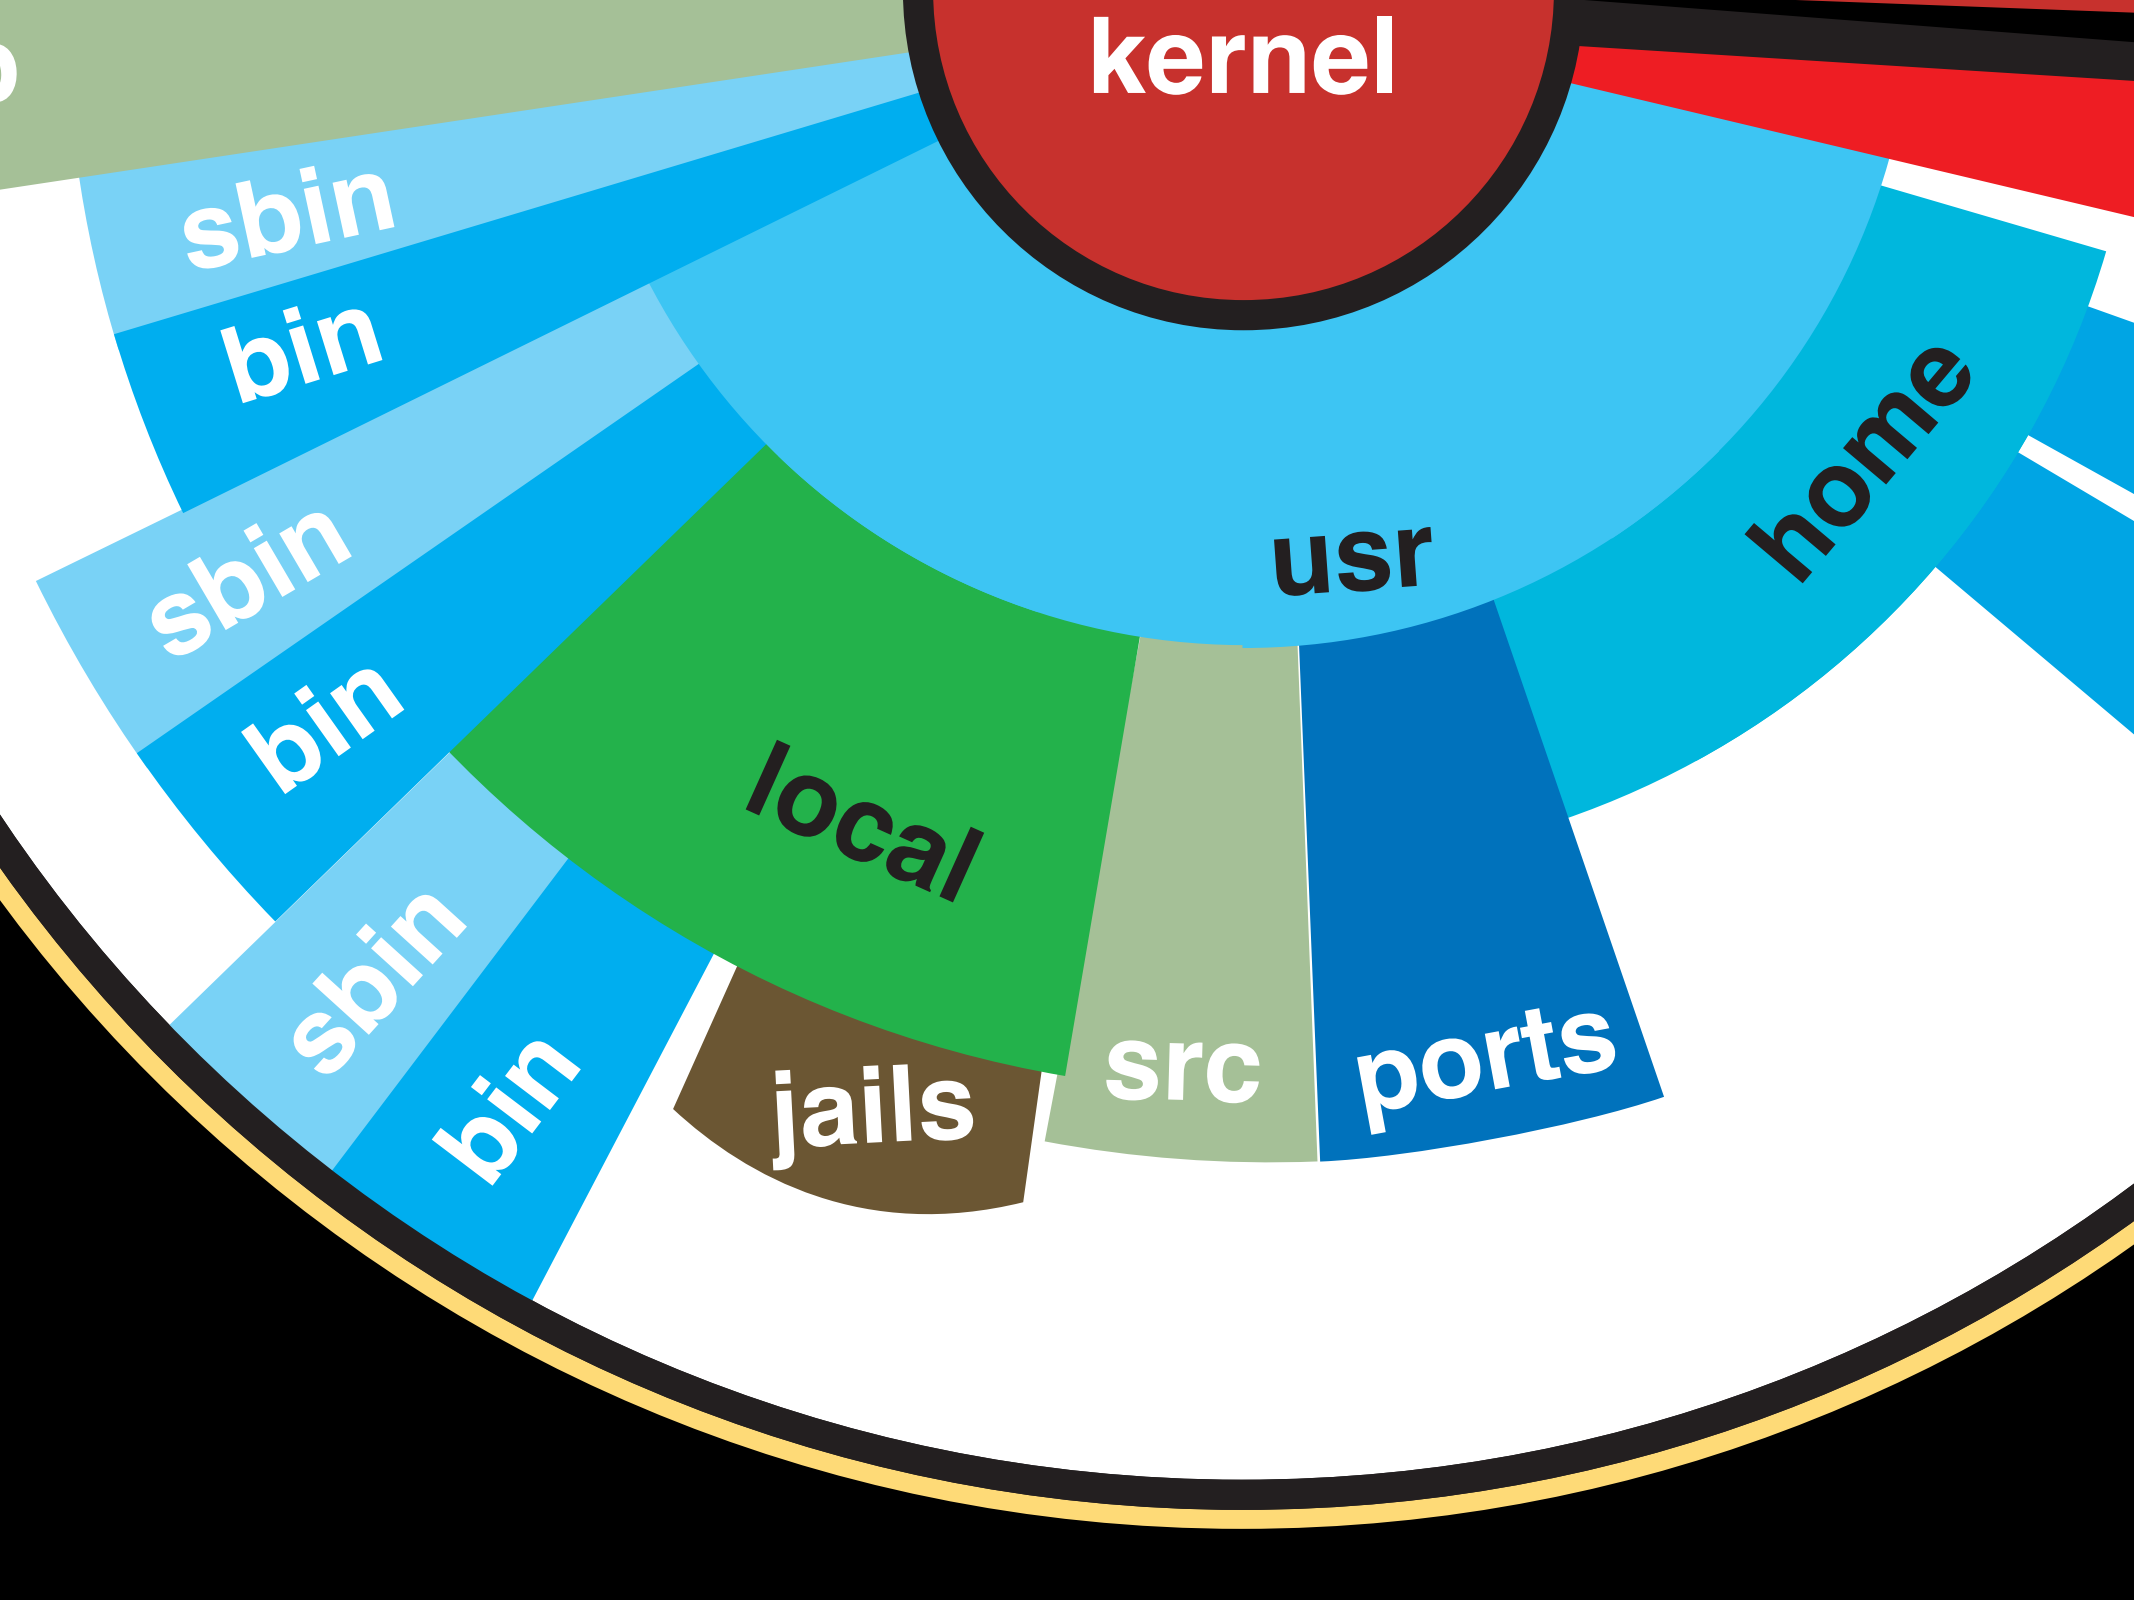
```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

# preflight- (build from src)



```
Terminal — ssh — 80x24
[ike@chicken ~]$ sudo bash
Password:
[root@chicken ~]# mkdir /usr/local/jails
[root@chicken ~]# mkdir /usr/local/jails/chick.diversaform.net
[root@chicken ~]# D=/usr/local/jails/chick.diversaform.net
[root@chicken ~]# cd /usr/src
[root@chicken /usr/src]# make world DESTDIR=$D
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```



**kernel**

**home**

**usr**

**ports**

**src**

**local**

**jails**

**bin**

**sbin**

**bin**

**sbin**

**bin**

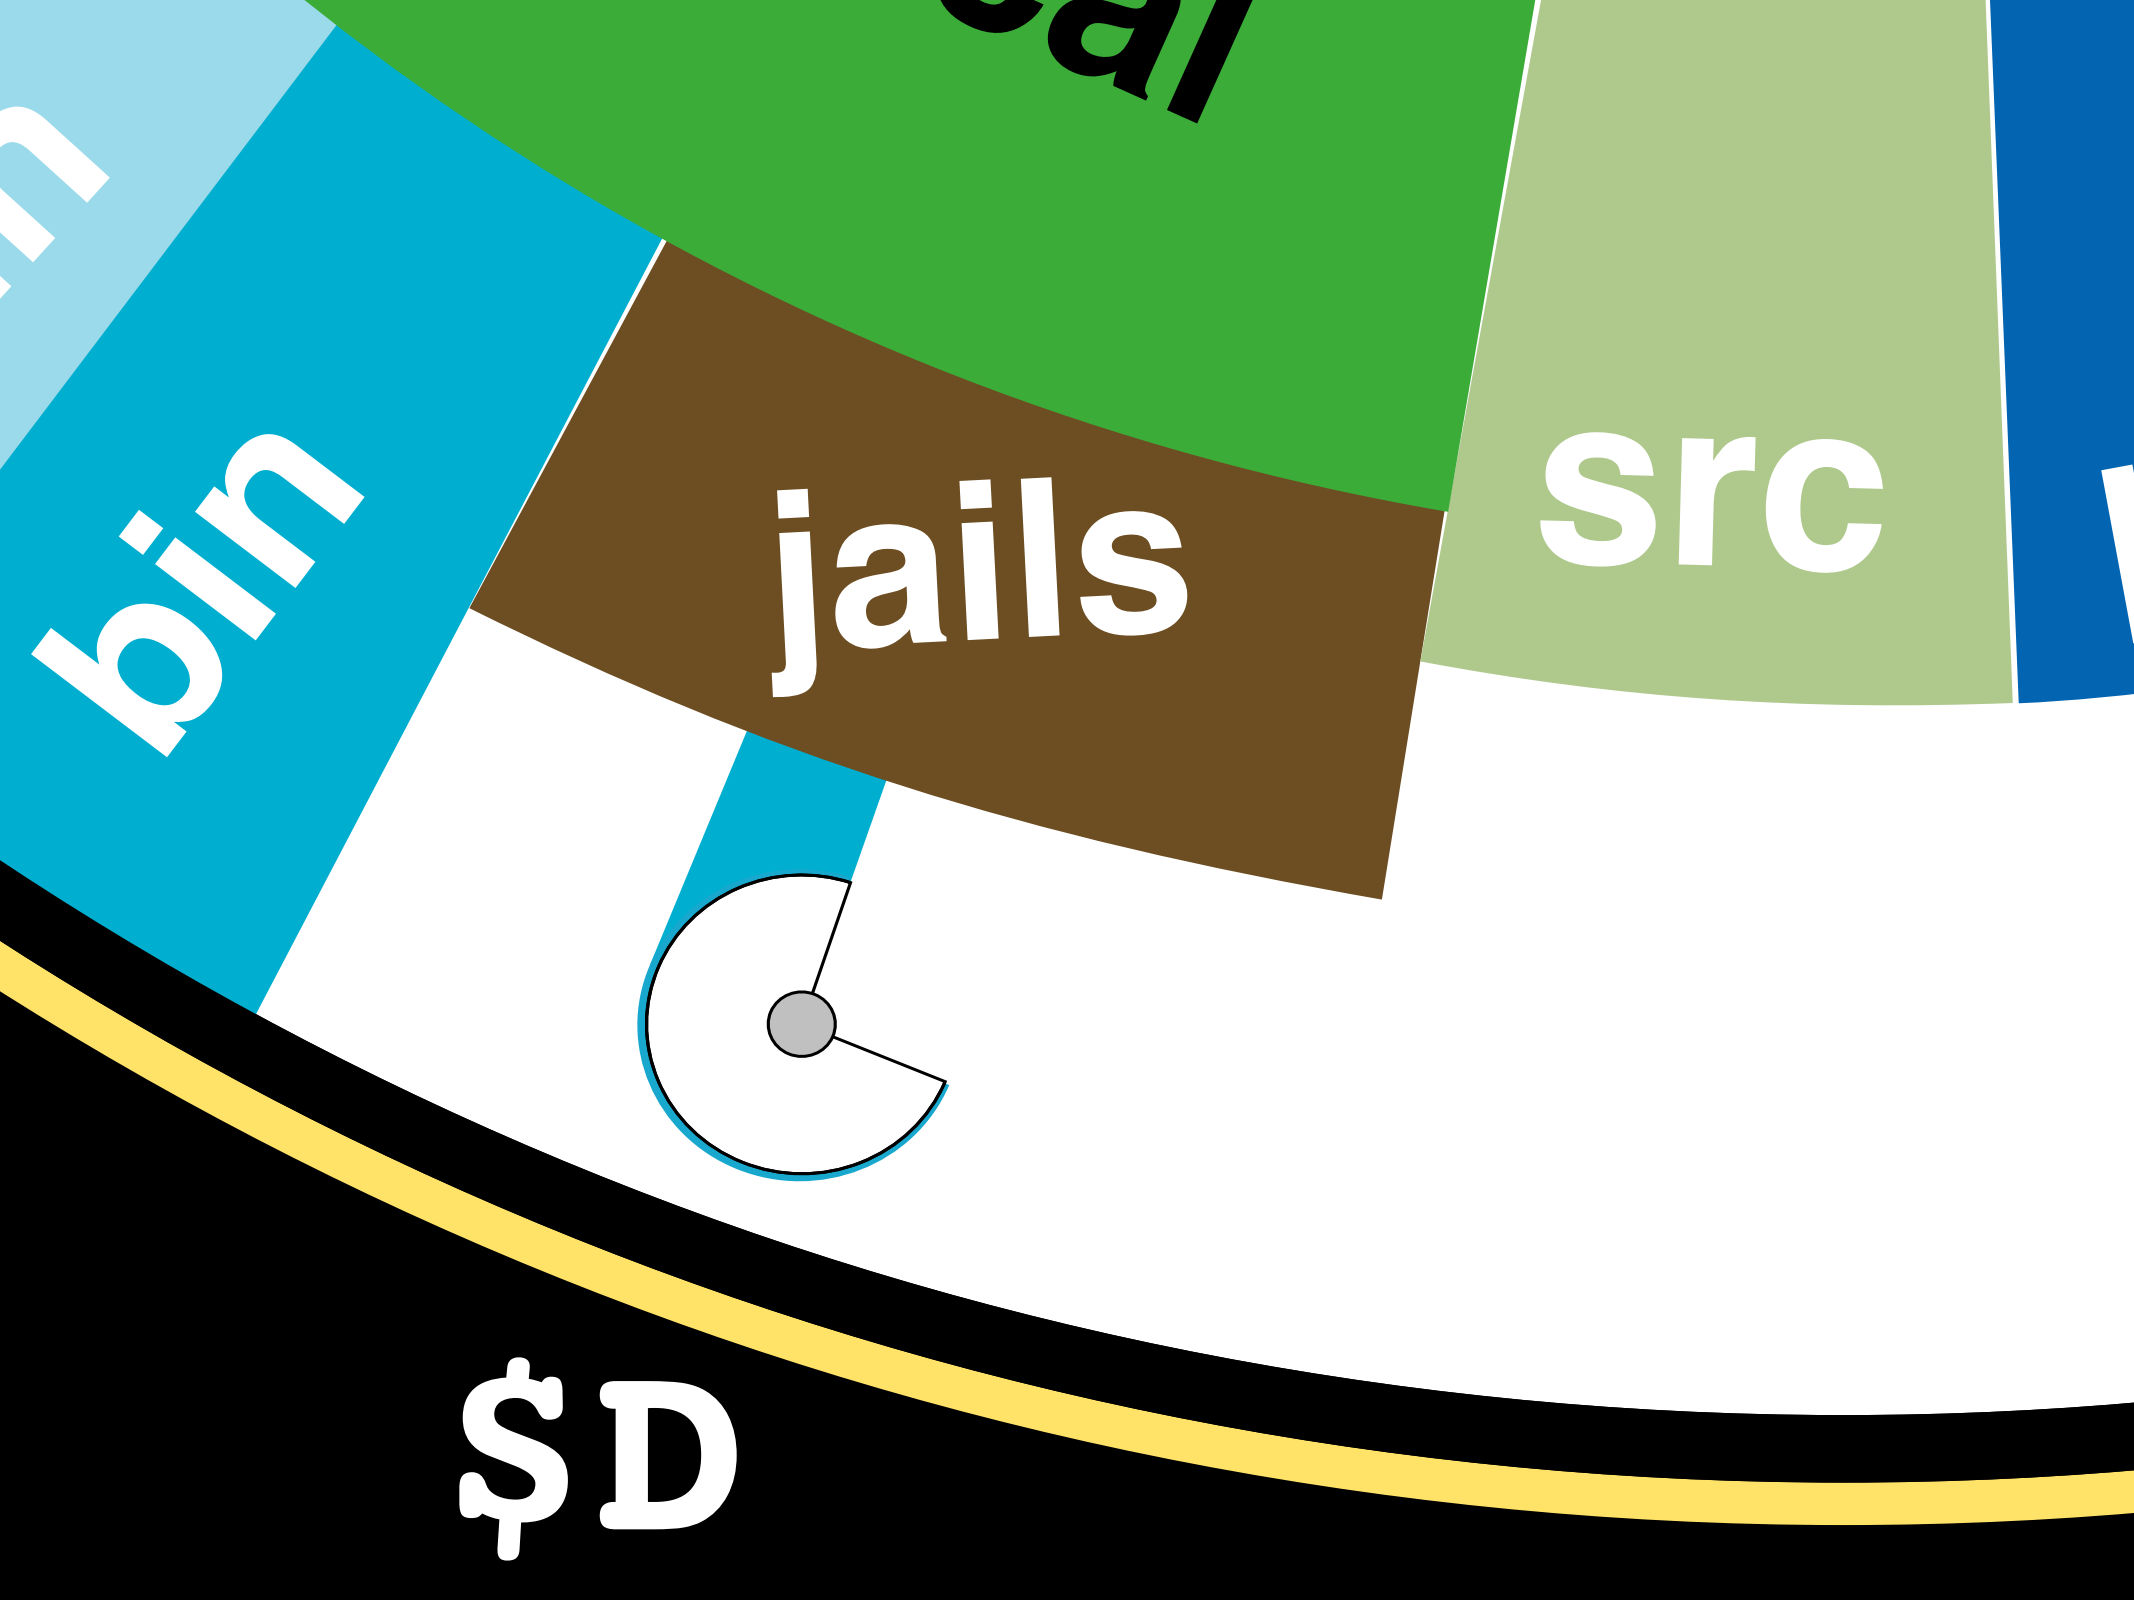
**sbin**

# preflight- (build from src)

```
Terminal — ssh — 80x24
[ike@chicken ~]$ sudo bash
Password:
[root@chicken ~]# mkdir /usr/local/jails
[root@chicken ~]# mkdir /usr/local/jails/chick.diversaform.net
[root@chicken ~]# D=/usr/local/jails/chick.diversaform.net
[root@chicken ~]# cd /usr/src
[root@chicken /usr/src]# make world DESTDIR=$D
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```





bin

jails

src

cal

\$D

# preflight- (build from src)

```
Terminal — ssh — 80x24
[ike@chicken ~]$ sudo bash
Password:
[root@chicken ~]# mkdir /usr/local/jails
[root@chicken ~]# mkdir /usr/local/jails/chick.diversaform.net
[root@chicken ~]# D=/usr/local/jails/chick.diversaform.net
[root@chicken ~]# cd /usr/src
[root@chicken /usr/src]# make world DESTDIR=$D
-----
>>> make world started on Mon Jul 10 14:39:21 EDT 2006
-----
>>> Rebuilding the temporary build tree
-----
rm -rf /usr/obj/usr/src/tmp
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/bin
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/games
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/include/c++/3.4
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/include/sys
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/lib
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/libexec
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/sbin
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/share/dict
mkdir -p /usr/obj/usr/src/tmp/legacy/usr/share/groff_font/devX100
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

compile!



# preflight- (build from src)


```
Terminal — ssh — 80x24
-----
[root@chicken /usr/src]# make distribution DESTDIR=$D
cd /usr/src/etc; MAKEOBJDIRPREFIX=/usr/obj MACHINE_ARCH=i386 MACmount_devfs devfs $D/dev
TYPE= GROFF_BIN_PATH=/usr/obj/usr/src/tmp/legacy/usr/bin GROFF_FONT_PATH=/usr/
obj/usr/src/tmp/legacy/usr/share/groff_font GROFF_TMAC_PATH=/usr/obj/usr/src/tm
p/legacy/usr/share/tmac PATH=/usr/obj/usr/src/tmp/legacy/usr/sbin:/usr/obj/usr/s
rc/tmp/legacy/usr/bin:/usr/obj/usr/src/tmp/legacy/usr/games:/usr/obj/usr/src/tmp
/usr/sbin:/usr/obj/usr/src/tmp/usr/bin:/usr/obj/usr/src/tmp/usr/games:/sbin:/bin
:/usr/sbin:/usr/bin /usr/obj/usr/src/make.i386/make distribution
cd /usr/src/etc; install -o root -g wheel -m 644 amd.map apmd.conf auth.conf
crontab csh.cshrc csh.login csh.logout devd.conf devfs.conf dhclient.conf diskd
ab fbtabs ftpusers gettytab group hosts hosts.allow hosts.equiv hosts.lpd inetd
.conf login.access login.conf mac.conf motd netconfig network.subr
syslog.conf portsnap.conf pf.conf pf.os phones profile protocols
nded rc.firewall rc.firewall6 rc.initdiskless rc.sendmail rc.shutdown
remote rpc services shells snmpd.config sysctl.conf syslog.conf us
i386/ttys /usr/src/etc/./gnu/usr.bin/man/manpath/manpath.config /
./usr.bin/mail/misc/mail.rc /usr/src/etc/./usr.bin/locate/locate
intcap /usr/local/jails/chick.diversaform.net/etc; cap_mkdb -l /usr
/chick.diversaform.net/etc/login.conf; install -o root -g wheel -m
t pccard_ether rc.suspend rc.resume /usr/local/jails/chick.diversafo
install -o root -g wheel -m 600 master.passwd nsmb.conf opieaccess
jails/chick.diversaform.net/etc; pwd_mkdb -L -i -p -d /usr/local/ja
verssaform.net/etc /usr/local/jails/chick.diversaform.net/etc/master
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
```


compile!



# preflight



# preflight





# preflight- (mount /dev)

```
Terminal — ssh — 80x24
chick.diversaform.net/etc/ppp
cd /usr/src/etc/mail; install -o root -g wheel -m 644 Makefile RE
nf access.sample virtusertable.sample mailertable.sample aliases
ls/chick.diversaform.net/etc/mail
+ ln -s mail/aliases /usr/local/jails/chick.diversaform.net/etc/aliases
install -o root -g operator -m 664 /dev/null /usr/local/jails/chick.diversaform
.net/etc/dumpdates
install -o nobody -g wheel -m 644 /dev/null /usr/local/jails/chick.diversaform.
net/var/db/locate.database
install -o root -g wheel -m 644 /usr/src/etc/minfree /usr/local/jails/chick.div
ersaform.net/var/crash
cd /usr/src/etc/..; install -o root -g wheel -m 444 COPYRIGHT /usr/local/jails/
chick.diversaform.net/
install -o root -g wheel -m 444 /usr/src/etc/./sys/i386/conf/GENERIC.hints /u
sr/local/jails/chick.diversaform.net/boot/device.hints
[root@chicken /usr/src]# mount_devfs devfs $D/dev
[root@chicken /usr/src]# cd $D
[root@chicken /usr/local/jails/chick.diversaform.net]# ln -s ./dev/null kernel
[root@chicken /usr/local/jails/chick.diversaform.net]# ls
.cshrc          boot           lib            rescue        tmp
.profile       dev           libexec       root          usr
COPYRIGHT      etc           mnt           sbin         var
bin            kernel        proc          sys
[root@chicken /usr/local/jails/chick.diversaform.net]#
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

# preflight- (mount /dev)

```
Terminal — ssh — 80x24
chick.diversaform.net/etc/ppp
cd /usr/src/etc/mail; install -o root -g wheel -m 644 Makefile RE
nf access.sample virtusertable.sample mailertable.sample aliases
ls/chick.diversaform.net/etc/mail
+ ln -s mail/aliases /usr/local/jails/chick.diversaform.net/etc/aliases
install -o root -g operator -m 664 /dev/null /usr/local/jails/chick.diversaform
.net/etc/dumpdates
install -o nobody -g wheel -m 644 /dev/null /usr/local/jails/chick.diversaform.
net/var/db/locate.database
install -o root -g wheel -m 644 /usr/src/etc/minfree /usr/local/jails/chick.div
ersaform.net/var/crash
cd /usr/src/etc/..; install -o root -g wheel -m 444 COPYRIGHT /usr/local/jails/
chick.diversaform.net/
install -o root -g wheel -m 444 /usr/src/etc/./sys/i386/conf/GENERIC.hints /u
sr/local/jails/chick.diversaform.net/boot/device.hints
[root@chicken /usr/src]# mount_devfs devfs $D/dev
[root@chicken /usr/src]# cd $D
[root@chicken /usr/local/jails/chick.diversaform.net]# ln -s ./dev/null kernel
[root@chicken /usr/local/jails/chick.diversaform.net]# ls
.cshrc          boot           lib            rescue        tmp
.profile       dev           libexec       root          usr
COPYRIGHT      etc          mnt           sbin         var
bin            kernel        proc          sys
[root@chicken /usr/local/jails/chick.diversaform.net]#
```

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

# preflight- (null kernel)

```
Terminal — ssh — 80x24
chick.diversaform.net/etc/ppp
cd /usr/src/etc/mail; install -o root -g wheel -m 644 Makefile README mailer.co
nf access.sample virtusertable.sample mailertable.sample aliases /usr/local/jai
ls/chick.diversaform.net/etc/mail
+ ln -s mail/aliases /usr/local/jails/chick.diversaform.net/etc/aliases
install -o root -g operator -m 664 /dev/null /usr/local/jails/chick.diversaform
.net/etc/dumpdates
install -o nobody -g wheel -m 644 /dev/null /usr/local/jails/chick.diversaform.
net/var/db/locate.database
install -o root -g wheel -m 644 /usr/src/etc/minfree /usr/local/jails/chick.div
ersaform.net/var/crash
cd /usr/src/etc/..; install -o root -g wheel -m 444 COPYRIGHT /usr/local/jails/
chick.diversaform.net/
install -o root -g wheel -m 444 /usr/src/etc/./sys/i386/conf/GENERIC.hints /u
sr/local/jails/chick.diversaform.net/boot/device.hints
[root@chicken /usr/src]# mount_devfs devfs $D/dev
[root@chicken /usr/src]# cd $D
[root@chicken /usr/local/jails/chick.diversaform.net]# ln -s ./dev/null kernel
[root@chicken /usr/local/jails/chick.diversaform.net]# ls
.cshrc          boot           lib            rescue        tmp
.profile       dev           libexec       root          usr
COPYRIGHT      etc           mnt           sbin         var
bin            kernel        proc          sys
[root@chicken /usr/local/jails/chick.diversaform.net]#
```

D=/here/is/the/jail  
cd /usr/src  
DIR=\$D  
n DESTDIR=\$D  
s \$D/dev


# preflight- (null kernel)

```
Terminal — ssh — 80x24
nf access.sample virtusertable.sample mailertable.sample aliases /usr/local/jails/chick.diversaform.net/etc/mail
+ ln -s mail/aliases /usr/local/jails/chick.diversaform.net/etc/aliases
install -o root -g operator -m 664 /dev/null /usr/local/jails/chick.diversaform.net/etc/dumpdates
install -o nobody -g wheel -m 644 /dev/null /usr/local/jails/chick.diversaform.net/var/db/locate.database
install -o root -g wheel -m 644 /usr/src/etc/minfree /usr/local/jails/chick.diversaform.net/var/crash
cd /usr/src/etc/..; install -o root -g wheel -m 444 COPYRIGHT /usr/local/jails/chick.diversaform.net/
install -o root -g wheel -m 444 /usr/src/etc/../../sys/i386/conf/GENERIC.hints /usr/local/jails/chick.diversaform.net/boot/device.hints
[root@chicken /usr/src]# mount_devfs devfs $D/dev
[root@chicken /usr/src]# cd $D
[root@chicken /usr/local/jails/chick.diversaform.net]# ln -s ../dev/null kernel
[root@chicken /usr/local/jails/chick.diversaform.net]# ls
.cshrc          boot           lib            rescue         tmp
.profile       dev            libexec       root           usr
COPYRIGHT      etc           mnt           sbin          var
bin            kernel        proc          sys
[root@chicken /usr/local/jails/chick.diversaform.net]# ls -lah kernel
lrwxr-xr-x  1 root  wheel   10B Jul 07 15:36 kernel -> ../dev/null
[root@chicken /usr/local/jails/chick.diversaform.net]#
```

D=/here/is/the/jail  
cd /usr/src


IR=\$D  
n DESTDIR=\$D  
s \$D/dev

# preflight





# preflight



# preflight

```
D=/here/is/the/jail
cd /usr/src
mkdir -p $D
make world DESTDIR=$D
make distribution DESTDIR=$D
mount_devfs devfs $D/dev
```

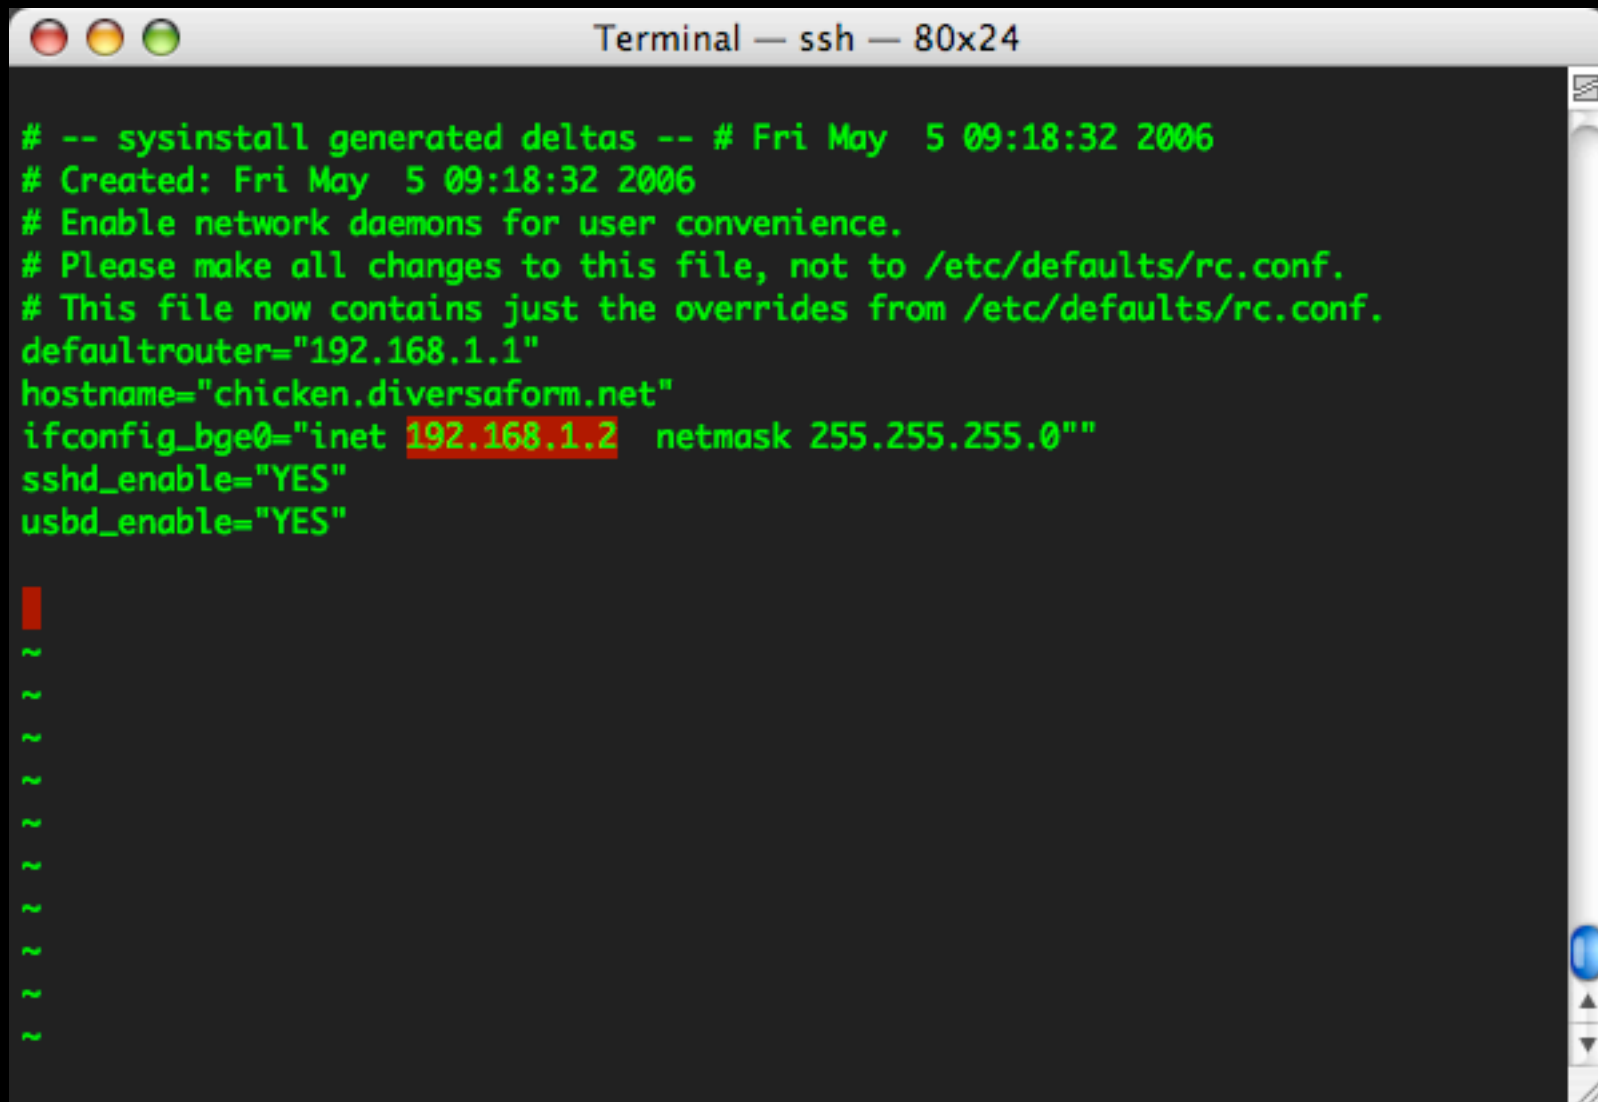
## Common Question:

Why isn't there an automated build system for this stage?

- Take care with the build procedure, it's better to automate things later, once you have basics setup.

(network, users, packages, time, etc.)

# preflight- (config host)



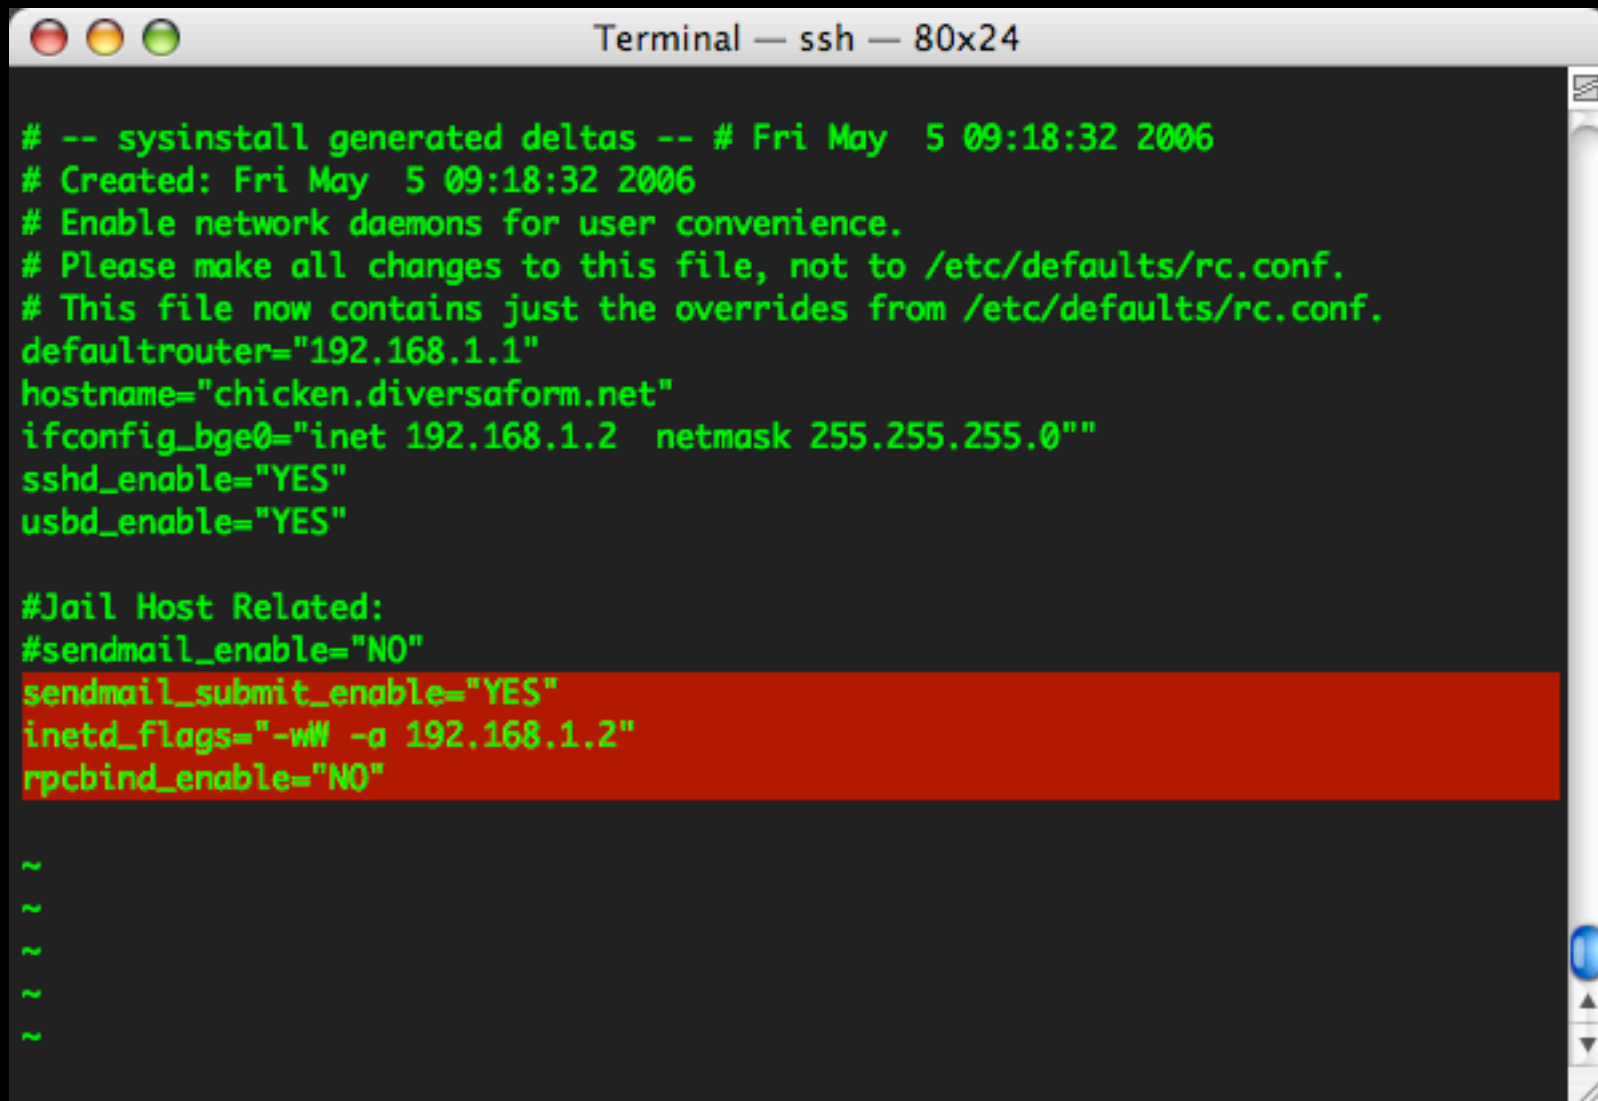
```
Terminal — ssh — 80x24

# -- sysinstall generated deltas -- # Fri May  5 09:18:32 2006
# Created: Fri May  5 09:18:32 2006
# Enable network daemons for user convenience.
# Please make all changes to this file, not to /etc/defaults/rc.conf.
# This file now contains just the overrides from /etc/defaults/rc.conf.
defaultrouter="192.168.1.1"
hostname="chicken.diversaform.net"
ifconfig_bge0="inet 192.168.1.2 netmask 255.255.255.0"
sshd_enable="YES"
usbd_enable="YES"

~
~
~
~
~
~
~
~
~
~
~
```

jailinghost:/etc/rc.conf (stock)

# preflight- (config host)



```
Terminal — ssh — 80x24

# -- sysinstall generated deltas -- # Fri May  5 09:18:32 2006
# Created: Fri May  5 09:18:32 2006
# Enable network daemons for user convenience.
# Please make all changes to this file, not to /etc/defaults/rc.conf.
# This file now contains just the overrides from /etc/defaults/rc.conf.
defaultrouter="192.168.1.1"
hostname="chicken.diversaform.net"
ifconfig_bge0="inet 192.168.1.2  netmask 255.255.255.0"
sshd_enable="YES"
usbd_enable="YES"

#Jail Host Related:
#sendmail_enable="NO"
sendmail_submit_enable="YES"
inetd_flags="-wW -a 192.168.1.2"
rpcbind_enable="NO"

~
~
~
~
~
```

jailinghost:/etc/rc.conf

# preflight- (master system)

A terminal window titled "Terminal — ssh — 80x24" displays the configuration file /etc/ssh/sshd\_config. The text is green on a dark background. The configuration includes comments about the strategy for options, FreeBSD-specific defaults, and the version addendum. The ListenAddress is set to 192.168.1.2, and the HostKey files are specified.

```
Terminal — ssh — 80x24

# The strategy used for options in the default sshd_config shipped with
# OpenSSH is to specify options with their default value where
# possible, but leave them commented. Uncommented options change a
# default value.

# Note that some of FreeBSD's defaults differ from OpenBSD's, and
# FreeBSD has a few additional options.

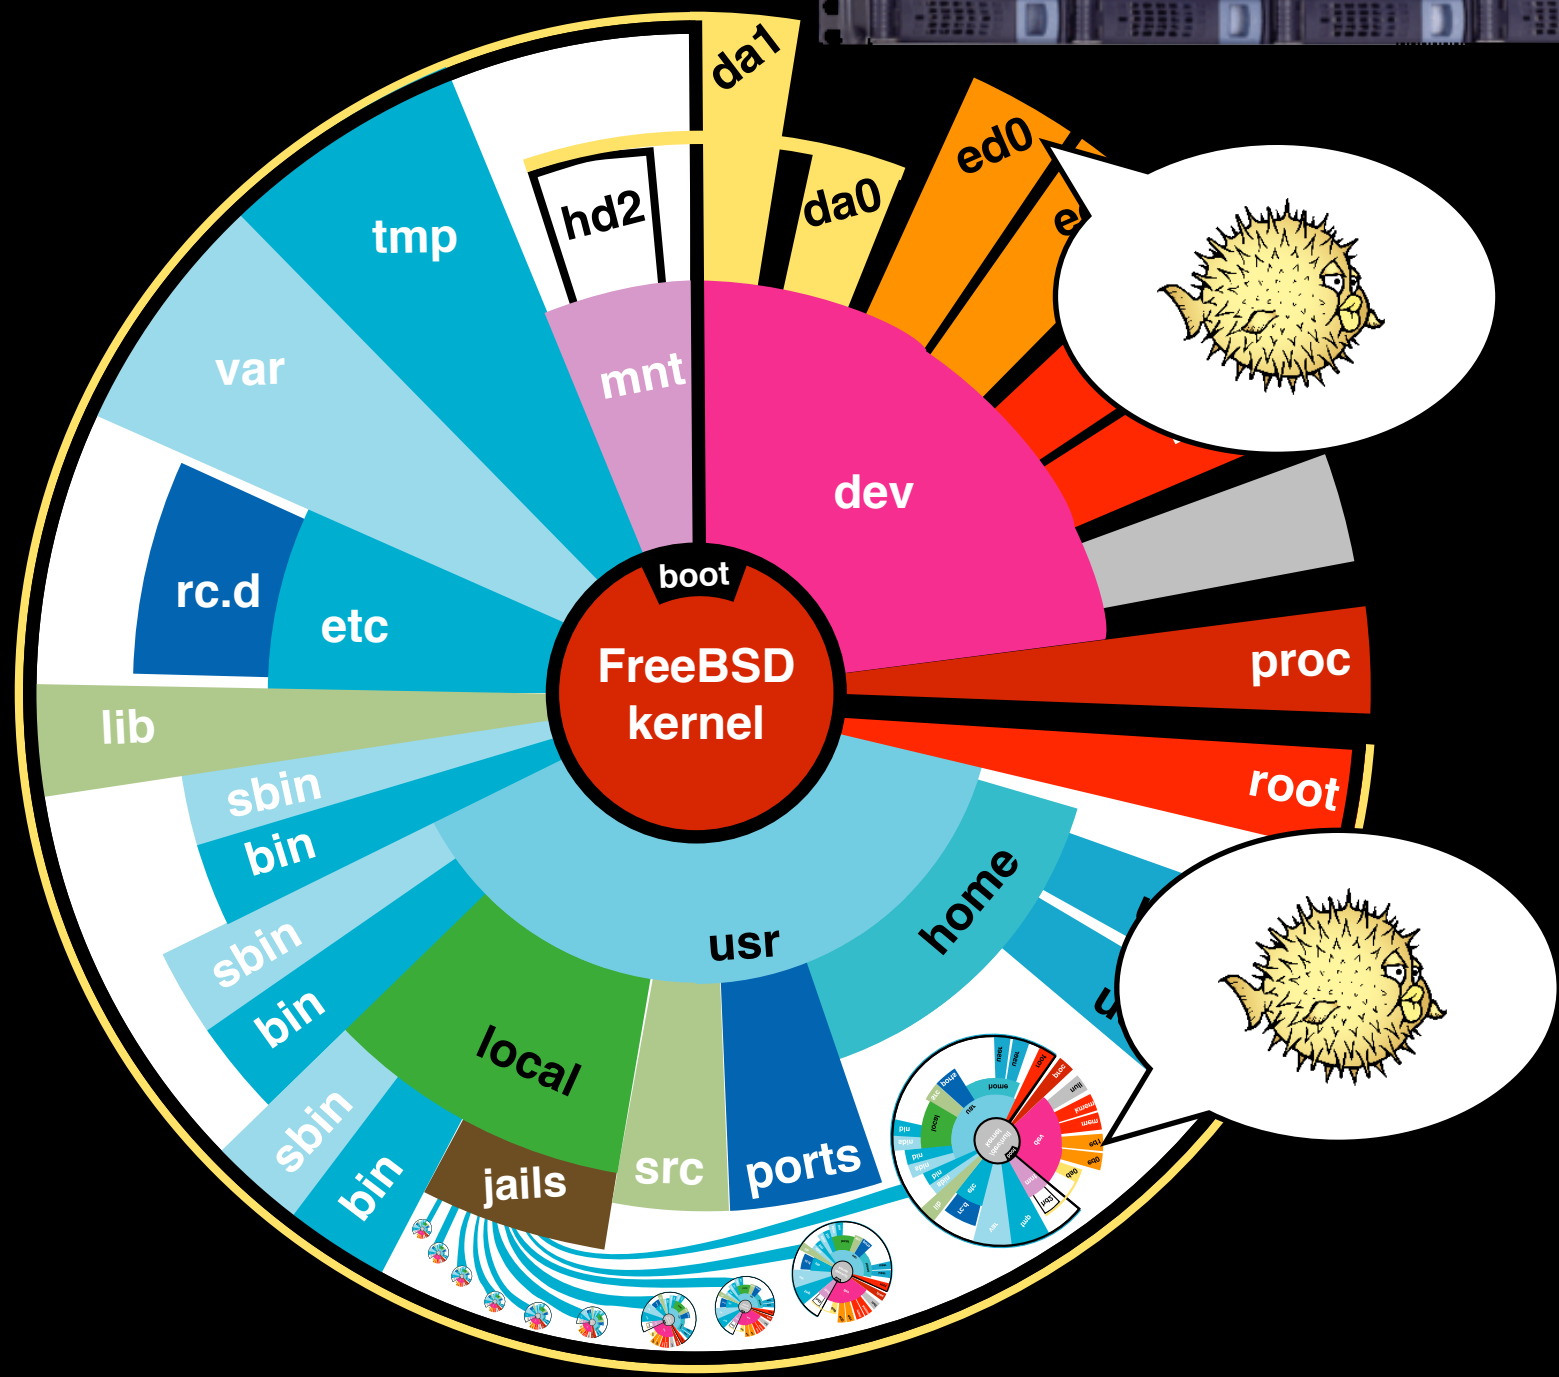
#VersionAddendum FreeBSD-20050903

Port 22
Protocol 2
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
ListenAddress 192.168.1.2

# HostKey for protocol version 1
#HostKey /etc/ssh/ssh_host_key
# HostKeys for protocol version 2
#HostKey /etc/ssh/ssh_host_dsa_key
```

jailinghost:/etc/ssh/sshd\_conf

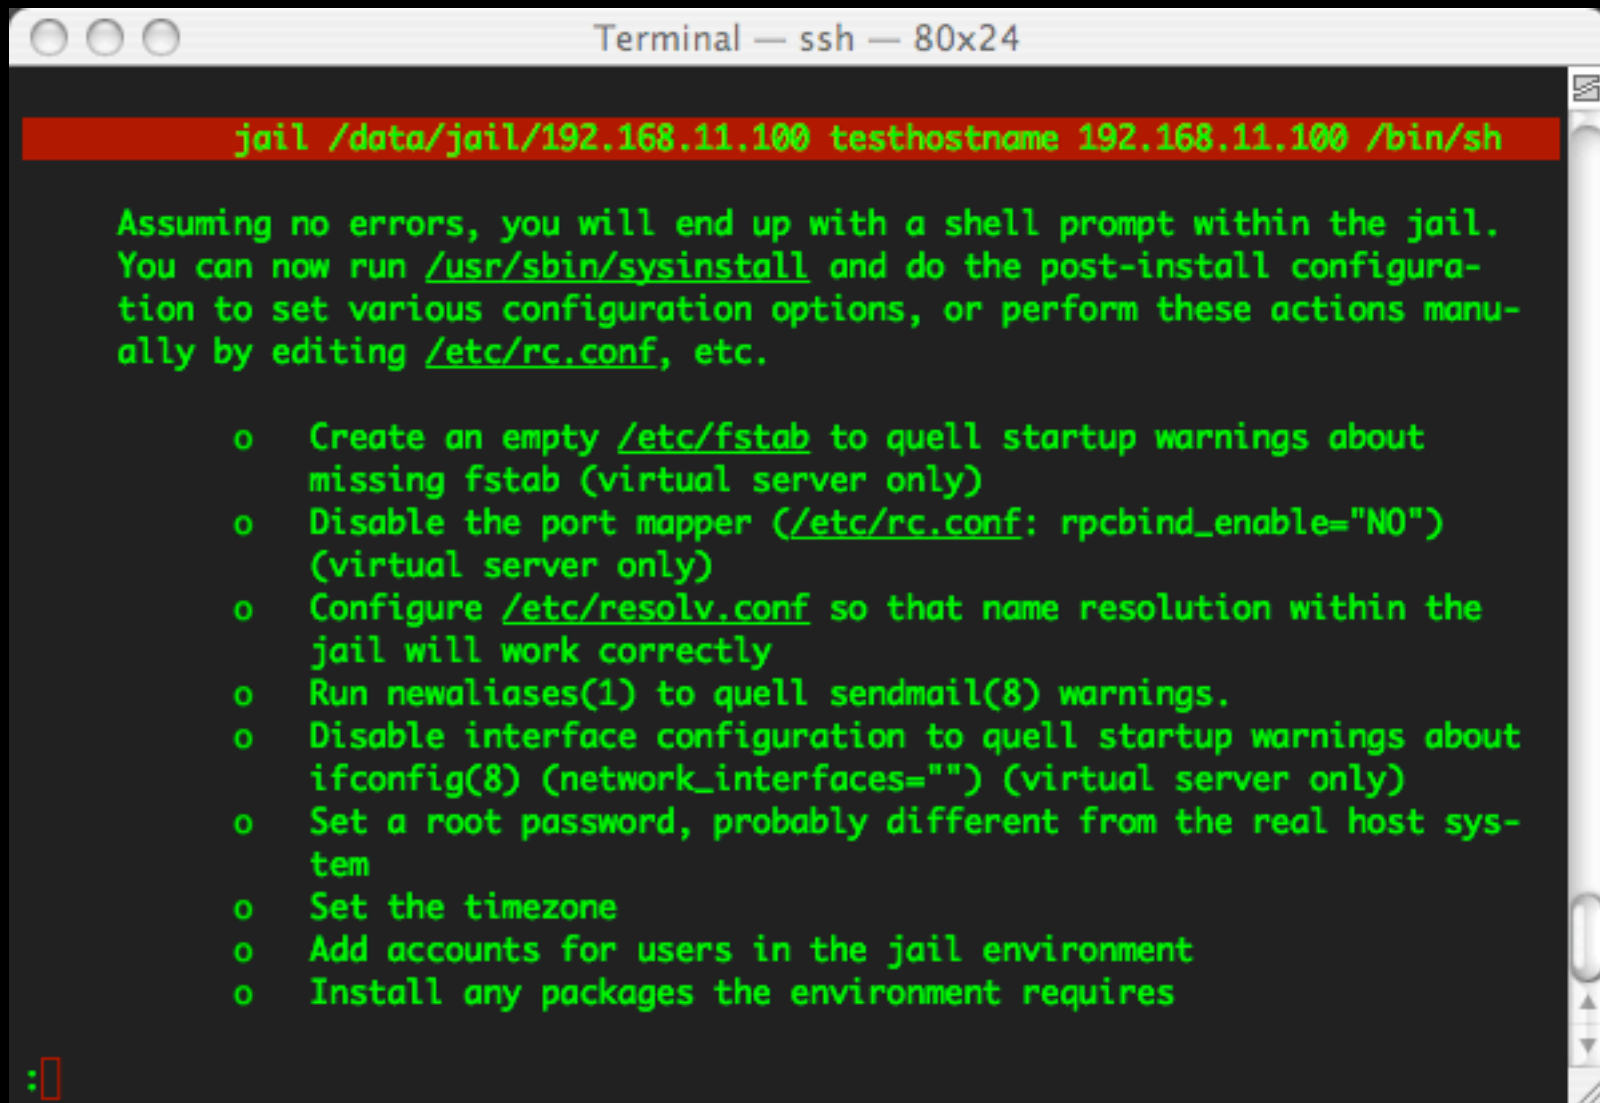




?

# configure - call jailed sh

(analagous to booting a machine in su mode)

A terminal window titled "Terminal — ssh — 80x24" with a red header bar containing the command "jail /data/jail/192.168.11.100 testhostname 192.168.11.100 /bin/sh". The terminal displays instructions for configuring a jail environment, including a list of tasks such as creating an empty /etc/fstab, disabling the port mapper, configuring /etc/resolv.conf, and setting a root password.

```
Terminal — ssh — 80x24
jail /data/jail/192.168.11.100 testhostname 192.168.11.100 /bin/sh

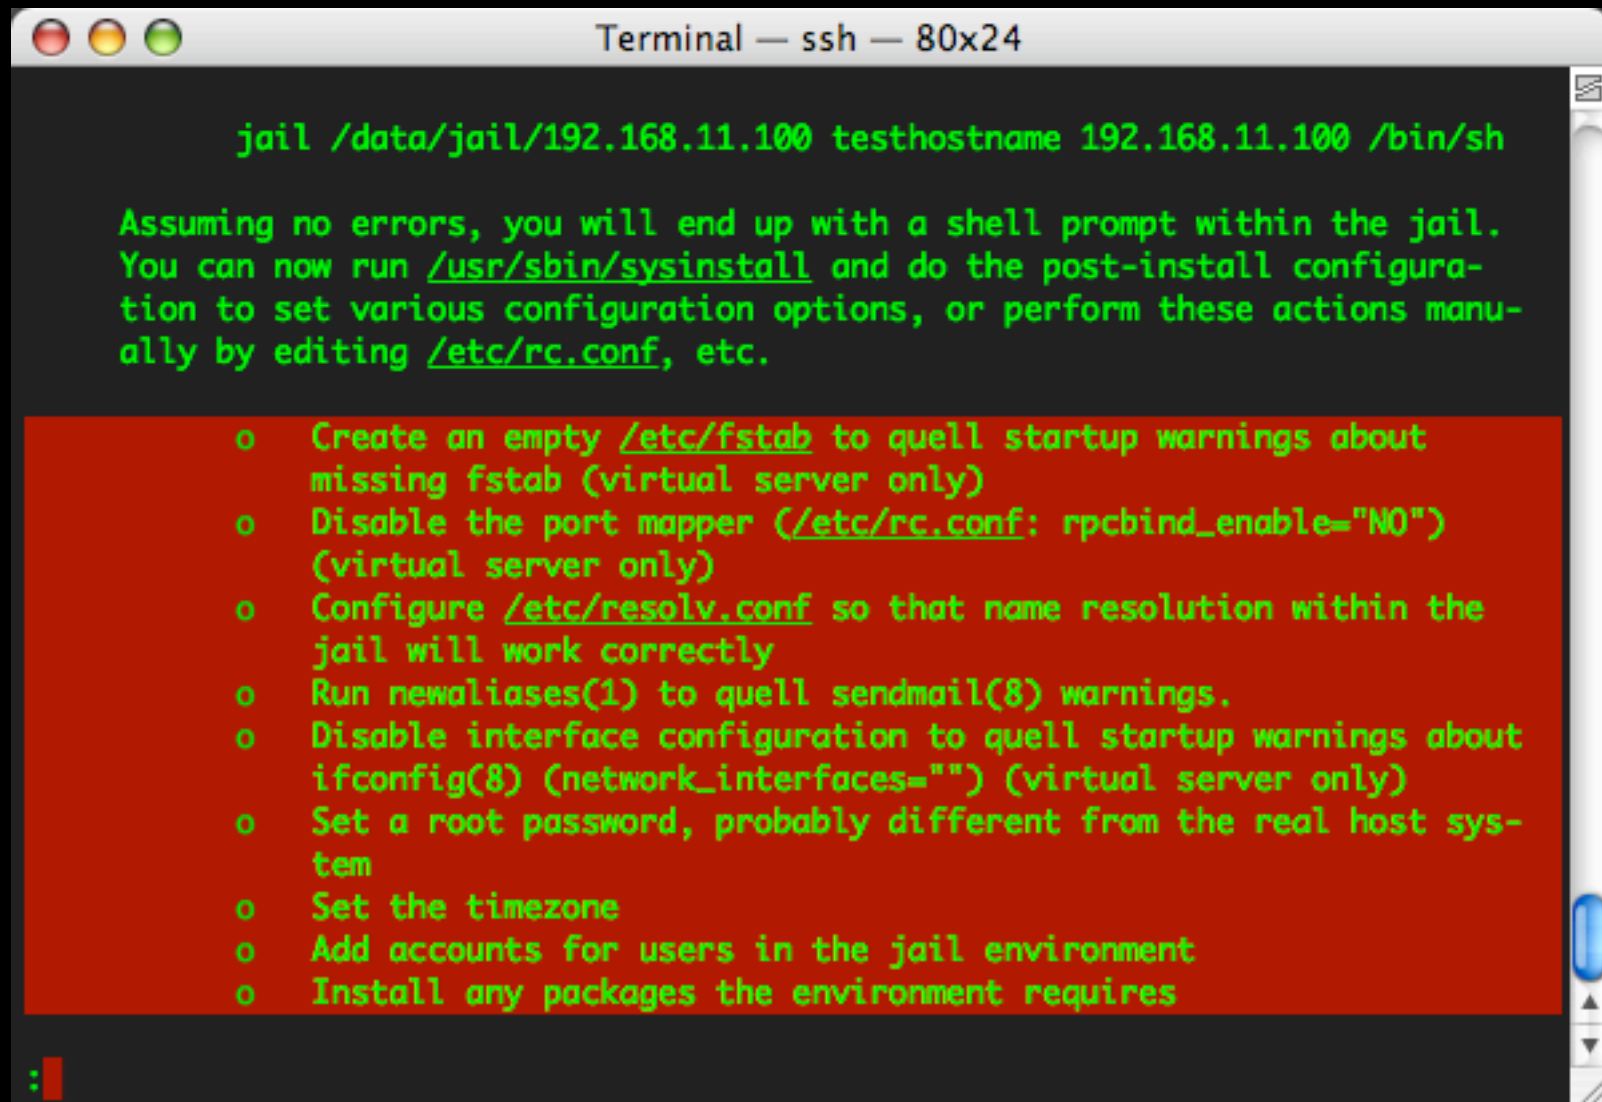
Assuming no errors, you will end up with a shell prompt within the jail.
You can now run /usr/sbin/sysinstall and do the post-install configura-
tion to set various configuration options, or perform these actions manu-
ally by editing /etc/rc.conf, etc.

o Create an empty /etc/fstab to quell startup warnings about
missing fstab (virtual server only)
o Disable the port mapper (/etc/rc.conf: rpcbind_enable="NO")
(virtual server only)
o Configure /etc/resolv.conf so that name resolution within the
jail will work correctly
o Run newaliases(1) to quell sendmail(8) warnings.
o Disable interface configuration to quell startup warnings about
ifconfig(8) (network_interfaces="") (virtual server only)
o Set a root password, probably different from the real host sys-
tem
o Set the timezone
o Add accounts for users in the jail environment
o Install any packages the environment requires

:|
```

# configure - call jailed sh

(analagous to booting a machine in su mode)



```
Terminal — ssh — 80x24

jail /data/jail/192.168.11.100 testhostname 192.168.11.100 /bin/sh

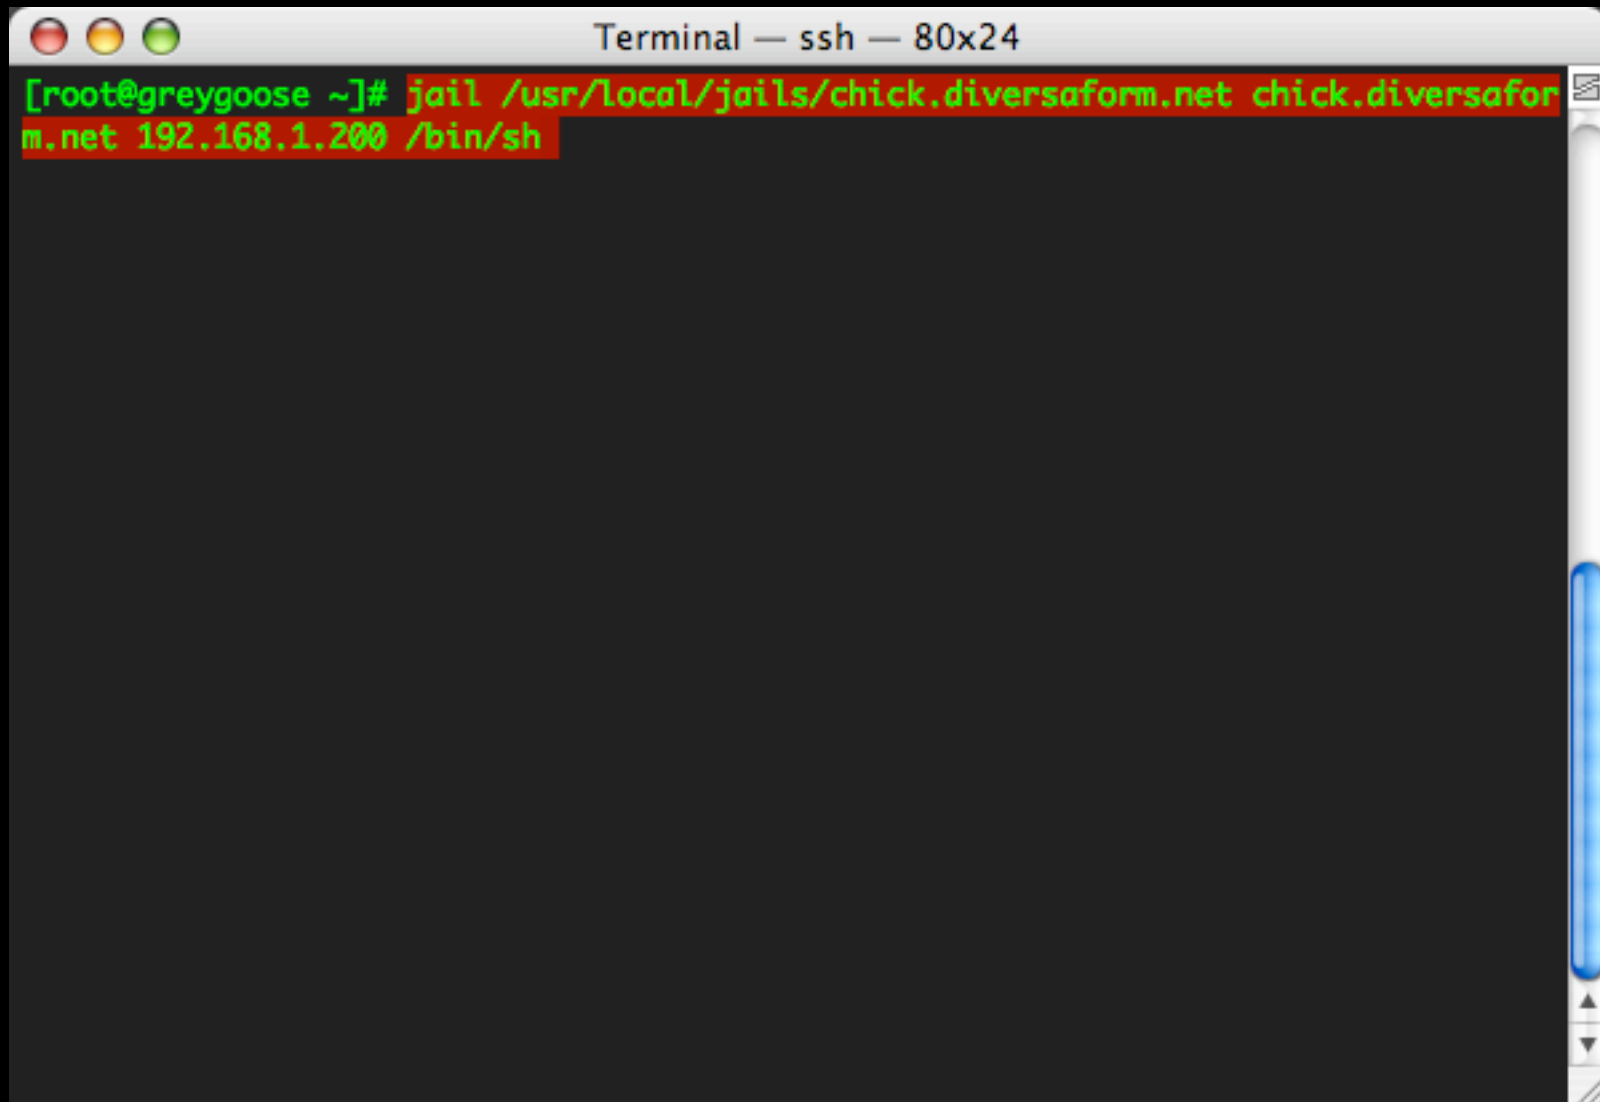
Assuming no errors, you will end up with a shell prompt within the jail.
You can now run /usr/sbin/sysinstall and do the post-install configura-
tion to set various configuration options, or perform these actions manu-
ally by editing /etc/rc.conf, etc.
```

- o Create an empty /etc/fstab to quell startup warnings about missing fstab (virtual server only)
- o Disable the port mapper (/etc/rc.conf: `rpcbind_enable="NO"`) (virtual server only)
- o Configure /etc/resolv.conf so that name resolution within the jail will work correctly
- o Run `newaliases(1)` to quell `sendmail(8)` warnings.
- o Disable interface configuration to quell startup warnings about `ifconfig(8)` (`network_interfaces=""`) (virtual server only)
- o Set a root password, probably different from the real host system
- o Set the timezone
- o Add accounts for users in the jail environment
- o Install any packages the environment requires

```
:|
```

# configure - call jailed sh

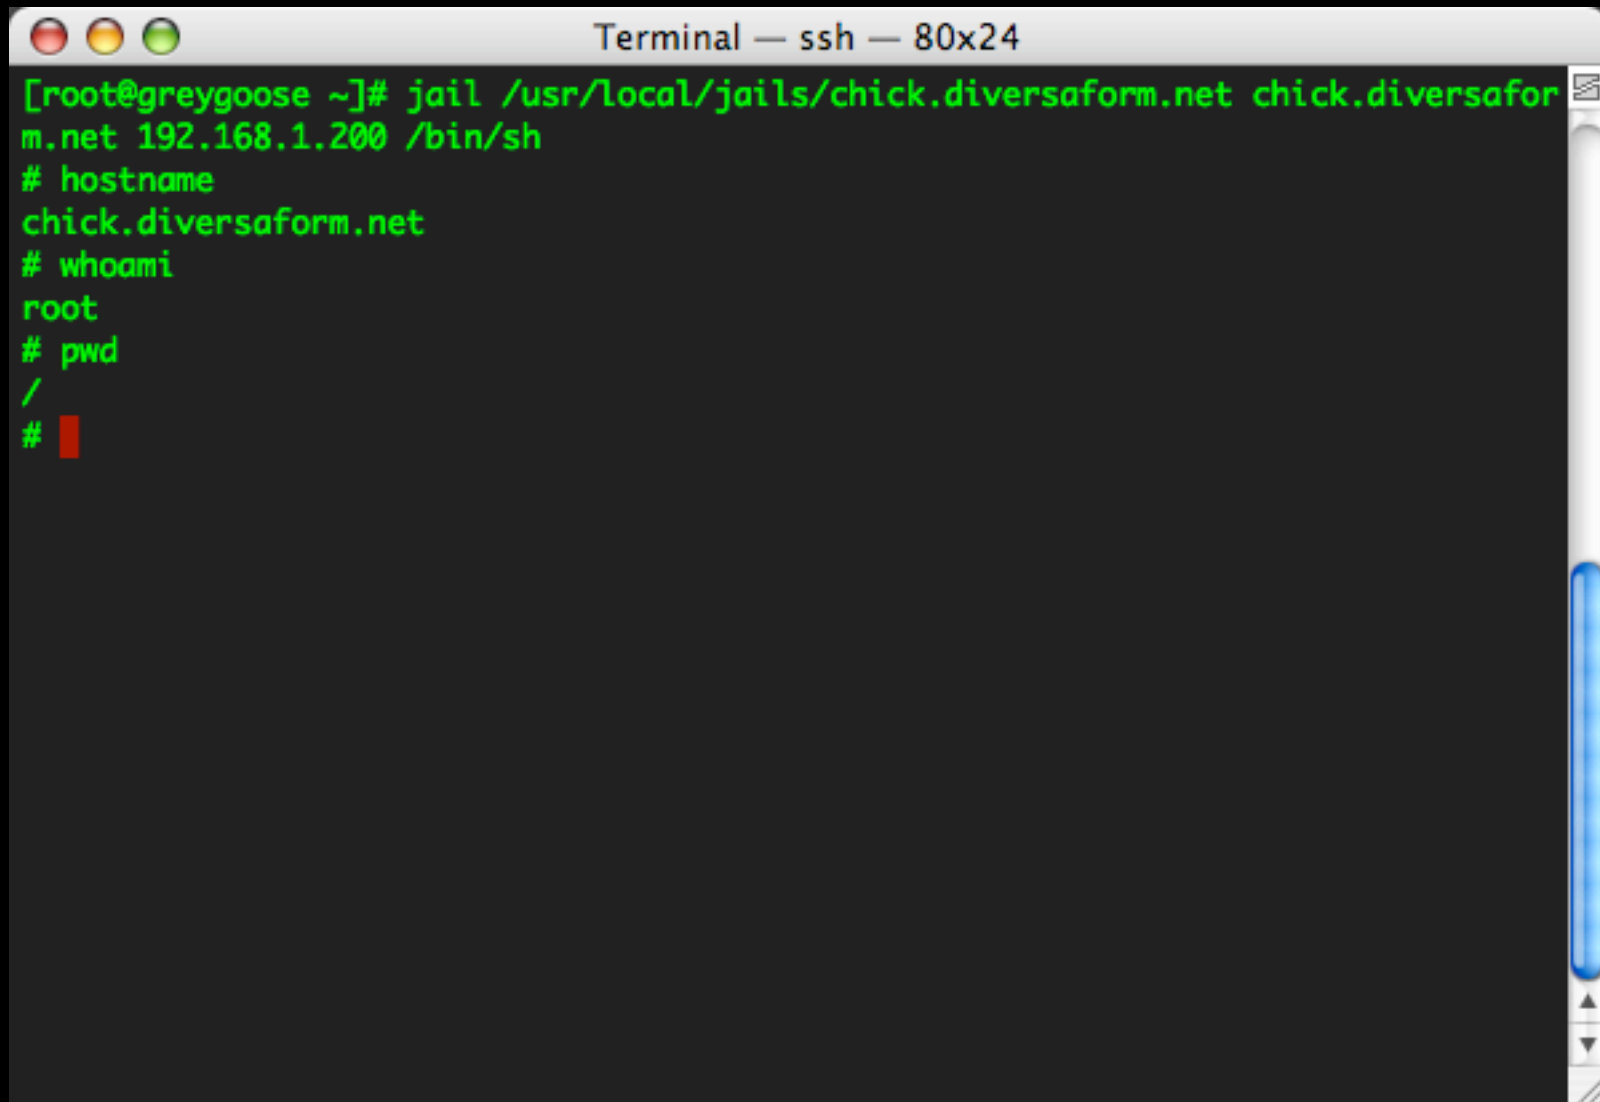
(analagous to booting a machine in su mode)

A terminal window titled "Terminal — ssh — 80x24" with standard macOS window controls (red, yellow, green buttons). The terminal content shows a root user at a host named "greygoose" in their home directory (~) typing a command to create a jail. The command is highlighted in red: `jail /usr/local/jails/chick.diversaform.net chick.diversaform.net 192.168.1.200 /bin/sh`. The rest of the terminal is empty.

```
[root@greygoose ~]# jail /usr/local/jails/chick.diversaform.net chick.diversaform.net 192.168.1.200 /bin/sh
```

# configure - call jailed sh

configure the jail, inside the jail

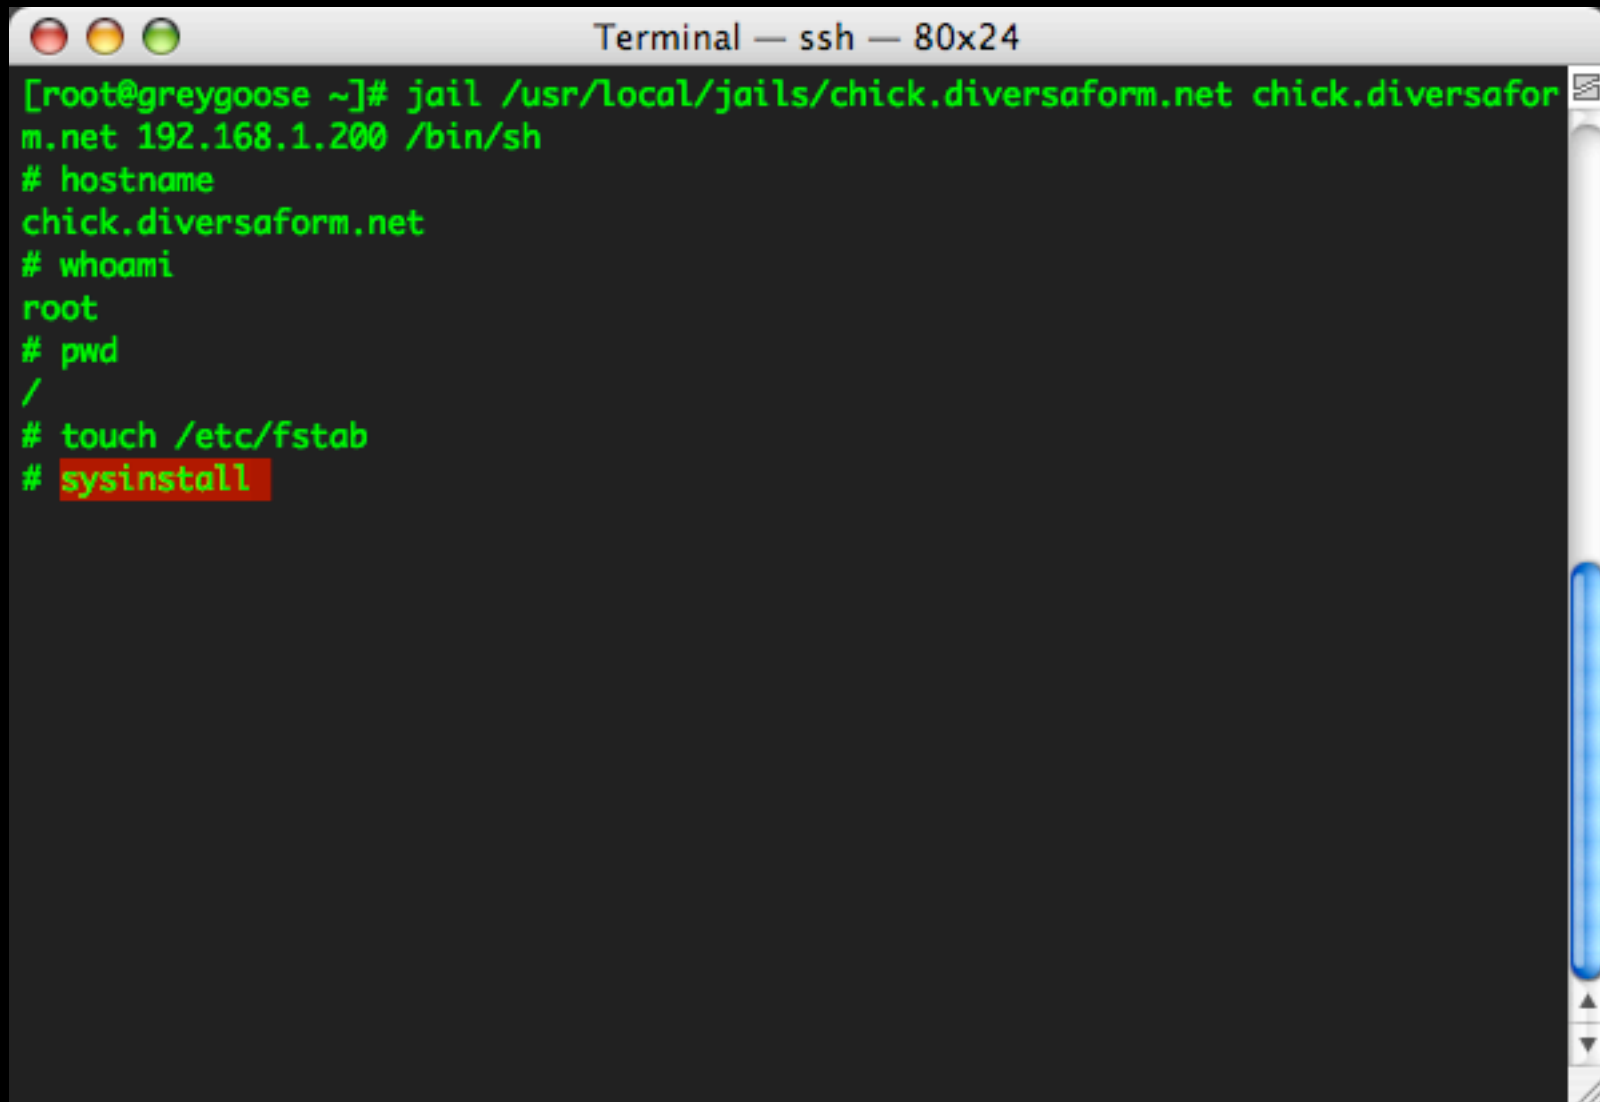
A terminal window titled "Terminal — ssh — 80x24" with standard macOS window controls (red, yellow, green buttons). The terminal shows a sequence of commands and their outputs in green text on a black background. The commands are: 'jail /usr/local/jails/chick.diversaform.net chick.diversaform.net 192.168.1.200 /bin/sh', '# hostname', '# whoami', and '# pwd'. The outputs are 'chick.diversaform.net', 'root', and '/' respectively. A red cursor is visible after the last command.

```
[root@greygoose ~]# jail /usr/local/jails/chick.diversaform.net chick.diversaform
m.net 192.168.1.200 /bin/sh
# hostname
chick.diversaform.net
# whoami
root
# pwd
/
# █
```



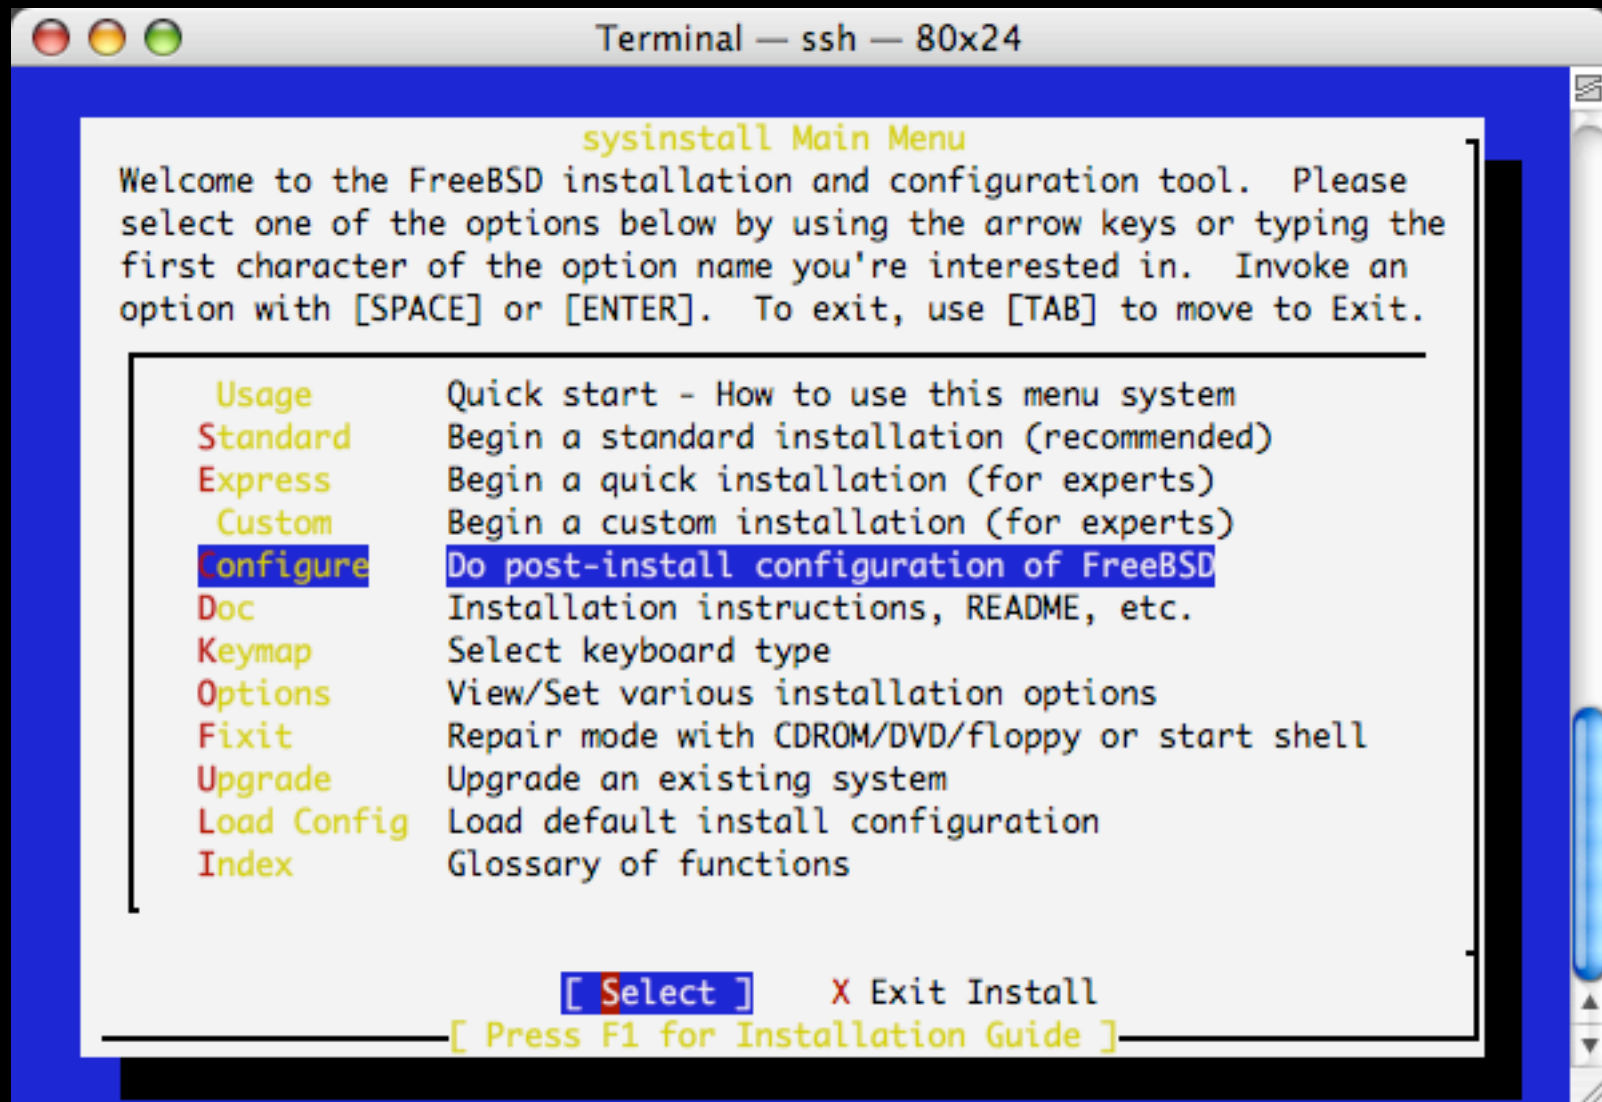
# configure - call jailed sh

configure the jail, inside the jail

A terminal window titled "Terminal — ssh — 80x24" with a dark background and light green text. The window shows a sequence of commands and their outputs. The first command is "jail /usr/local/jails/chick.diversaform.net chick.diversaform.net 192.168.1.200 /bin/sh". Subsequent commands include "hostname" (output: chick.diversaform.net), "whoami" (output: root), "pwd" (output: /), "touch /etc/fstab", and "sysinstall" (highlighted in red).

```
Terminal — ssh — 80x24
[root@greygoose ~]# jail /usr/local/jails/chick.diversaform.net chick.diversaform
m.net 192.168.1.200 /bin/sh
# hostname
chick.diversaform.net
# whoami
root
# pwd
/
# touch /etc/fstab
# sysinstall
```

# configure - call jailed sh sysctl, whee!



```
Terminal — ssh — 80x24

sysinstall Main Menu

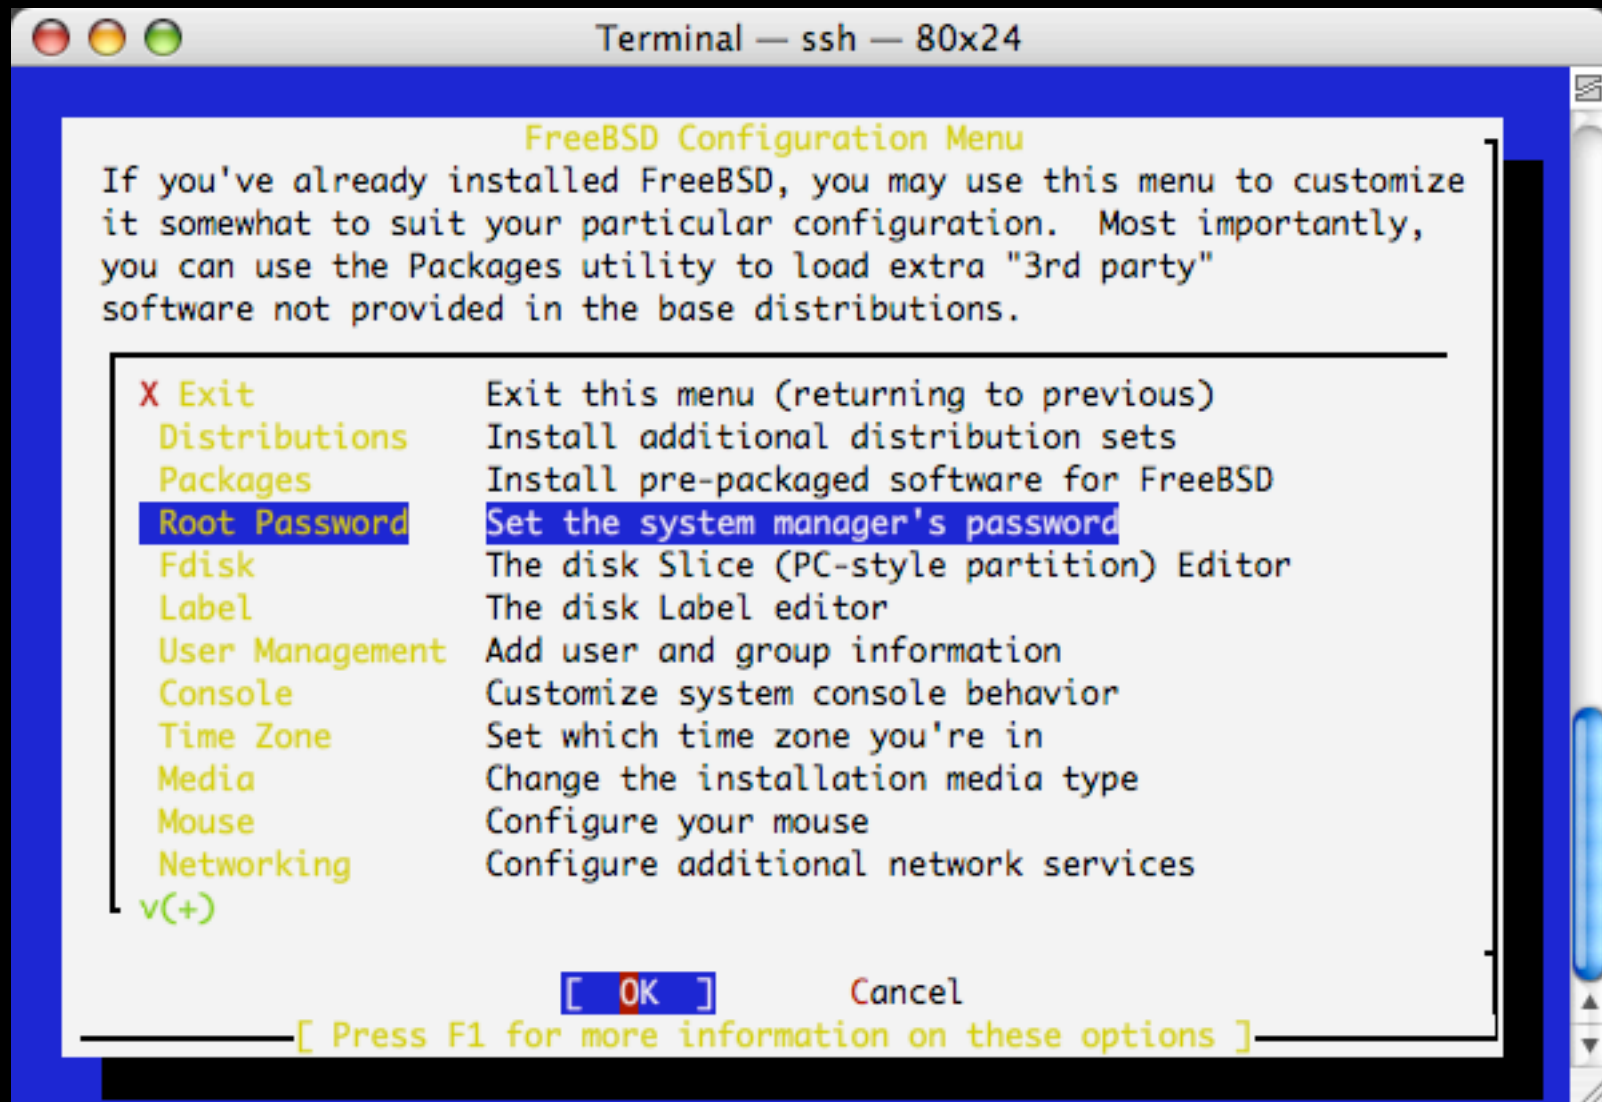
Welcome to the FreeBSD installation and configuration tool. Please
select one of the options below by using the arrow keys or typing the
first character of the option name you're interested in. Invoke an
option with [SPACE] or [ENTER]. To exit, use [TAB] to move to Exit.

Usage      Quick start - How to use this menu system
Standard   Begin a standard installation (recommended)
Express    Begin a quick installation (for experts)
Custom     Begin a custom installation (for experts)
Configure  Do post-install configuration of FreeBSD
Doc        Installation instructions, README, etc.
Keymap     Select keyboard type
Options    View/Set various installation options
Fixit      Repair mode with CDROM/DVD/floppy or start shell
Upgrade    Upgrade an existing system
Load Config Load default install configuration
Index      Glossary of functions

[ Select ]      X Exit Install
[ Press F1 for Installation Guide ]
```


# configure - call jailed sh

## root pw



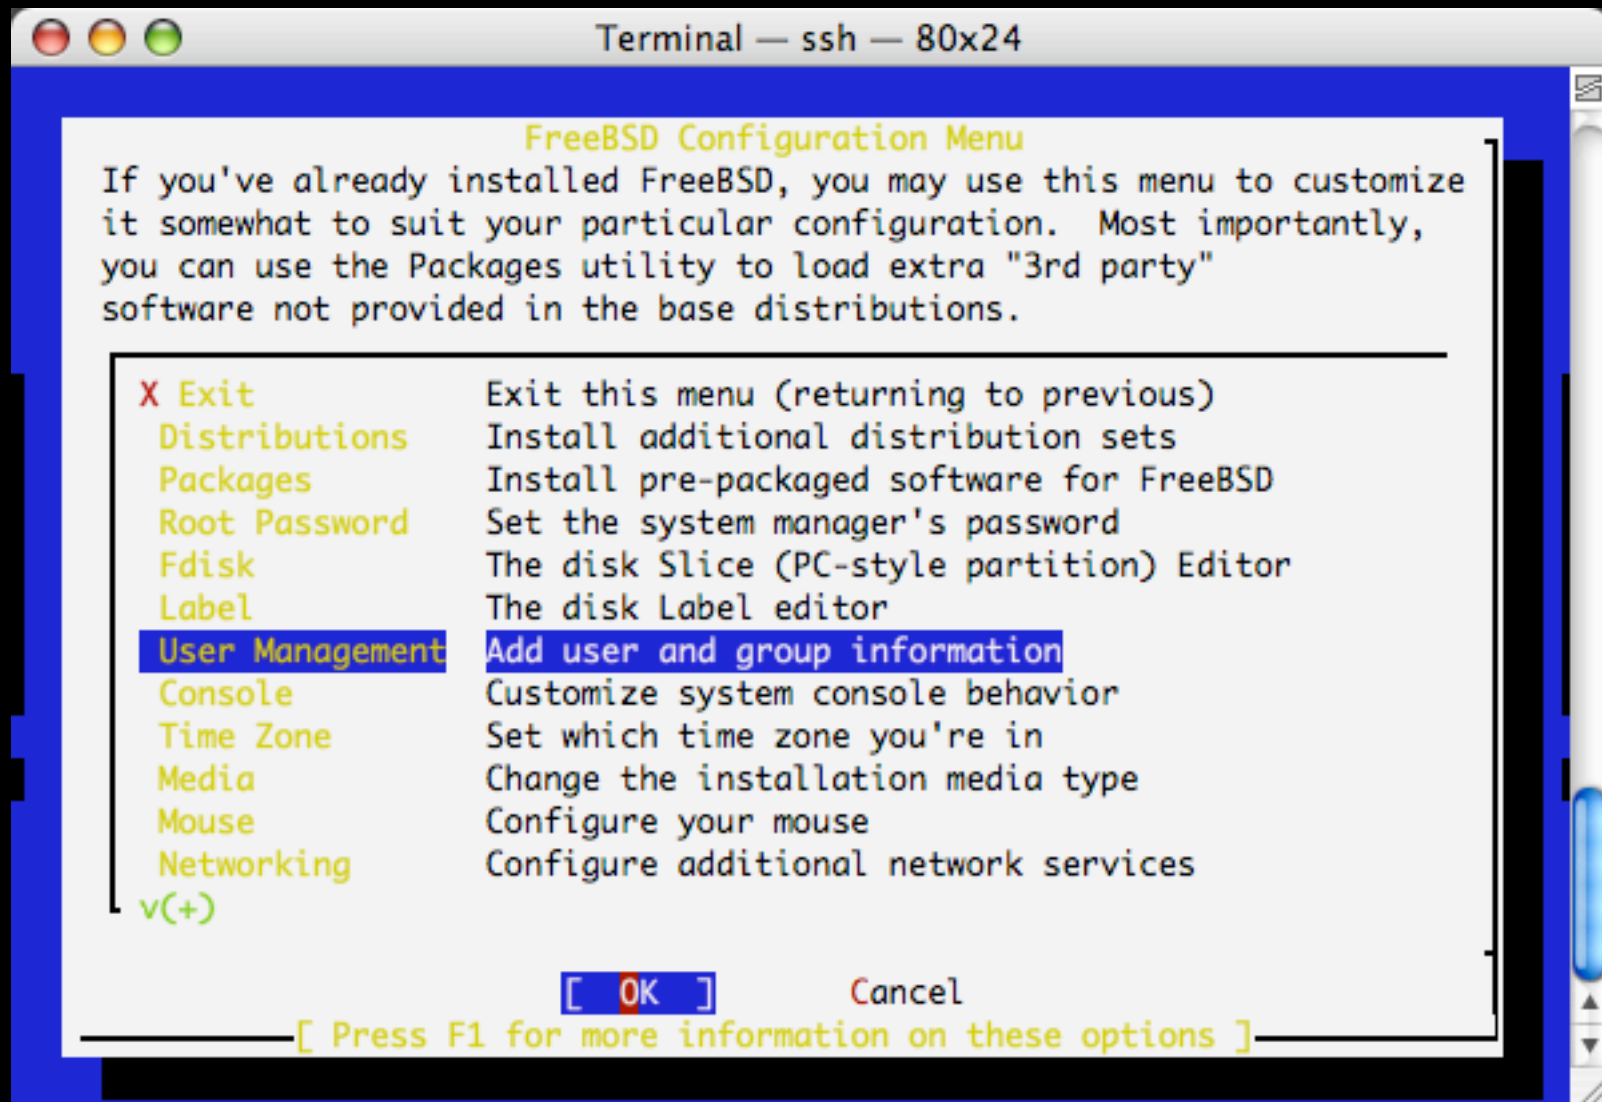
# configure - call jailed sh

root pw



# configure - call jailed sh

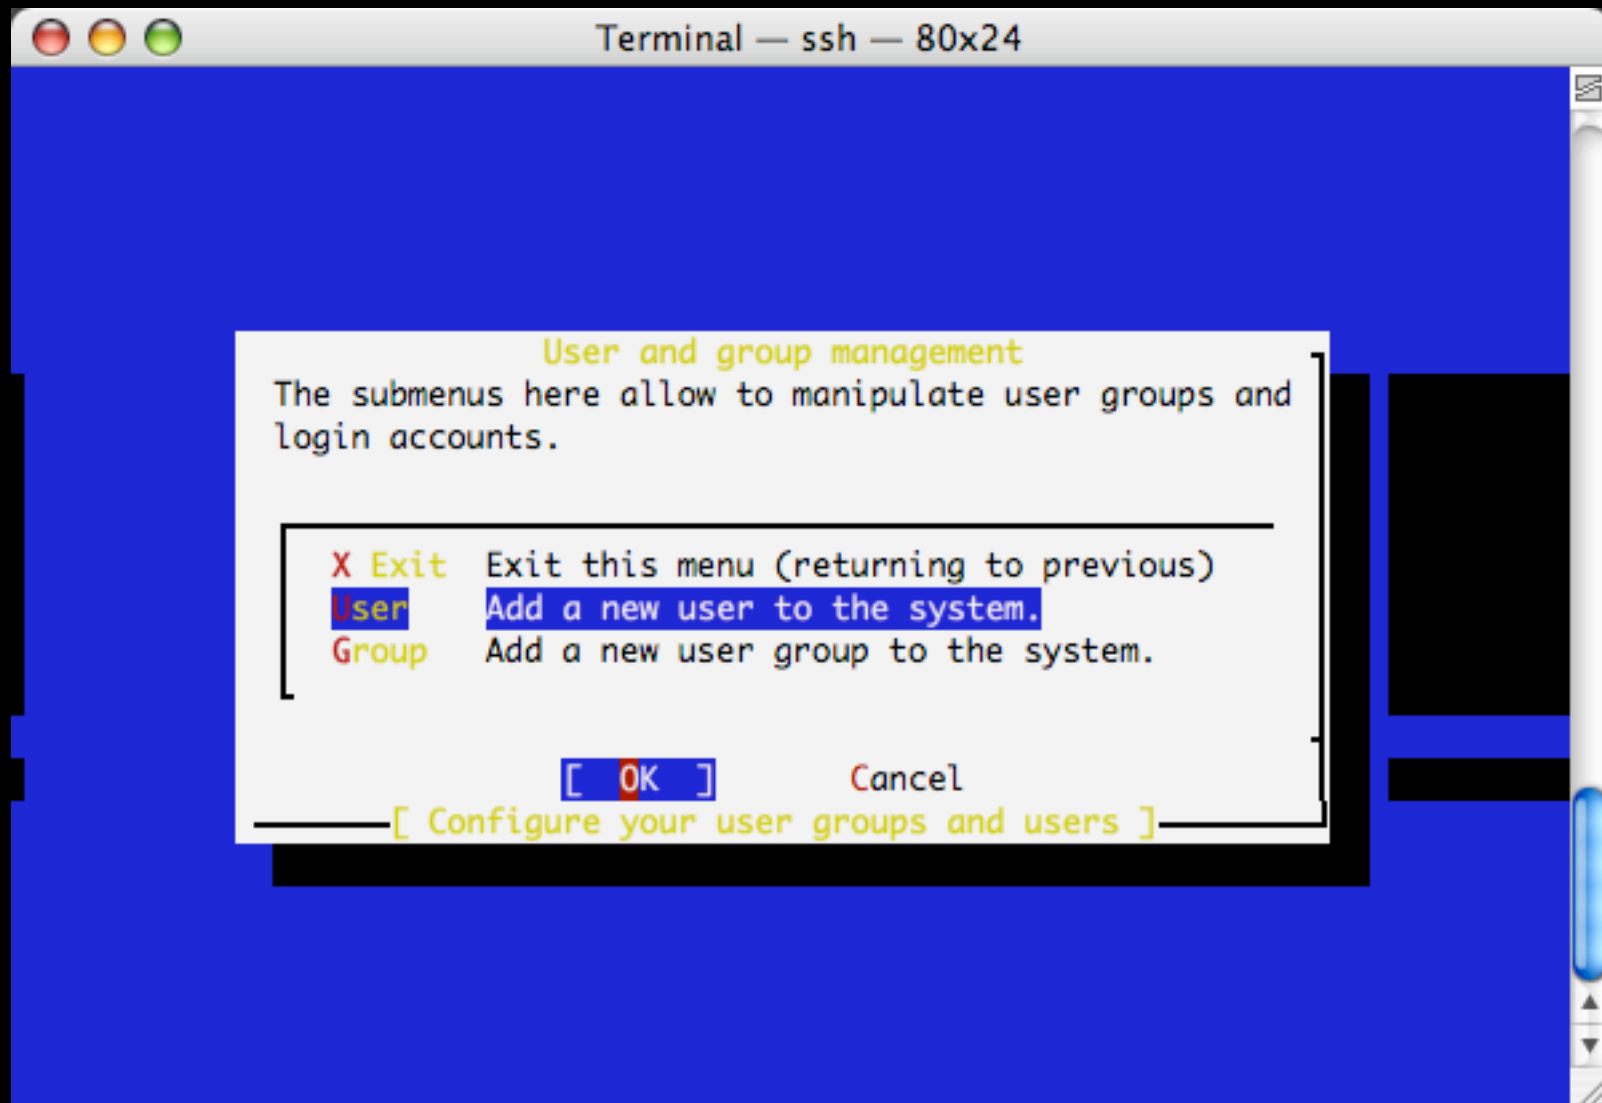
## add users





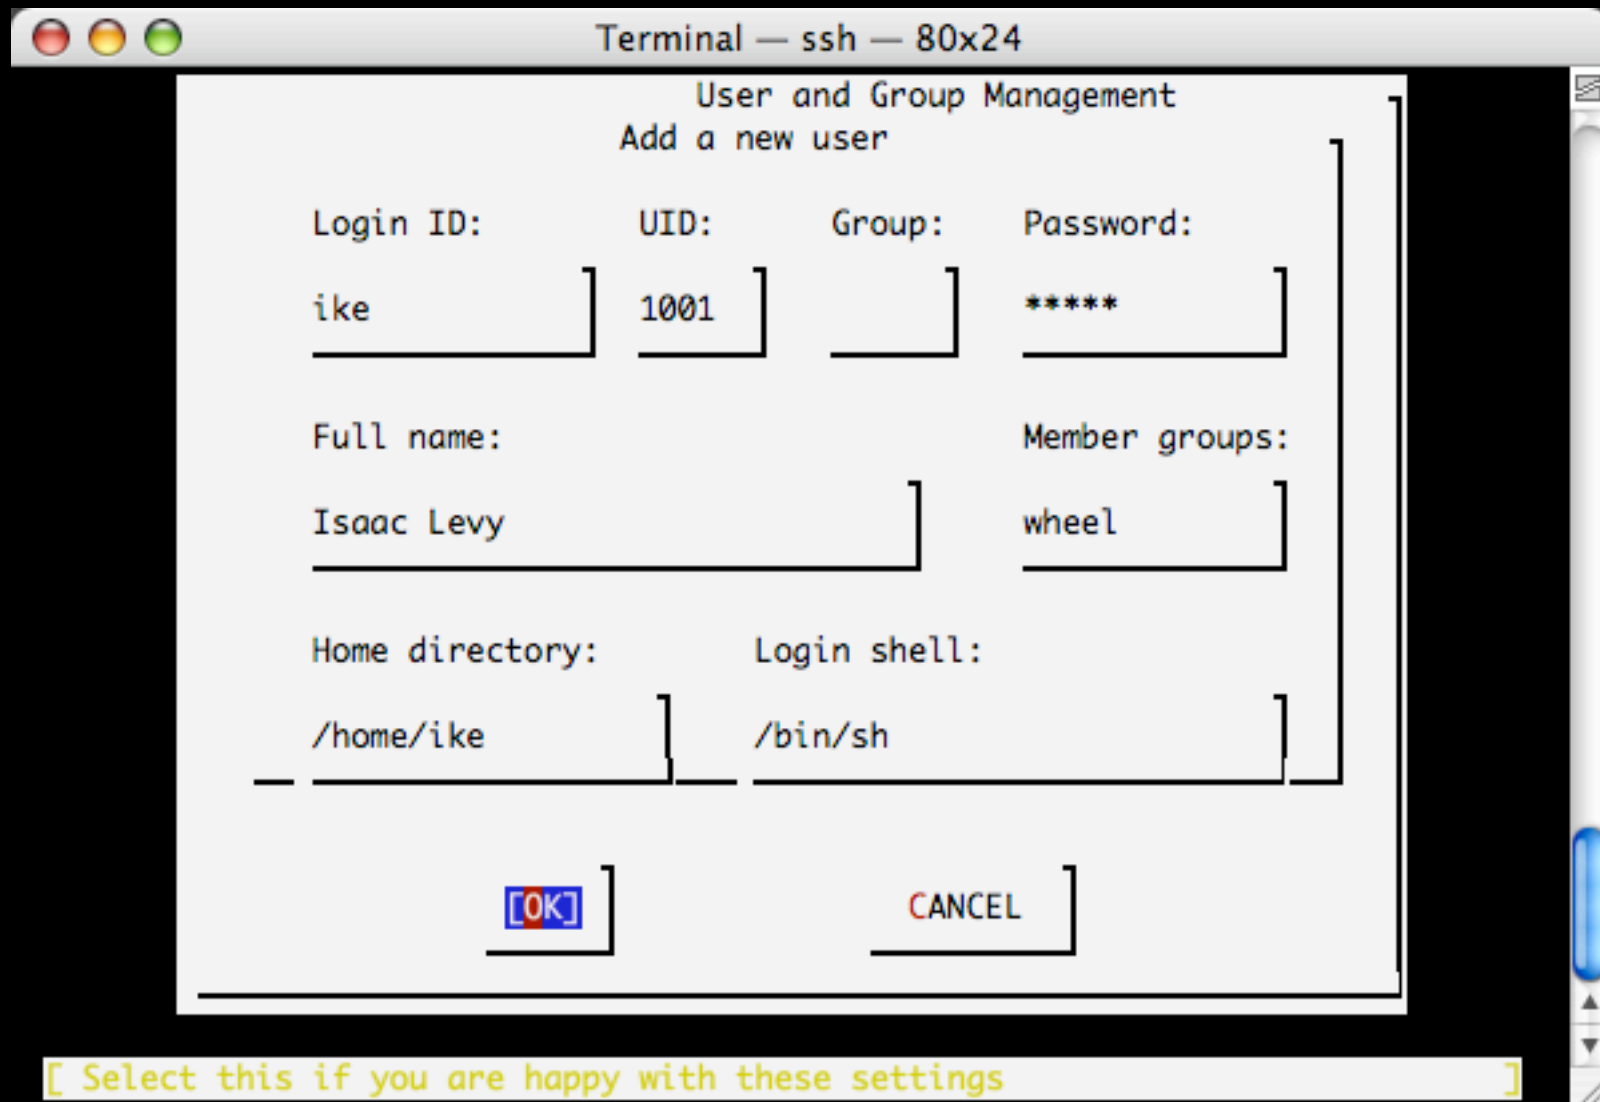
# configure - call jailed sh

## add users



# configure - call jailed sh

## add users



Terminal — ssh — 80x24

User and Group Management  
Add a new user

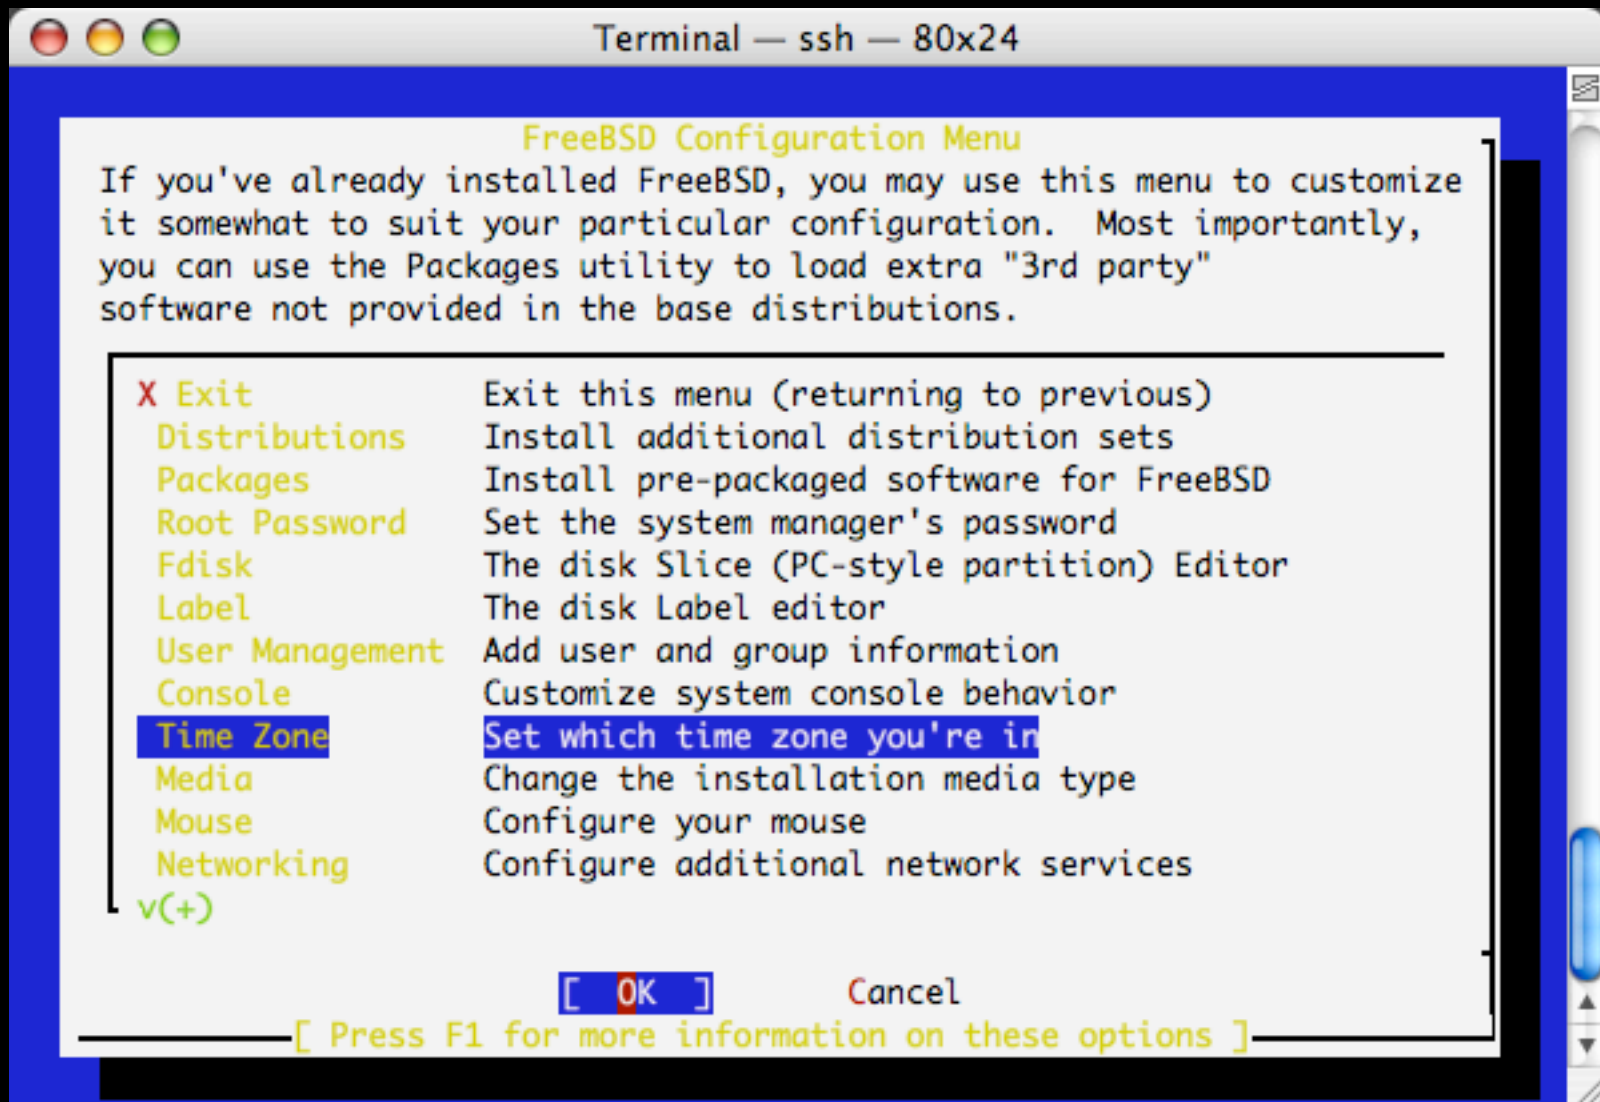
Login ID:	UID:	Group:	Password:
ike	1001		*****
Full name:		Member groups:	
Isaac Levy		wheel	
Home directory:	Login shell:		
/home/ike	/bin/sh		

[ OK ] [ CANCEL ]

[ Select this if you are happy with these settings ]

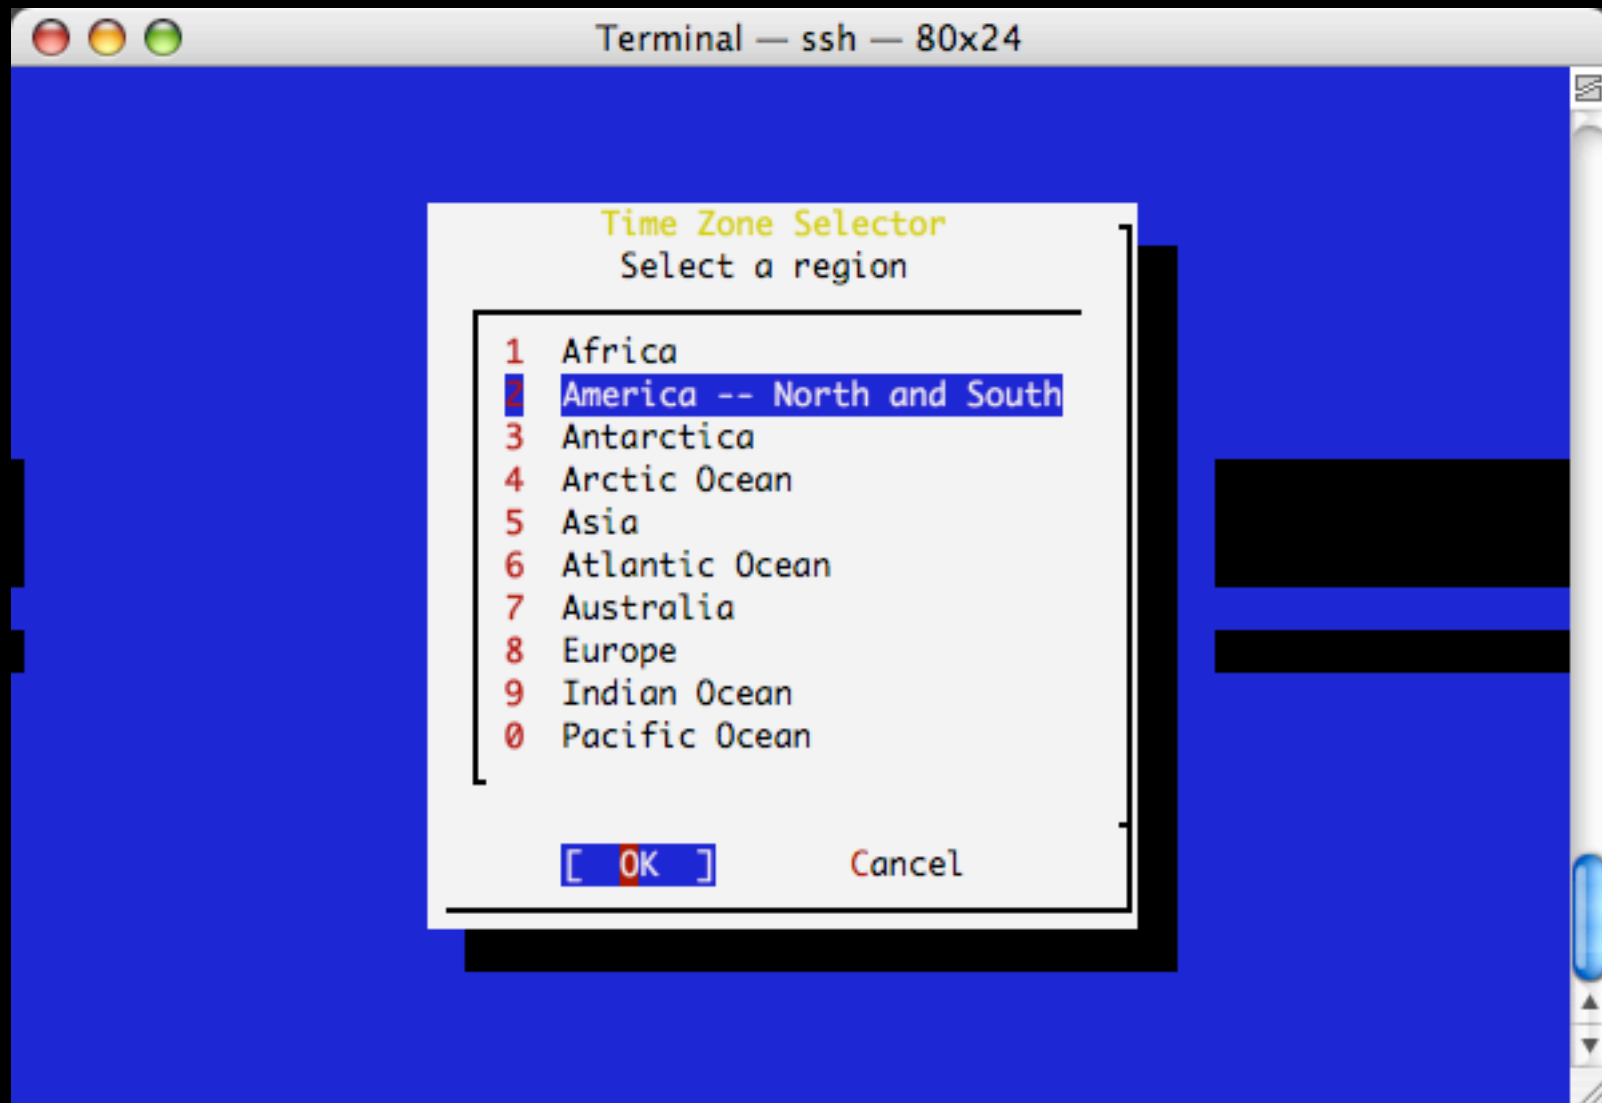
# configure - call jailed sh

## set timezone



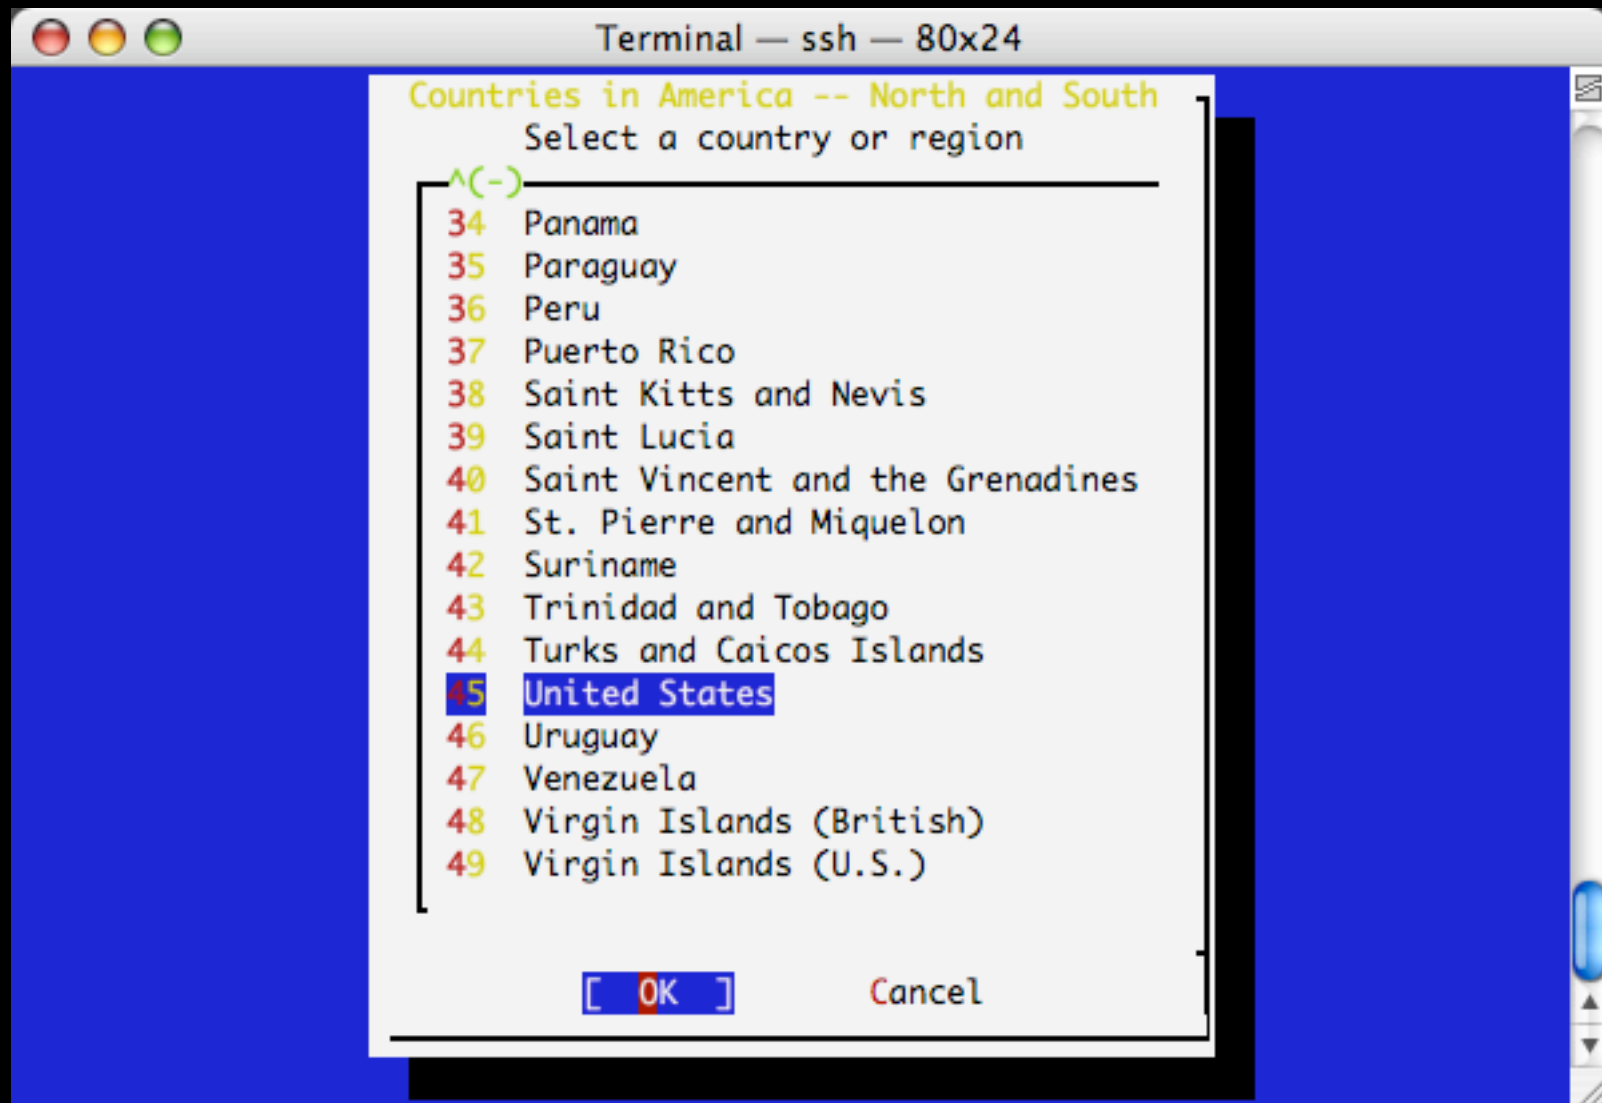
# configure - call jailed sh

## set timezone



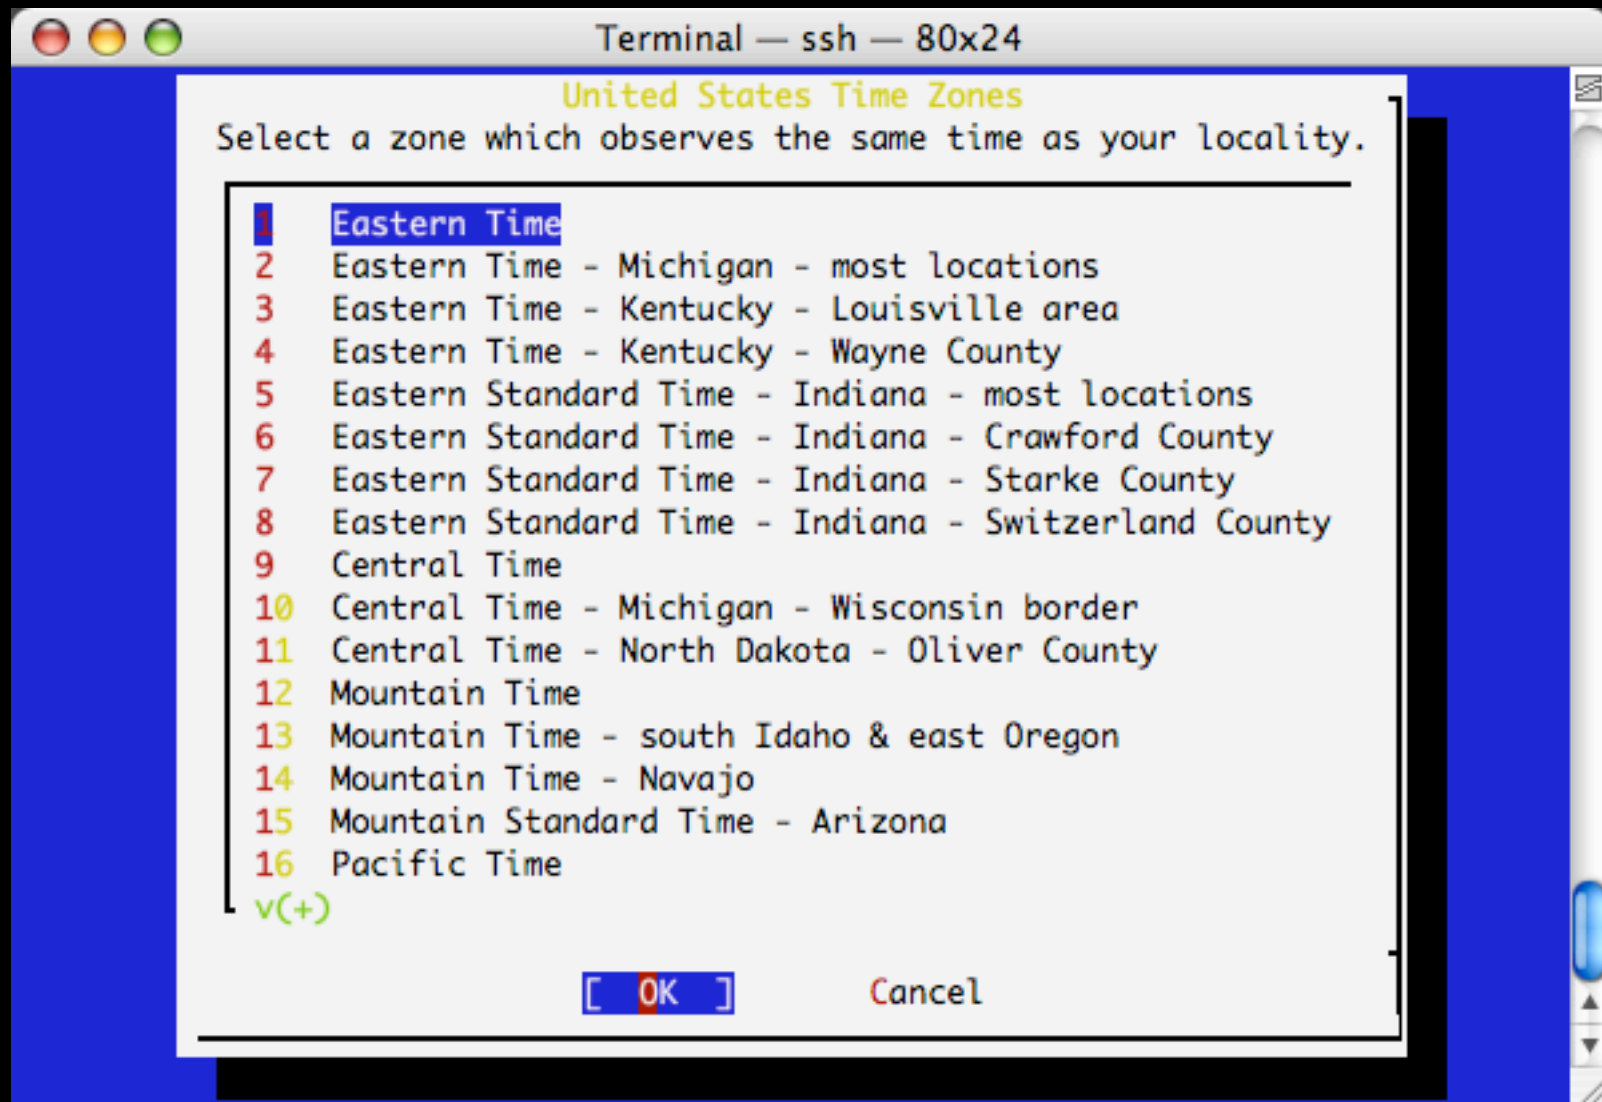
# configure - call jailed sh

## set timezone



# configure - call jailed sh


## set timezone





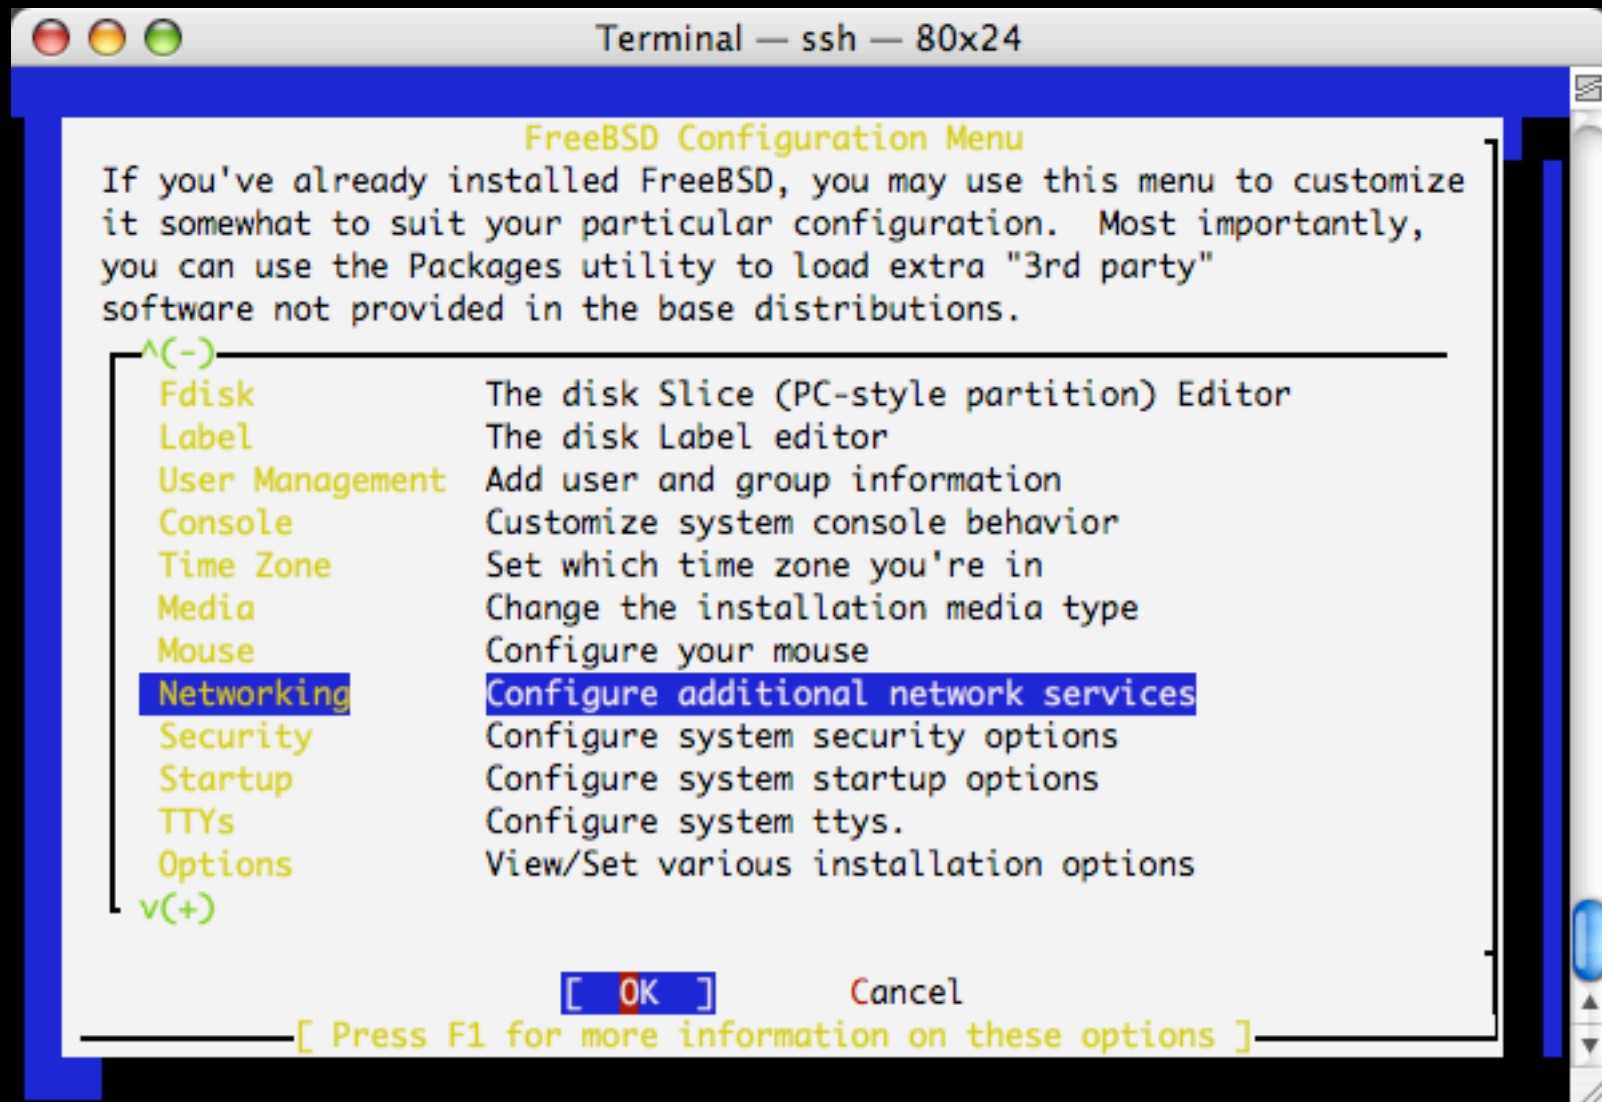
# configure - call jailed sh

## set timezone



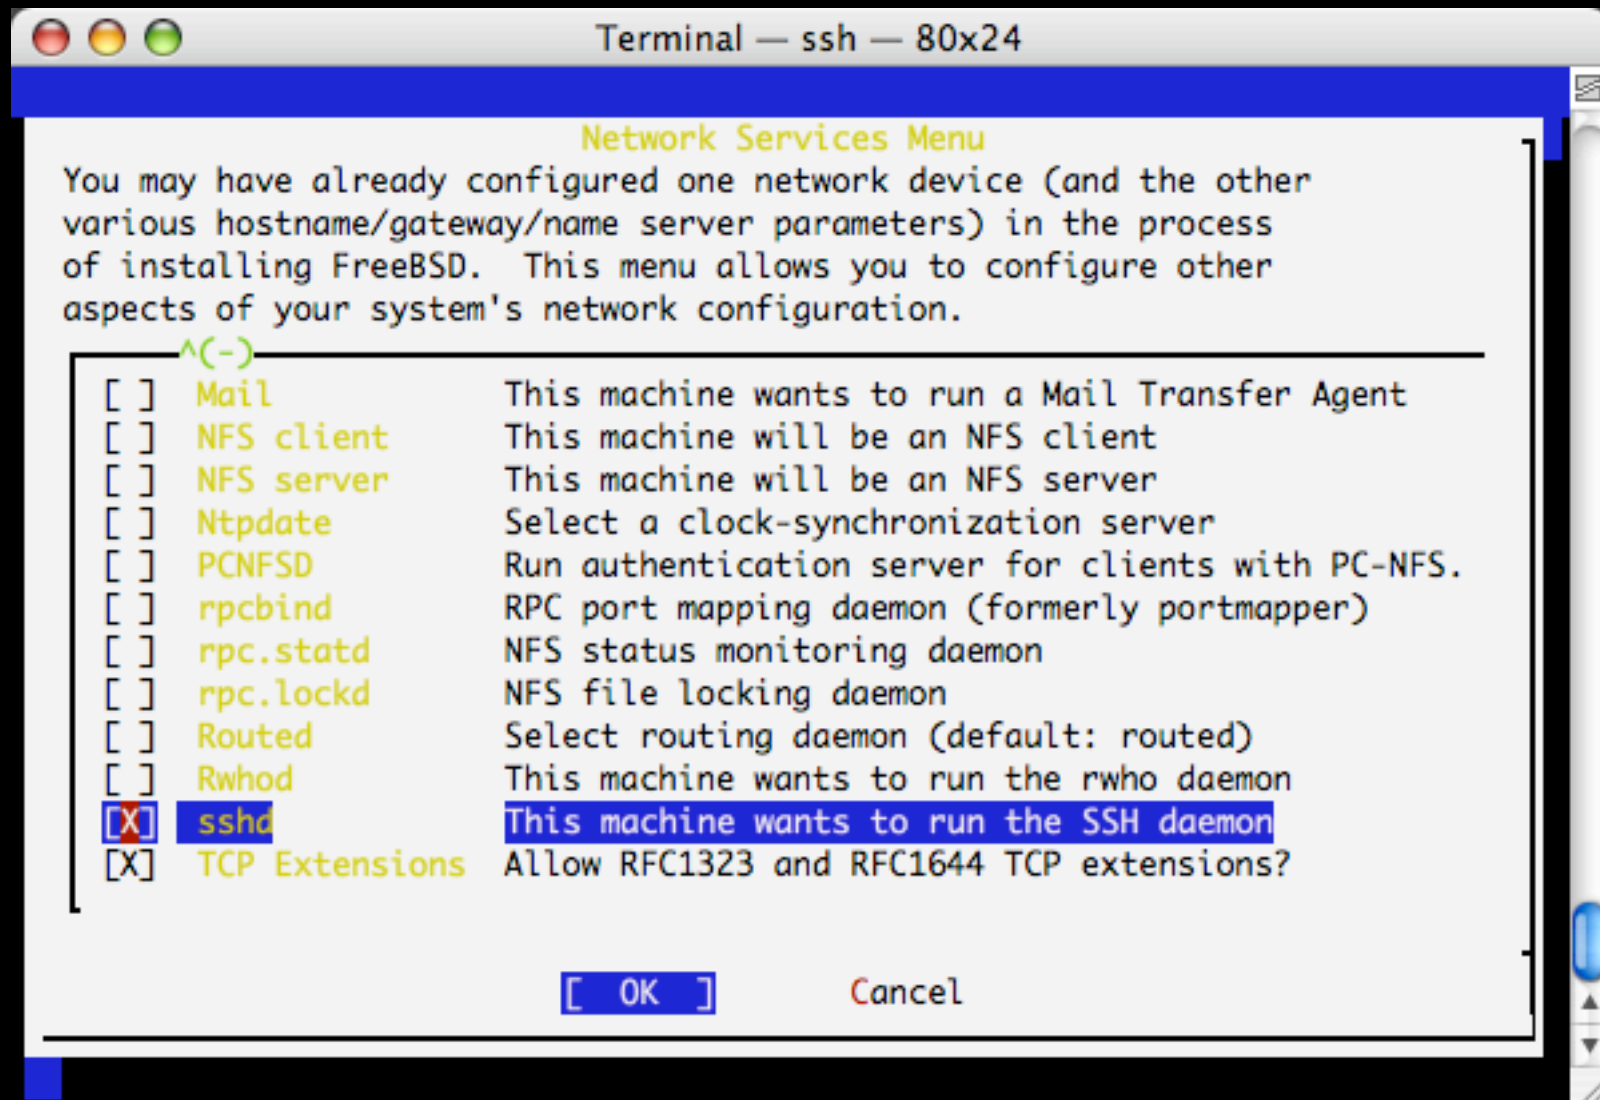
# configure - call jailed sh

## network options...



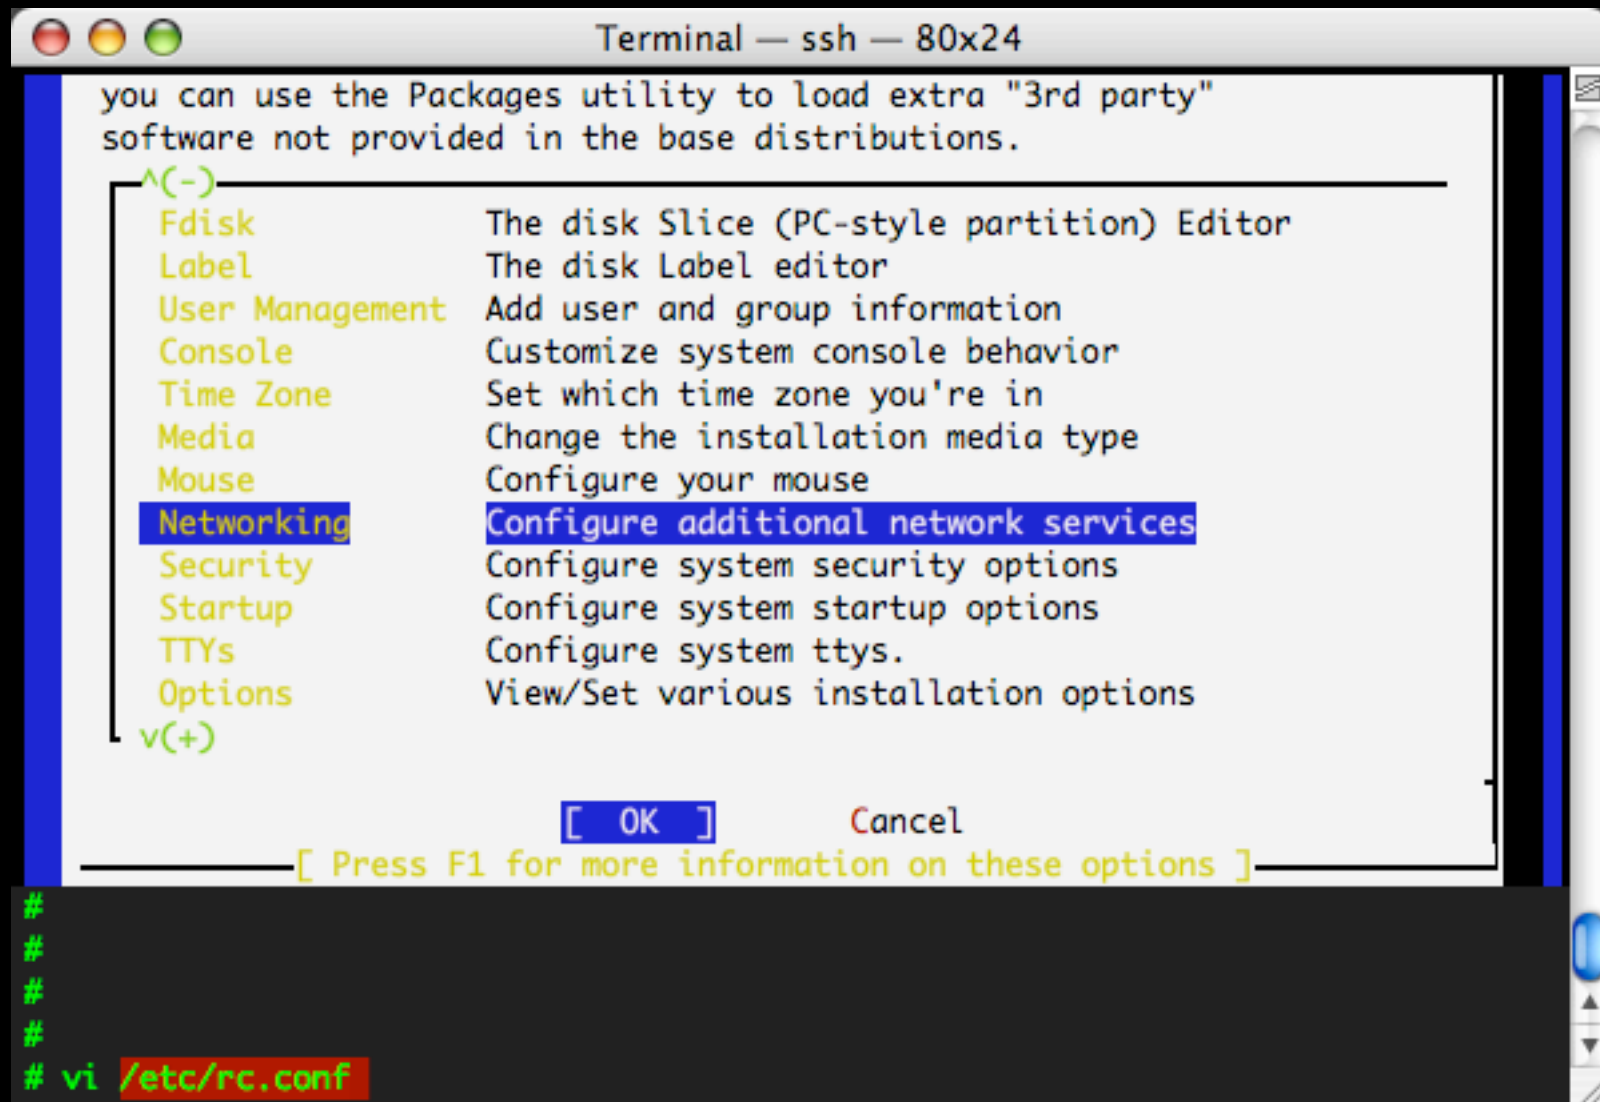
# configure - call jailed sh

run ssh, important



# configure - call jailed sh

check rc.conf in jail



The image shows a terminal window titled "Terminal — ssh — 80x24". The terminal displays the output of the 'configure' utility, which lists various system configuration options. The 'Networking' option is highlighted in blue. Below the list, there are buttons for 'OK' and 'Cancel', and a prompt to press F1 for more information. At the bottom of the terminal, the prompt '# vi /etc/rc.conf' is visible, indicating that the user has opened the rc.conf file in the vi editor.

```
you can use the Packages utility to load extra "3rd party"
software not provided in the base distributions.
^(-)
Fdisk          The disk Slice (PC-style partition) Editor
Label          The disk Label editor
User Management Add user and group information
Console        Customize system console behavior
Time Zone      Set which time zone you're in
Media          Change the installation media type
Mouse          Configure your mouse
Networking     Configure additional network services
Security       Configure system security options
Startup        Configure system startup options
TTYs           Configure system ttys.
Options        View/Set various installation options
v(+)

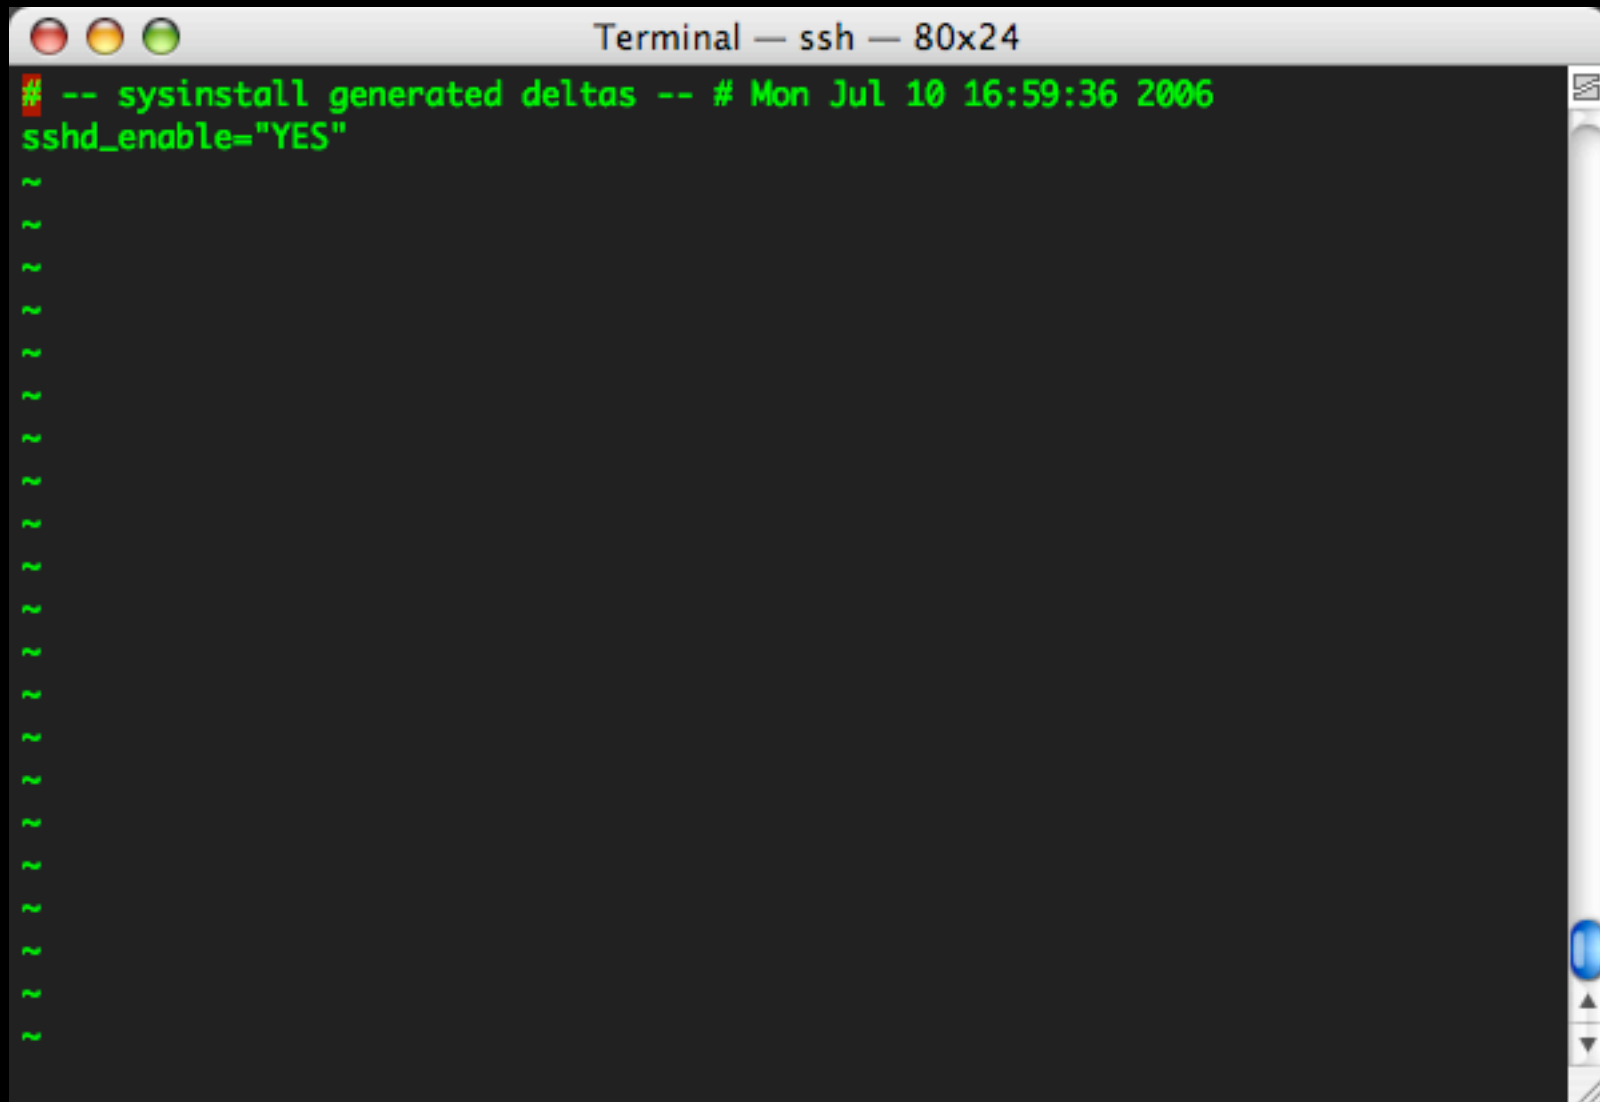
[ OK ]      Cancel

[ Press F1 for more information on these options ]

#
#
#
#
# vi /etc/rc.conf
```

# configure - call jailed sh

check rc.conf in jail


A screenshot of a terminal window titled "Terminal - ssh - 80x24". The terminal shows the output of a system boot process. The first line is a green prompt character followed by "-- sysinstall generated deltas -- # Mon Jul 10 16:59:36 2006". The second line is "sshd\_enable='YES'". This is followed by a vertical column of approximately 20 tilde (~) characters, indicating the progress of the boot process. The terminal window has a standard macOS-style title bar with three colored buttons (red, yellow, green) on the left and a scroll bar on the right.

```
Terminal - ssh - 80x24
# -- sysinstall generated deltas -- # Mon Jul 10 16:59:36 2006
sshd_enable="YES"
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
```






# configure



# configure





# configure - assign ip alias

(use ifconfig)

```
Terminal — ssh — 80x24
[root@chicken ~]# ifconfig bge0 inet alias 192.168.1.200/32
[root@chicken ~]# ifconfig
bge0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
    inet6 fe80::2e0:81ff:fe34:bf8c%bge0 prefixlen 64 scopeid 0x1
    inet 192.168.1.2 netmask 0xffffffff broadcast 192.168.1.255
    inet 192.168.1.200 netmask 0xffffffff broadcast 192.168.1.200
    ether 00:e0:81:34:bf:8c
    media: Ethernet autoselect (1000baseTX <full-duplex>)
    status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
    inet 127.0.0.1 netmask 0xff000000
[root@chicken ~]#
```

# configure - assign ip alias

(ip for the jail)

```
Terminal — ssh — 80x24
[root@chicken ~]# ifconfig bge0 inet alias 192.168.1.200/32
[root@chicken ~]# ifconfig
bge0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
    inet6 fe80::2e0:81ff:fe34:bf8c%bge0 prefixlen 64 scopeid 0x1
    inet 192.168.1.2 netmask 0xffffffff broadcast 192.168.1.255
    inet 192.168.1.200 netmask 0xffffffff broadcast 192.168.1.200
    ether 00:e0:81:34:bf:8c
    media: Ethernet autoselect (1000baseTX <full-duplex>)
    status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
    inet 127.0.0.1 netmask 0xff000000
[root@chicken ~]#
```

# configure - assign ip alias

(original ip for the host machine)

```
Terminal — ssh — 80x24
[root@chicken ~]# ifconfig bge0 inet alias 192.168.1.200/32
[root@chicken ~]# ifconfig
bge0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
  options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
  inet6 fe80::2e0:81ff:fe34:bf8c%bge0 prefixlen 64 scopeid 0x1
  inet 192.168.1.2 netmask 0xffffffff broadcast 192.168.1.255
  inet 192.168.1.200 netmask 0xffffffff broadcast 192.168.1.200
  ether 00:e0:81:34:bf:8c
  media: Ethernet autoselect (1000baseTX <full-duplex>)
  status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
  inet6 ::1 prefixlen 128
  inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
  inet 127.0.0.1 netmask 0xff000000
[root@chicken ~]#
```



# configure - call jailed sh

(analagous to booting a machine in su mode)

```
Terminal — ssh — 80x24

ether 00:e0:81:34:bf:8c
media: Ethernet autoselect (1000baseTX <full-duplex>)
status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
    inet 127.0.0.1 netmask 0xff000000
[root@chicken ~]# mount -t procfs proc /usr/local/jails/chick.diversaform.net/pr
oc
[root@chicken ~]# df -h
Filesystem      Size  Used Avail Capacity  Mounted on
/dev/da0s1a    496M   57M  399M    13%      /
devfs           1.0K   1.0K   0B    100%    /dev
/dev/da0s1e    496M   12K  456M     0%    /tmp
/dev/da0s1f    63G   2.0G   56G     3%    /usr
/dev/da0s1d   4.2G   39M   3.8G     1%    /var
devfs           1.0K   1.0K   0B    100%    /usr/local/jails/chick.diversaform
.net/dev
procfs          4.0K   4.0K   0B    100%    /usr/local/jails/chick.diversaform
.net/proc
[root@chicken ~]#
```


# configure - call jailed sh

(analagous to booting a machine in su mode)


```
Terminal — ssh — 80x24

ether 00:e0:81:34:bf:8c
media: Ethernet autoselect (1000baseTX <full-duplex>)
status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x6
    inet 127.0.0.1 netmask 0xff000000
[root@chicken ~]# mount -t procfs proc /usr/local/jails/chick.diversaform.net/pr
oc
[root@chicken ~]# df -h
Filesystem      Size  Used  Avail Capacity  Mounted on
/dev/da0s1a    496M   57M   399M    13%      /
devfs           1.0K   1.0K    0B   100%    /dev
/dev/da0s1e    496M   12K   456M    0%      /tmp
/dev/da0s1f     63G   2.0G   56G     3%      /usr
/dev/da0s1d    4.2G   39M   3.8G    1%      /var
devfs           1.0K   1.0K    0B   100%    /usr/local/jails/chick.diversaform
.net/dev
procfs          4.0K   4.0K    0B   100%    /usr/local/jails/chick.diversaform
.net/proc
[root@chicken ~]#
```

# preflight




# preflight



192.168.1.2

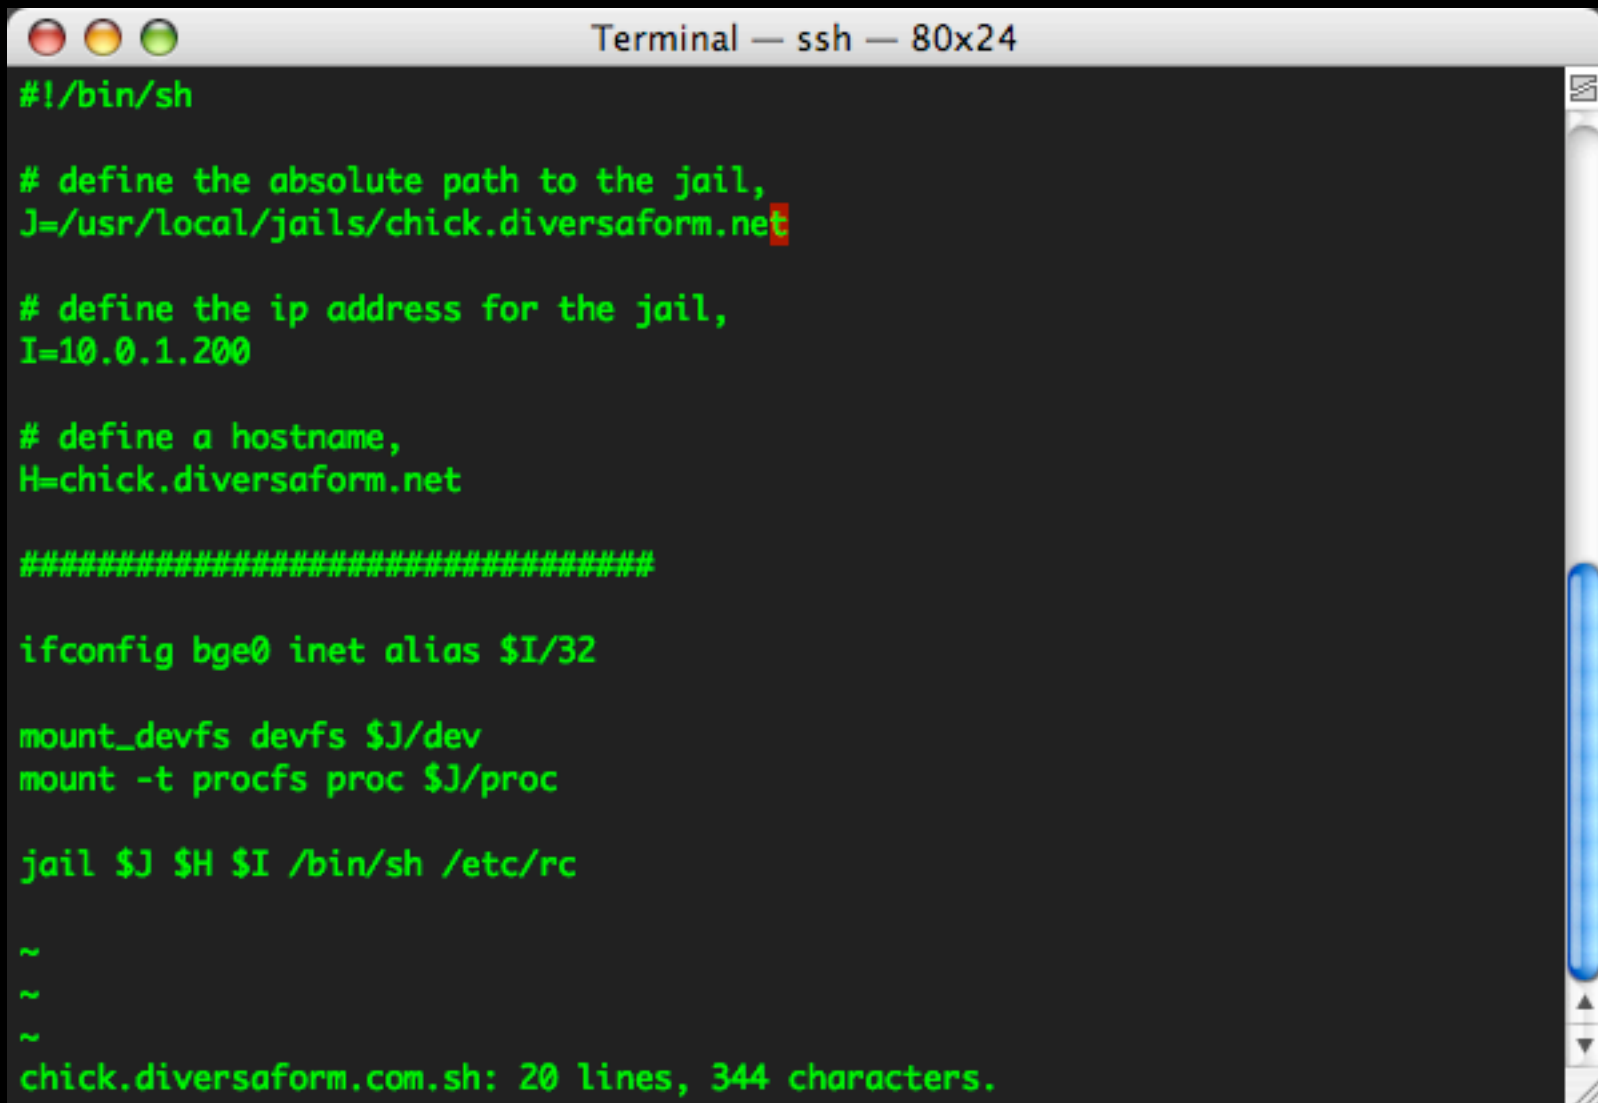
192.168.1.200

192.168.1.x



# start tangent! (script),

remember how I said rc.d is usually a bad idea?



```
Terminal — ssh — 80x24
#!/bin/sh

# define the absolute path to the jail,
J=/usr/local/jails/chick.diversaform.net

# define the ip address for the jail,
I=10.0.1.200

# define a hostname,
H=chick.diversaform.net

#####

ifconfig bge0 inet alias $I/32

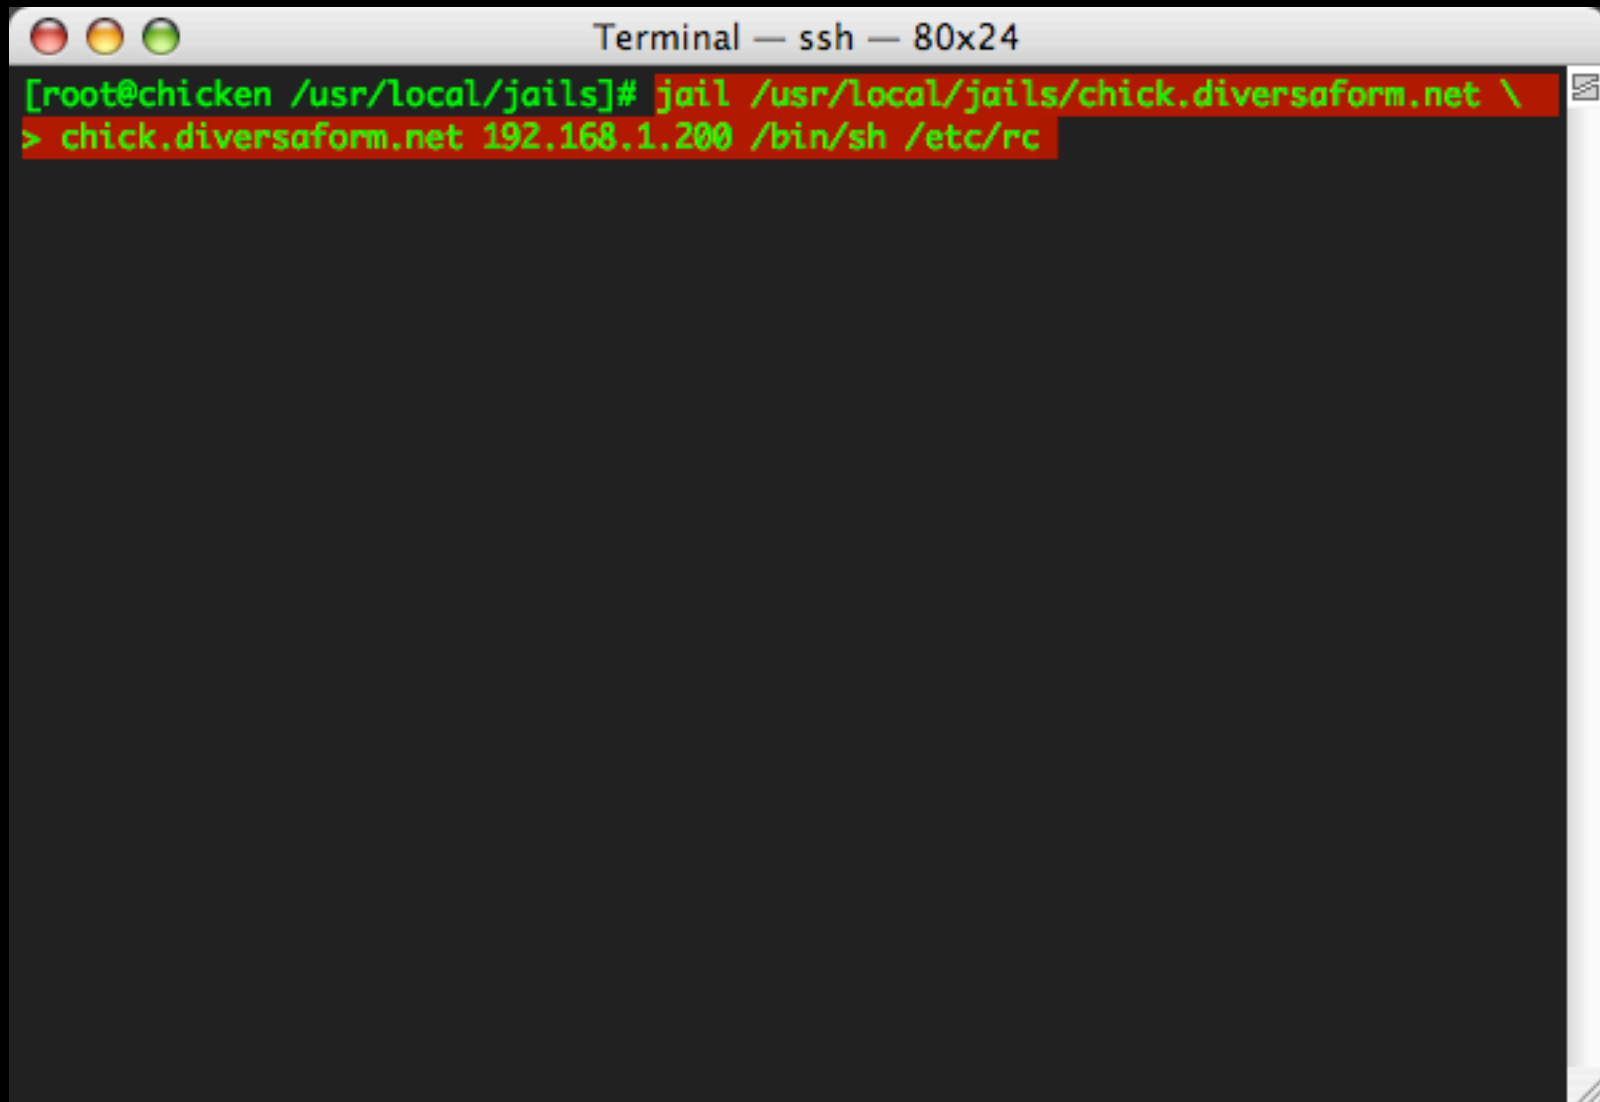
mount_devfs devfs $J/dev
mount -t procfs proc $J/proc

jail $J $H $I /bin/sh /etc/rc

~
~
~
chick.diversaform.com.sh: 20 lines, 344 characters.
```

# start!

we're gonna start the jail manually here....

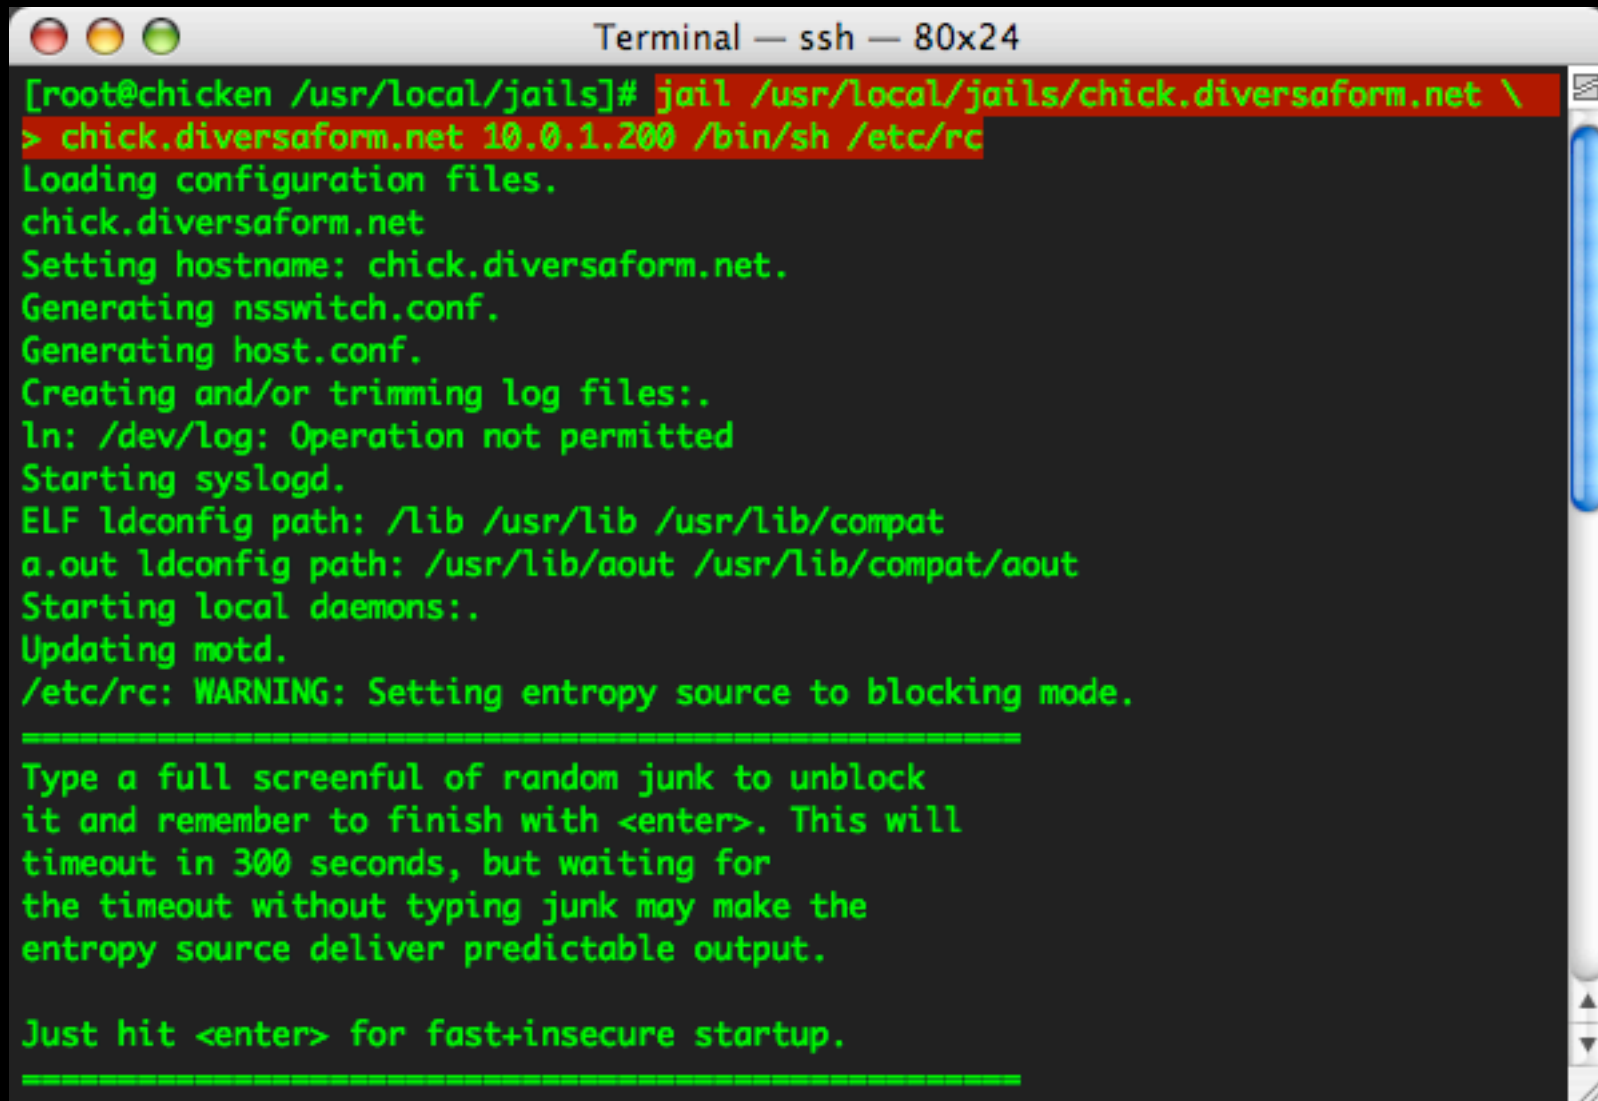
A terminal window titled "Terminal — ssh — 80x24" with three window control buttons (red, yellow, green) in the top-left corner. The terminal content is highlighted in red and shows the command: 

```
[root@chicken /usr/local/jails]# jail /usr/local/jails/chick.diversaform.net \  
> chick.diversaform.net 192.168.1.200 /bin/sh /etc/rc
```



# start!

we're gonna start the jail manually here....

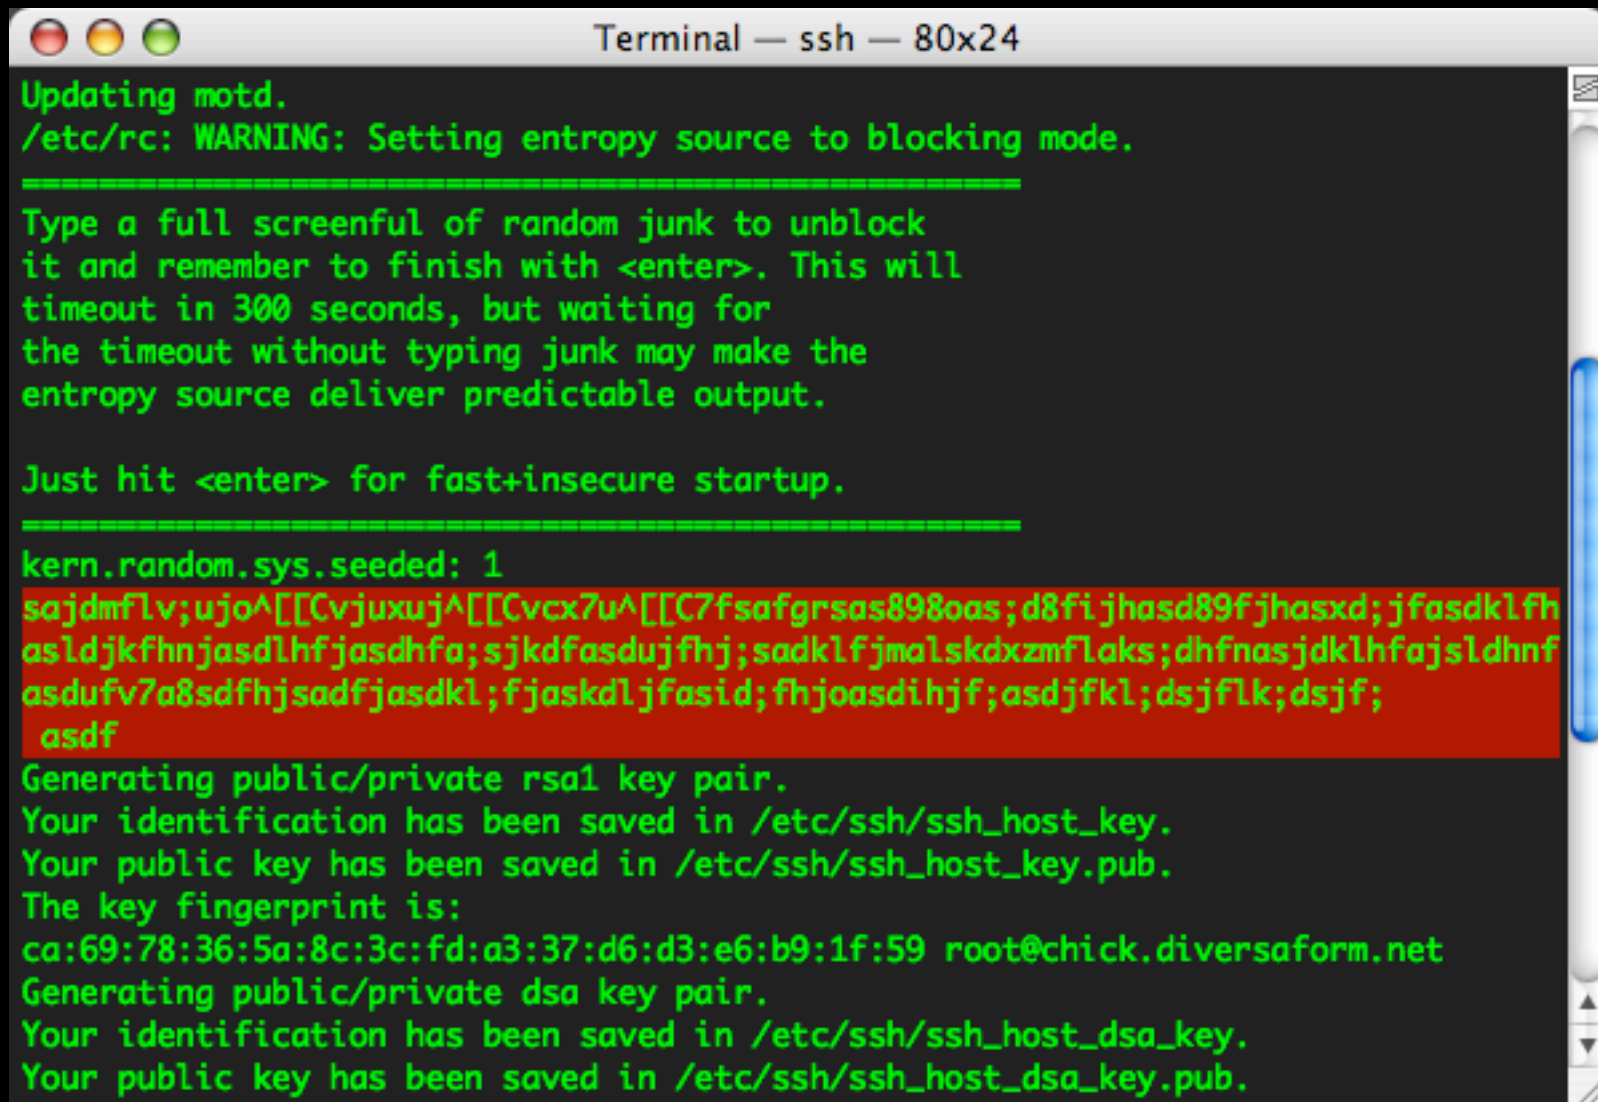
A terminal window titled "Terminal — ssh — 80x24" showing the manual startup of a jail. The user enters the command "jail /usr/local/jails/chick.diversaform.net \> chick.diversaform.net 10.0.1.200 /bin/sh /etc/rc". The output shows the jail loading configuration files, setting the hostname to "chick.diversaform.net", generating nsswitch.conf and host.conf, and starting local daemons. A warning message indicates that the entropy source is set to blocking mode. The terminal also displays instructions for unblocking the entropy source by typing random junk, and a note to hit <enter> for fast+insecure startup.

```
[root@chicken /usr/local/jails]# jail /usr/local/jails/chick.diversaform.net \
> chick.diversaform.net 10.0.1.200 /bin/sh /etc/rc
Loading configuration files.
chick.diversaform.net
Setting hostname: chick.diversaform.net.
Generating nsswitch.conf.
Generating host.conf.
Creating and/or trimming log files:.
ln: /dev/log: Operation not permitted
Starting syslogd.
ELF ldconfig path: /lib /usr/lib /usr/lib/compat
a.out ldconfig path: /usr/lib/aout /usr/lib/compat/aout
Starting local daemons:.
Updating motd.
/etc/rc: WARNING: Setting entropy source to blocking mode.
=====
Type a full screenful of random junk to unblock
it and remember to finish with <enter>. This will
timeout in 300 seconds, but waiting for
the timeout without typing junk may make the
entropy source deliver predictable output.

Just hit <enter> for fast+insecure startup.
=====
```

# start!

type some random junk to seed entropy,



```
Terminal — ssh — 80x24
Updating motd.
/etc/rc: WARNING: Setting entropy source to blocking mode.
=====
Type a full screenful of random junk to unblock
it and remember to finish with <enter>. This will
timeout in 300 seconds, but waiting for
the timeout without typing junk may make the
entropy source deliver predictable output.

Just hit <enter> for fast+insecure startup.
=====
kern.random.sys.seeded: 1
sajdmflv;ujo^[[Cvjuxuj^[[Cvcx7u^[[C7fsafgrsas898oas;d8fi jhasd89fjhasxd;jfasdklfh
asldjkhfnjasdlhfjasdhfa;sjkdfasdujfhj;sadklfjmałskdxzmfłaks;dhfnasjdklfajsl dhnf
asdufv7a8sdfhjsadfjasdkl;fjaskdljfasid;fhjoasdi hjf;asdjflk;dsjflk;dsjf;
asdf
Generating public/private rsa1 key pair.
Your identification has been saved in /etc/ssh/ssh_host_key.
Your public key has been saved in /etc/ssh/ssh_host_key.pub.
The key fingerprint is:
ca:69:78:36:5a:8c:3c:fd:a3:37:d6:d3:e6:b9:1f:59 root@chick.diversaform.net
Generating public/private dsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_dsa_key.
Your public key has been saved in /etc/ssh/ssh_host_dsa_key.pub.
```

# start!

## jail finished starting

```
Terminal — ssh — 80x24
asdf
Generating public/private rsa1 key pair.
Your identification has been saved in /etc/ssh/ssh_host_key.
Your public key has been saved in /etc/ssh/ssh_host_key.pub.
The key fingerprint is:
ca:69:78:36:5a:8c:3c:fd:a3:37:d6:d3:e6:b9:1f:59 root@chick.diversaform.net
Generating public/private dsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_dsa_key.
Your public key has been saved in /etc/ssh/ssh_host_dsa_key.pub.
The key fingerprint is:
d1:14:6e:3a:5d:e2:c7:c3:eb:4b:ce:4e:ee:b5:86:8b root@chick.diversaform.net
Generating public/private rsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_rsa_key.
Your public key has been saved in /etc/ssh/ssh_host_rsa_key.pub.
The key fingerprint is:
fc:8e:0c:a5:79:95:d3:d7:05:14:cb:86:f7:45:1f:a9 root@chick.diversaform.net
Starting sshd.
sendmail_submit: /etc/mail/aliases.db not present, generating
/etc/mail/aliases: 27 aliases, longest 10 bytes, 275 bytes total
Starting cron.
Local package initialization:.

Mon Jul 10 18:42:07 EDT 2006
[root@chicken /usr/local/jails]#
```

# running

jls(8) lists running jails, gives a jail ID

```
Terminal — ssh — 80x24
The key fingerprint is:
ca:69:78:36:5a:8c:3c:fd:a3:37:d6:d3:e6:b9:1f:59 root@chick.diversaform.net
Generating public/private dsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_dsa_key.
Your public key has been saved in /etc/ssh/ssh_host_dsa_key.pub.
The key fingerprint is:
d1:14:6e:3a:5d:e2:c7:c3:eb:4b:ce:4e:ee:b5:86:8b root@chick.diversaform.net
Generating public/private rsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_rsa_key.
Your public key has been saved in /etc/ssh/ssh_host_rsa_key.pub.
The key fingerprint is:
fc:8e:0c:a5:79:95:d3:d7:05:14:cb:86:f7:45:1f:a9 root@chick.diversaform.net
Starting sshd.
sendmail_submit: /etc/mail/aliases.db not present, generating
/etc/mail/aliases: 27 aliases, longest 10 bytes, 275 bytes total
Starting cron.
Local package initialization:.

Mon Jul 10 18:42:07 EDT 2006
[root@chicken /usr/local/jails]# jls
  JID  IP Address      Hostname                Path
   5   10.0.1.200     chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]#
```

# using the jail

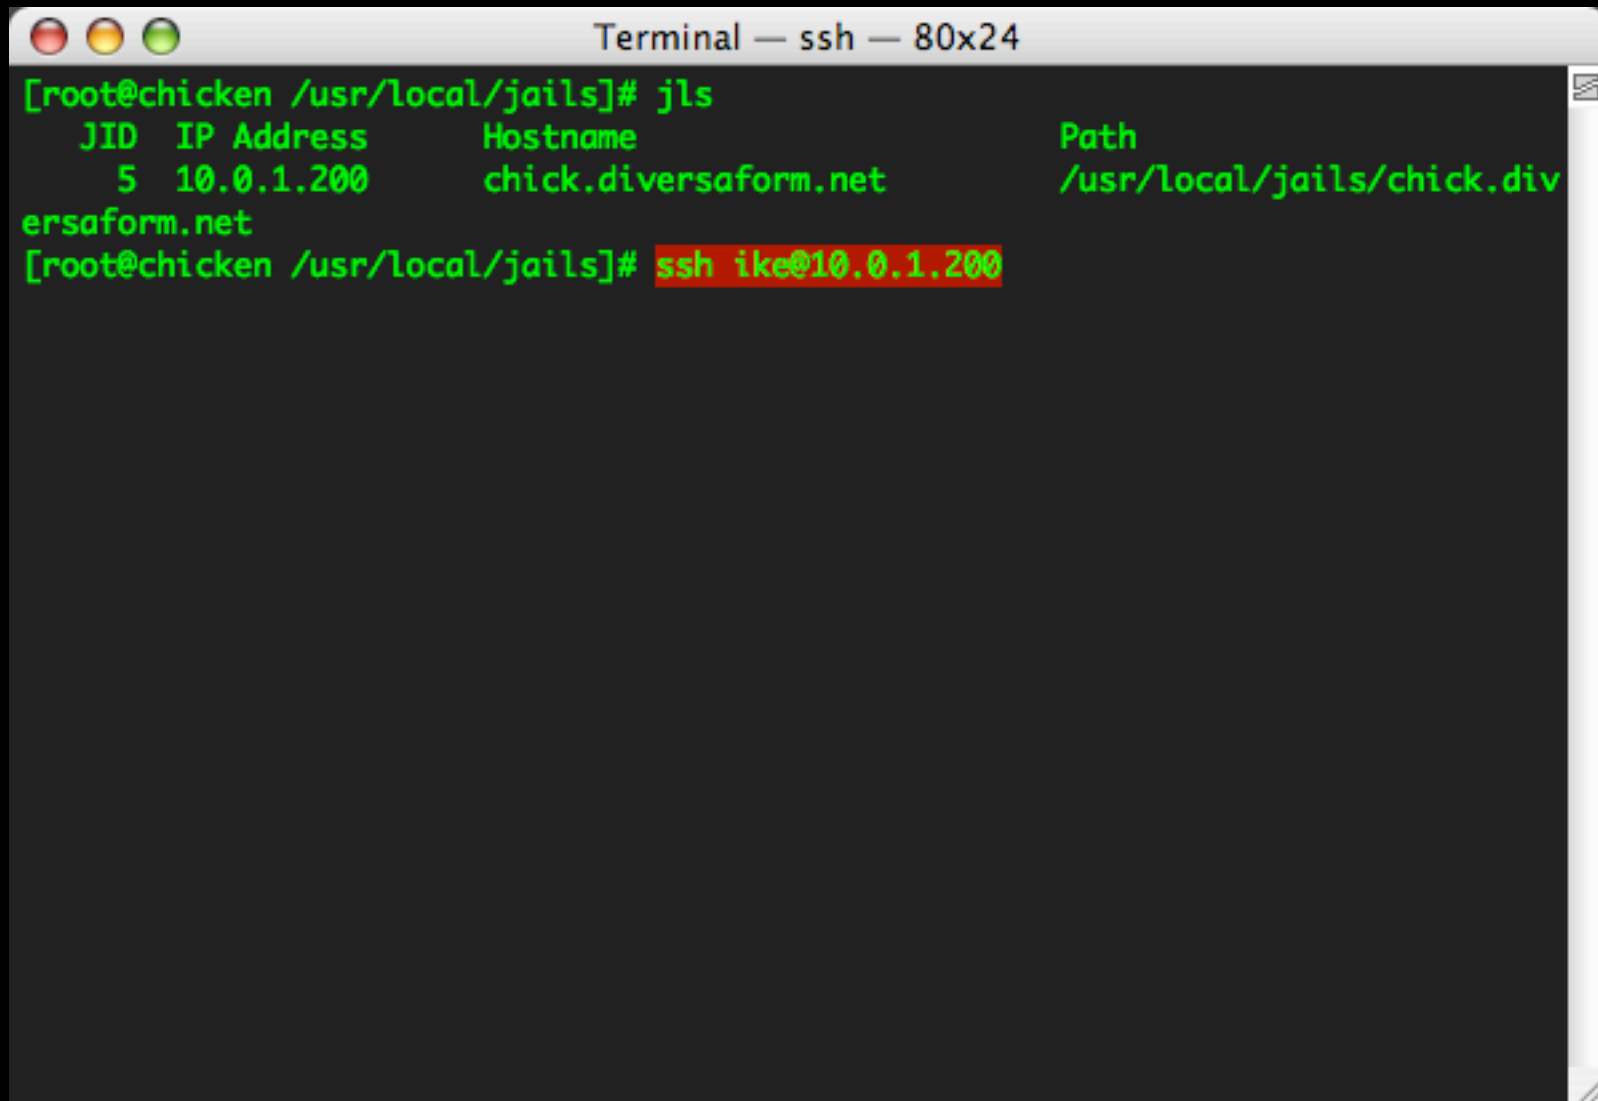
ssh into the jail, treat it like a server.

```
Terminal — ssh — 80x24
Your identification has been saved in /etc/ssh/ssh_host_dsa_key.
Your public key has been saved in /etc/ssh/ssh_host_dsa_key.pub.
The key fingerprint is:
d1:14:6e:3a:5d:e2:c7:c3:eb:4b:ce:4e:ee:b5:86:8b root@chick.diversaform.net
Generating public/private rsa key pair.
Your identification has been saved in /etc/ssh/ssh_host_rsa_key.
Your public key has been saved in /etc/ssh/ssh_host_rsa_key.pub.
The key fingerprint is:
fc:8e:0c:a5:79:95:d3:d7:05:14:cb:86:f7:45:1f:a9 root@chick.diversaform.net
Starting sshd.
sendmail_submit: /etc/mail/aliases.db not present, generating
/etc/mail/aliases: 27 aliases, longest 10 bytes, 275 bytes total
Starting cron.
Local package initialization:.

Mon Jul 10 18:42:07 EDT 2006
[root@chicken /usr/local/jails]# jls
      JID  IP Address      Hostname                Path
      5   10.0.1.200     chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# ssh 10.0.1.200
The authenticity of host '10.0.1.200 (10.0.1.200)' can't be established.
DSA key fingerprint is d1:14:6e:3a:5d:e2:c7:c3:eb:4b:ce:4e:ee:b5:86:8b.
Are you sure you want to continue connecting (yes/no)?
```

# using the jail

ssh into the jail, treat it like a server.



A terminal window titled "Terminal — ssh — 80x24" showing the following output:

```
[root@chicken /usr/local/jails]# jls
```

JID	IP Address	Hostname	Path
5	10.0.1.200	chick.diversaform.net	/usr/local/jails/chick.diversaform.net

```
[root@chicken /usr/local/jails]# ssh ike@10.0.1.200
```



# using the jail

ssh into the jail, treat it like a server.

```
Terminal — ssh — 80x24
[root@chicken /usr/local/jails]# jls
  JID  IP Address      Hostname                Path
   5   10.0.1.200     chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# ssh ike@10.0.1.200
The authenticity of host '10.0.1.200 (10.0.1.200)' can't be established.
DSA key fingerprint is d1:14:6e:3a:5d:e2:c7:c3:eb:4b:ce:4e:ee:b5:86:8b.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.0.1.200' (DSA) to the list of known hosts.
Password:
Copyright (c) 1980, 1983, 1986, 1988, 1990, 1991, 1993, 1994
  The Regents of the University of California.  All rights reserved.

FreeBSD 6.1-RELEASE-p3 (DIVERSAFORM_NET-4-SMP) #3: Mon Jul 10 11:43:08 EDT 2006

Welcome to FreeBSD!

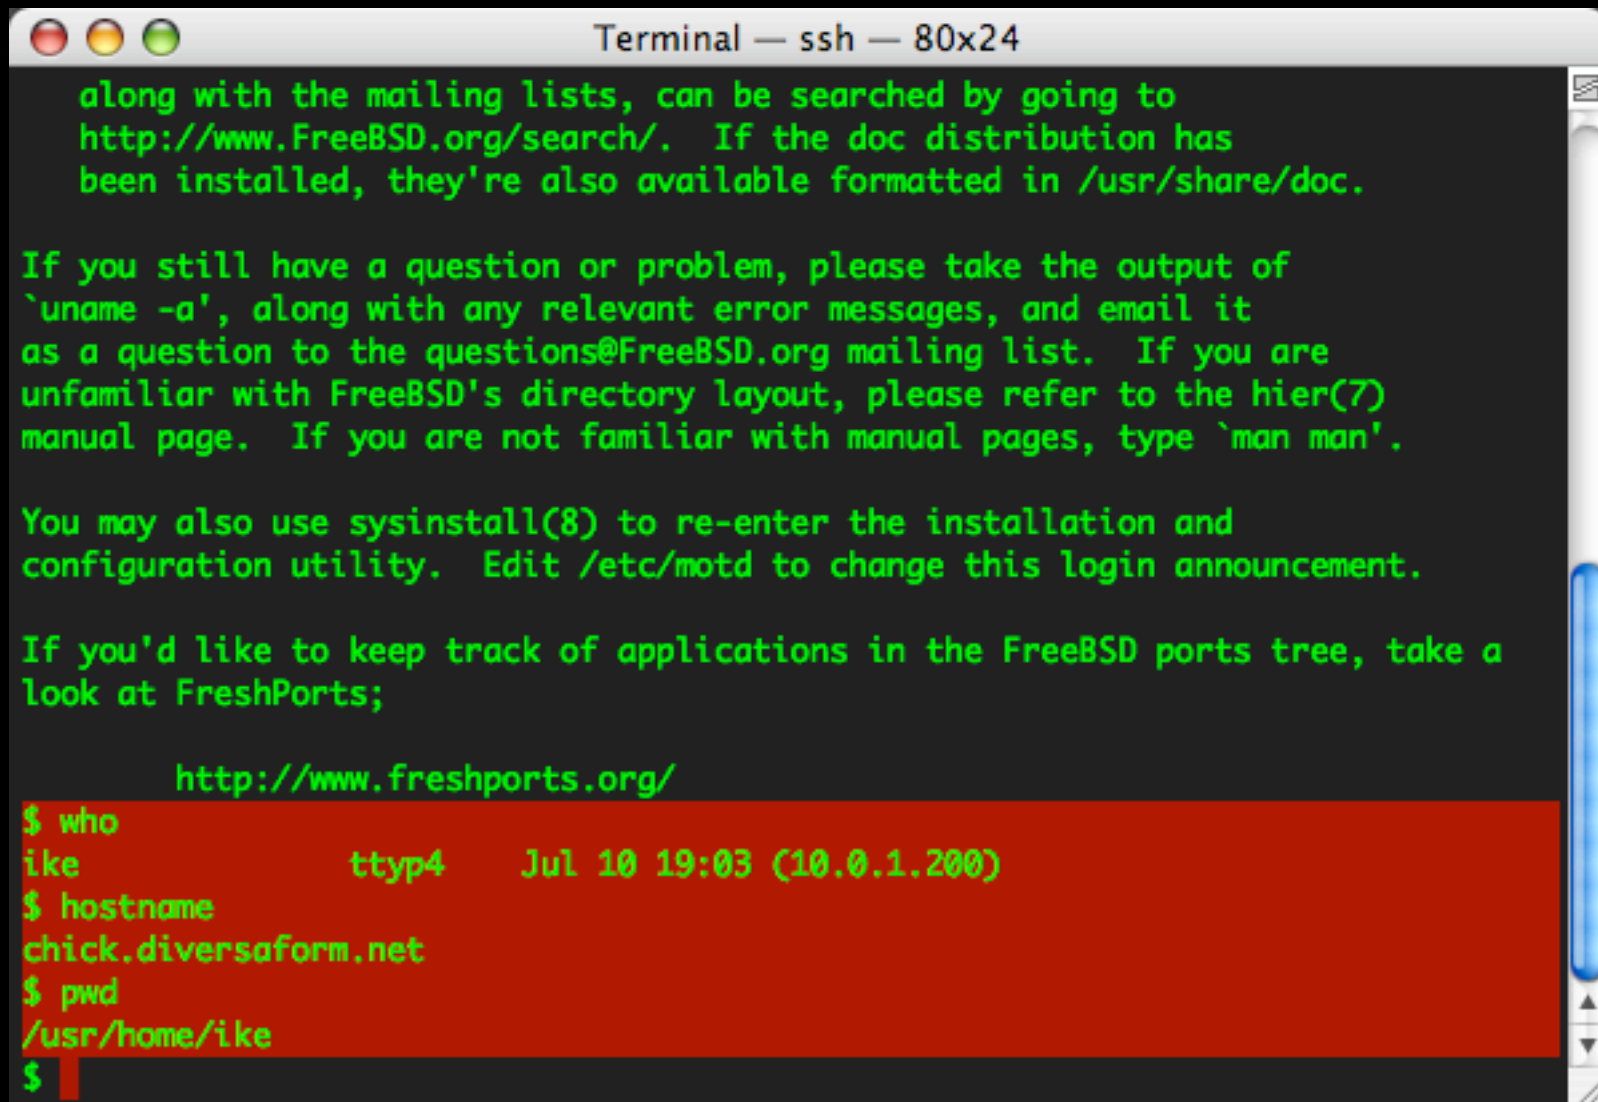
Before seeking technical support, please use the following resources:

o Security advisories and updated errata information for all releases are
  at http://www.FreeBSD.org/releases/ - always consult the ERRATA section
  for your release first as it's updated frequently.

o The Handbook and FAQ documents are at http://www.FreeBSD.org/ and,
```

# inside the jail

## just like any new server



```
Terminal — ssh — 80x24

along with the mailing lists, can be searched by going to
http://www.FreeBSD.org/search/.  If the doc distribution has
been installed, they're also available formatted in /usr/share/doc.

If you still have a question or problem, please take the output of
`uname -a`, along with any relevant error messages, and email it
as a question to the questions@FreeBSD.org mailing list.  If you are
unfamiliar with FreeBSD's directory layout, please refer to the hier(7)
manual page.  If you are not familiar with manual pages, type `man man`.

You may also use sysinstall(8) to re-enter the installation and
configuration utility.  Edit /etc/motd to change this login announcement.

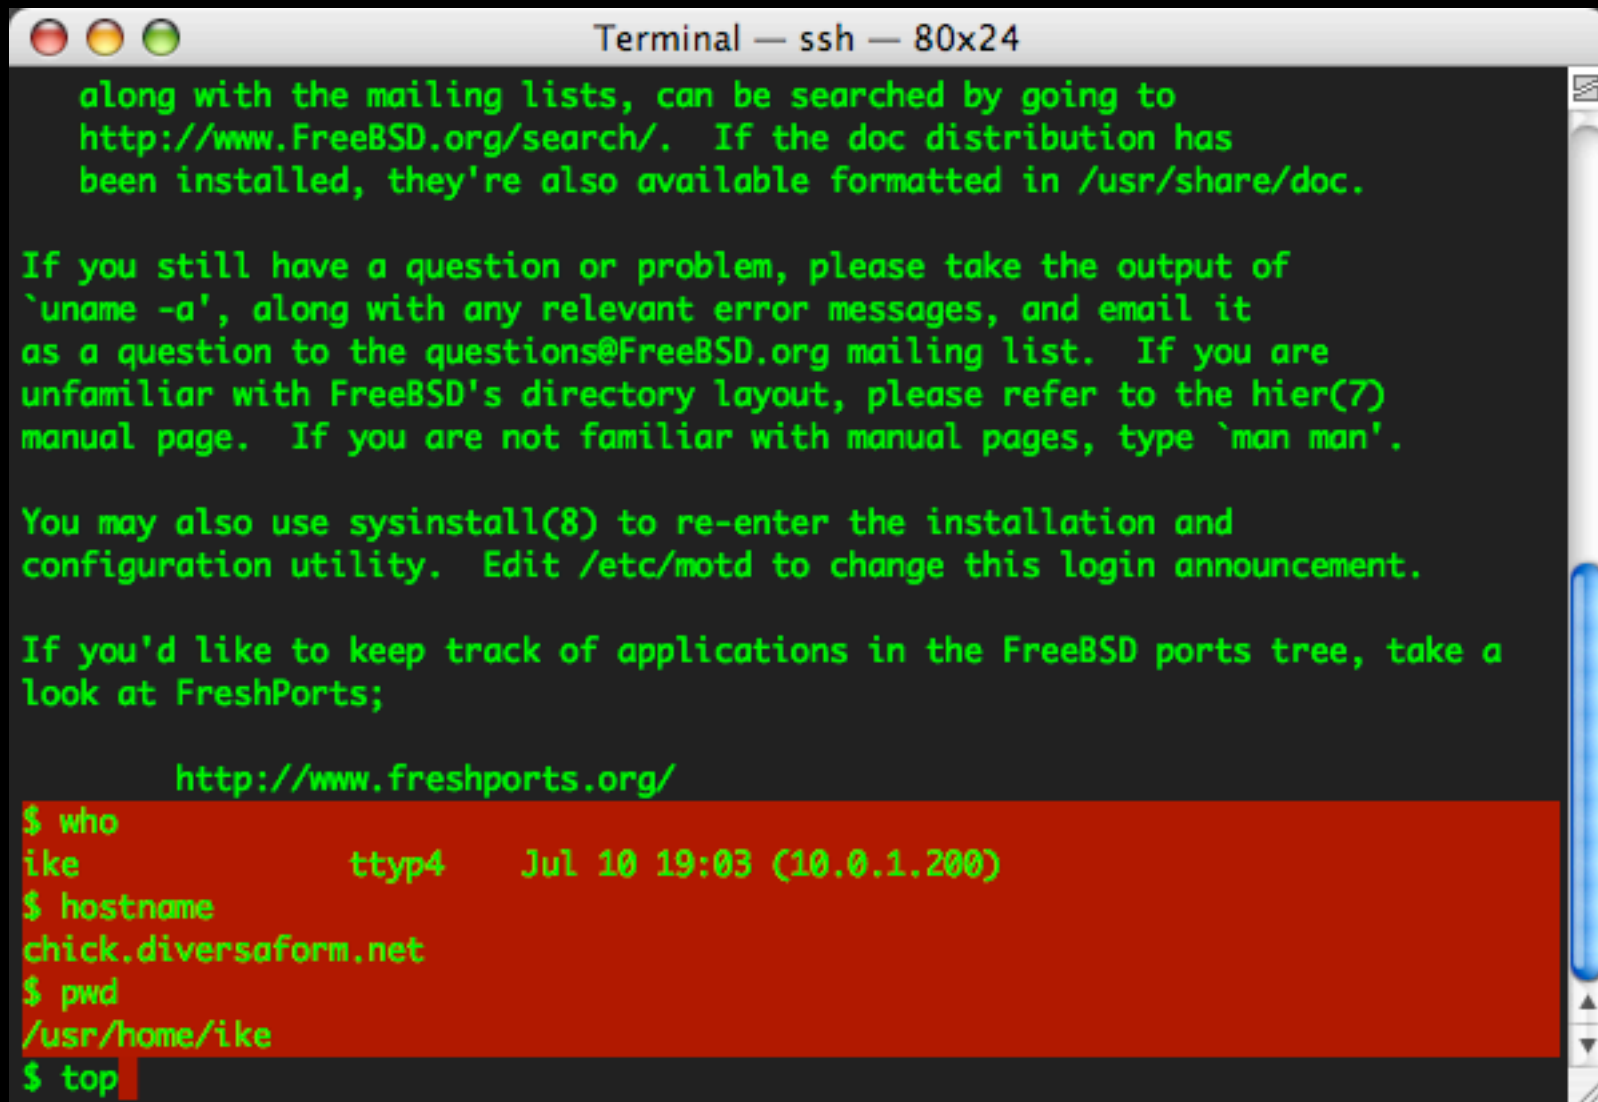
If you'd like to keep track of applications in the FreeBSD ports tree, take a
look at FreshPorts;

    http://www.freshports.org/

$ who
ike                ttyp4      Jul 10 19:03 (10.0.1.200)
$ hostname
chick.diversaform.net
$ pwd
/usr/home/ike
$
```

# inside the jail

## just like any new server



```
Terminal — ssh — 80x24

along with the mailing lists, can be searched by going to
http://www.FreeBSD.org/search/.  If the doc distribution has
been installed, they're also available formatted in /usr/share/doc.

If you still have a question or problem, please take the output of
`uname -a`, along with any relevant error messages, and email it
as a question to the questions@FreeBSD.org mailing list.  If you are
unfamiliar with FreeBSD's directory layout, please refer to the hier(7)
manual page.  If you are not familiar with manual pages, type `man man`.

You may also use sysinstall(8) to re-enter the installation and
configuration utility.  Edit /etc/motd to change this login announcement.

If you'd like to keep track of applications in the FreeBSD ports tree, take a
look at FreshPorts;

    http://www.freshports.org/

$ who
ike          ttyp4      Jul 10 19:03 (10.0.1.200)
$ hostname
chick.diversaform.net
$ pwd
/usr/home/ike
$ top
```

# inside the jail

just like any new server

```
Terminal — ssh — 80x24
last pid: 1237; load averages: 0.00, 0.00, 0.00 up 0+02:26:54 19:06:22
9 processes: 1 running, 8 sleeping
CPU states: 0.0% user, 0.0% nice, 0.0% system, 0.0% interrupt, 100% idle
Mem: 14M Active, 11M Inact, 26M Wired, 17M Buf, 3334M Free
Swap: 4096M Total, 4096M Free

```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	C	TIME	WCPU	COMMAND
1222	root	1	4	0	6104K	3108K	sbwait	0	0:00	0.00%	sshd
1095	root	1	96	0	3420K	2828K	select	0	0:00	0.00%	sendmail
1226	ike	1	8	0	1680K	1364K	wait	0	0:00	0.00%	sh
1225	ike	1	96	0	6080K	3124K	select	0	0:00	0.00%	sshd
1024	root	1	96	0	1300K	948K	select	0	0:00	0.00%	syslogd
1105	root	1	8	0	1312K	1044K	nanslp	0	0:00	0.00%	cron
1237	ike	1	96	0	2152K	1416K	CPU0	0	0:00	0.00%	top
1088	root	1	96	0	3356K	2828K	select	0	0:00	0.00%	sshd
1099	smsp	1	20	0	3300K	2716K	pause	0	0:00	0.00%	sendmail

# inside the jail

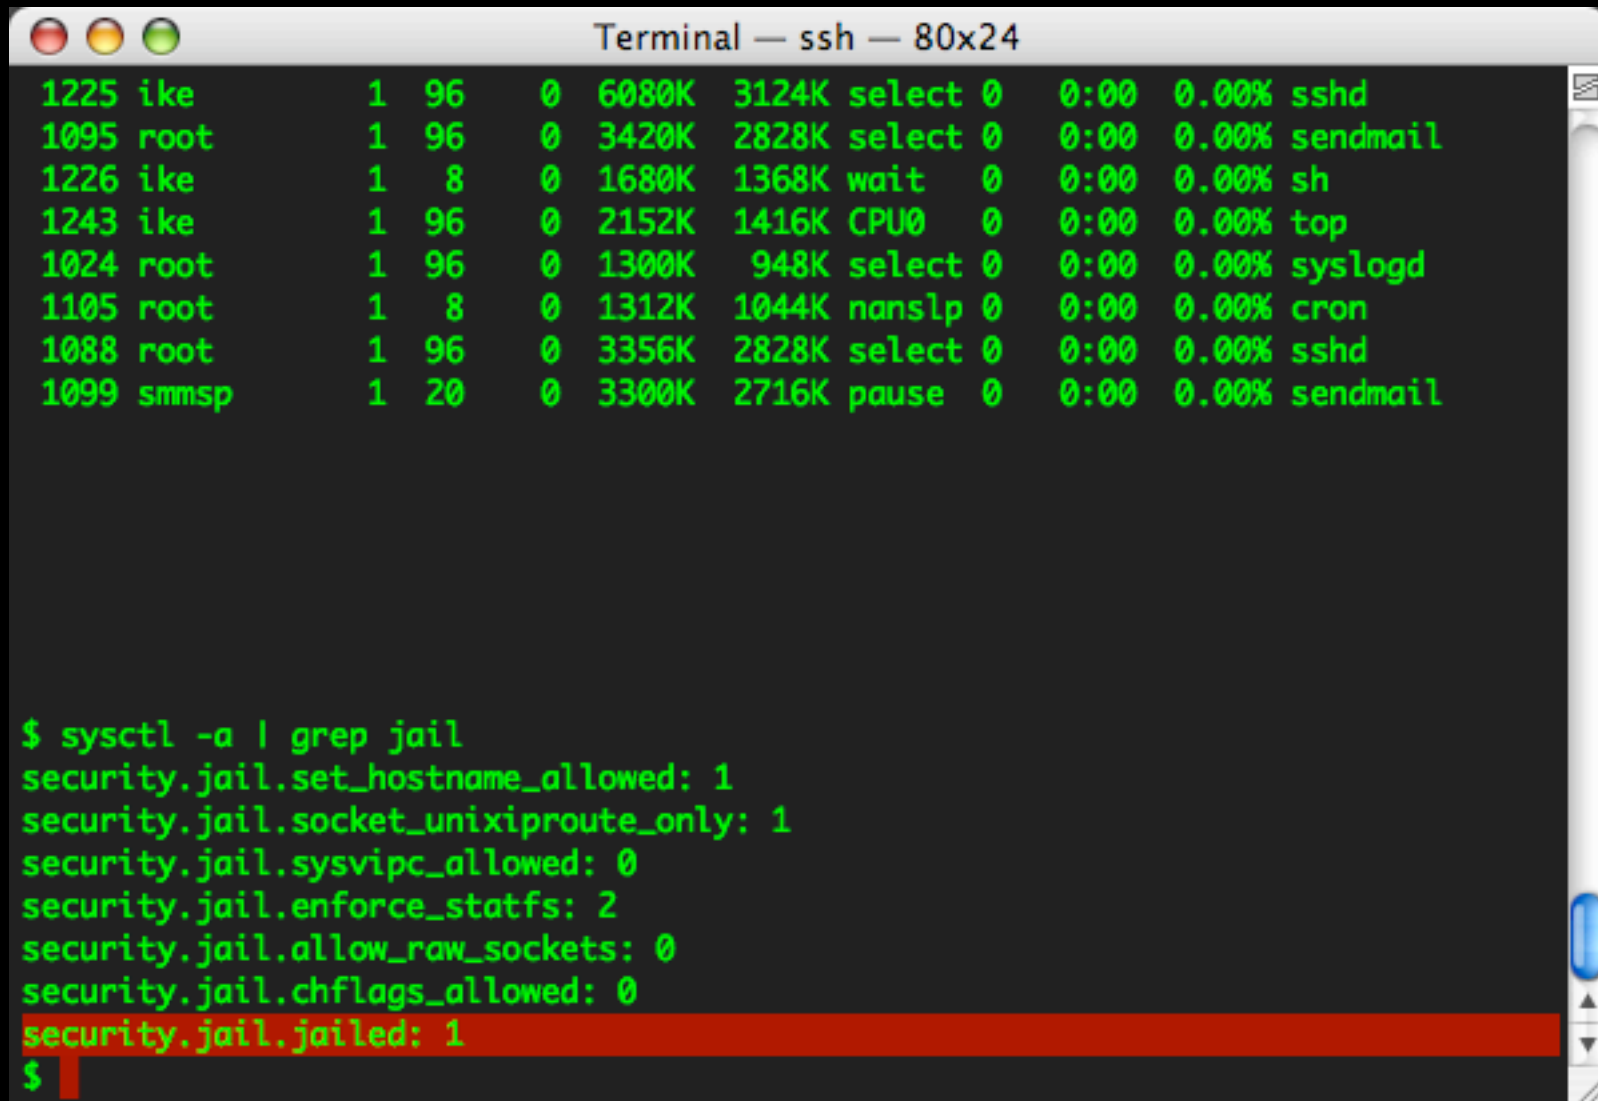
you have root!

```
Terminal — ssh — 80x24
$ hostname
chick.diversaform.net
$ ifconfig
bge0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      options=1b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING>
      inet 10.0.1.200 netmask 0xffffffff broadcast 10.0.1.200
      ether 00:e0:81:34:bf:8c
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
bge1: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      options=1b<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING>
      ether 00:e0:81:34:bf:8d
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
plip0: flags=108810<POINTOPOINT,SIMPLEX,MULTICAST,NEEDSGIANT> mtu 1500
pflog0: flags=0<> mtu 33208
pfsync0: flags=0<> mtu 2020
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
$ su
Password:
chick# whoami
root
chick#
```



# inside the jail

how do you know you are inside a jail?



The image shows a terminal window titled "Terminal — ssh — 80x24". The terminal displays a list of processes in a table format, followed by the output of the command `$ sysctl -a | grep jail`. The output shows several sysctl variables related to jail security, with `security.jail.jailed: 1` highlighted in red.

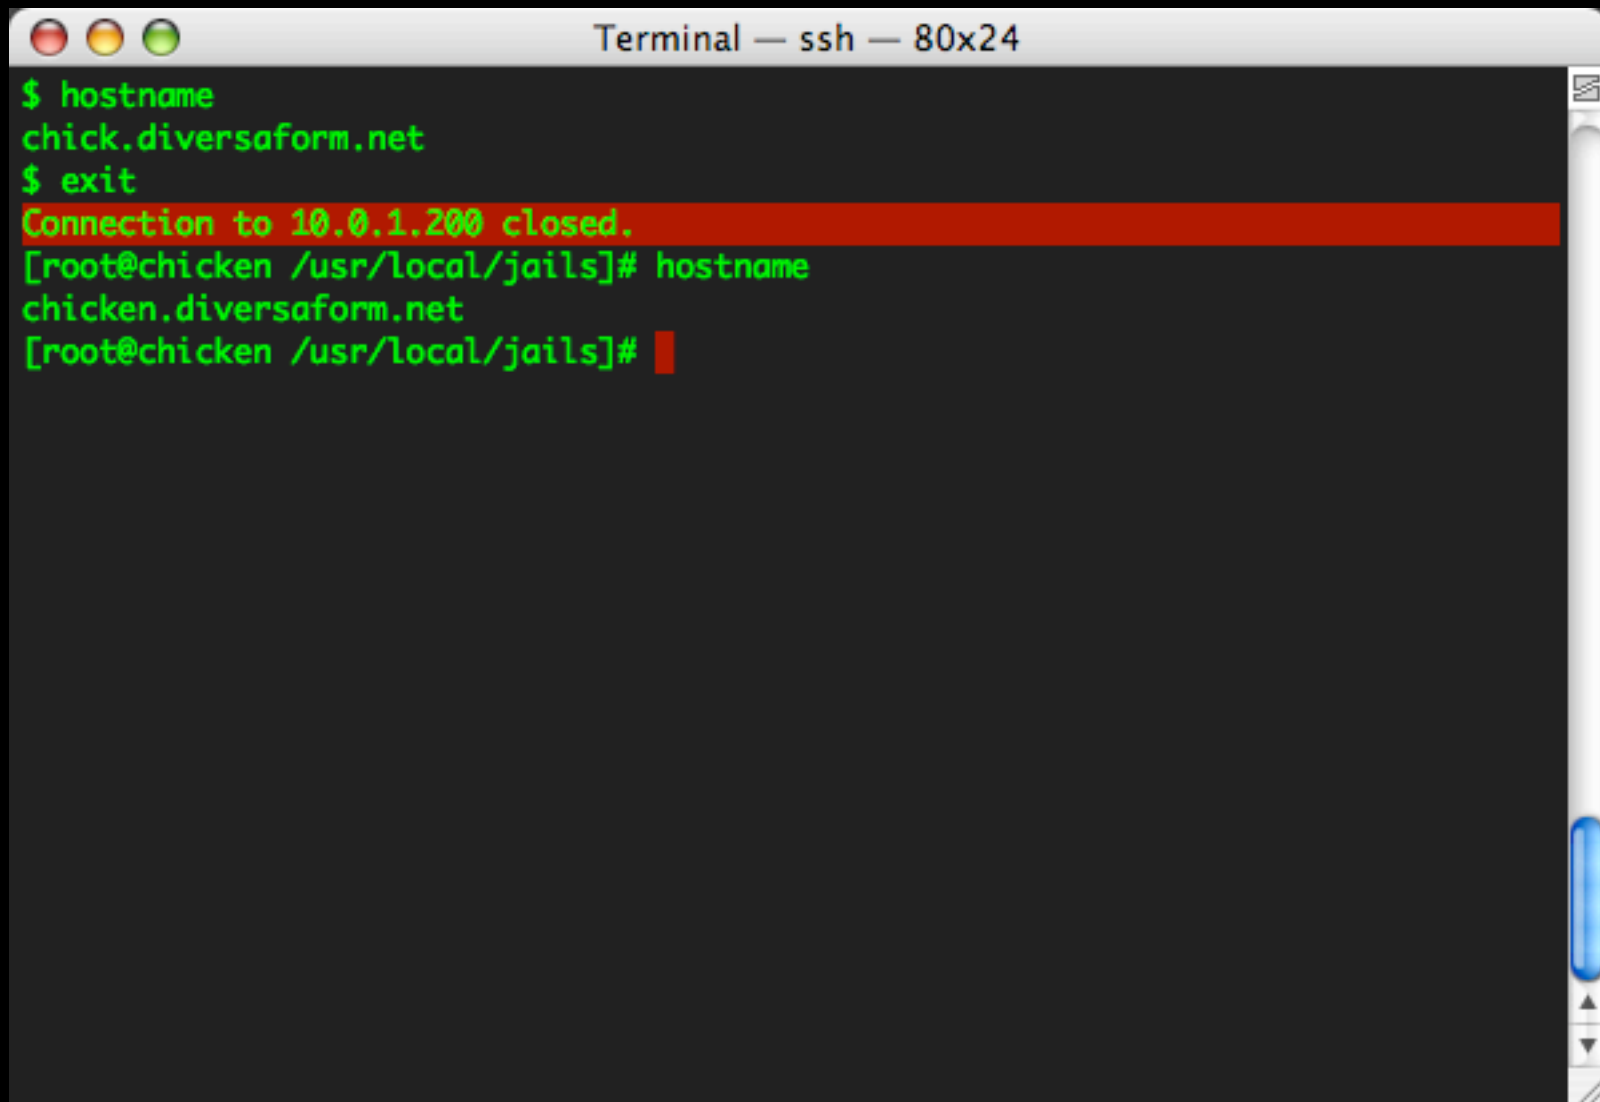
```
1225 ike      1  96    0  6080K  3124K select 0    0:00  0.00% sshd
1095 root     1  96    0  3420K  2828K select 0    0:00  0.00% sendmail
1226 ike      1   8     0  1680K  1368K wait   0    0:00  0.00% sh
1243 ike      1  96    0  2152K  1416K CPU0   0    0:00  0.00% top
1024 root     1  96    0  1300K   948K select 0    0:00  0.00% syslogd
1105 root     1   8     0  1312K  1044K nanslp 0    0:00  0.00% cron
1088 root     1  96    0  3356K  2828K select 0    0:00  0.00% sshd
1099 smmsp    1  20    0  3300K  2716K pause  0    0:00  0.00% sendmail

$ sysctl -a | grep jail
security.jail.set_hostname_allowed: 1
security.jail.socket_unixiproute_only: 1
security.jail.sysvipc_allowed: 0
security.jail.enforce_statfs: 2
security.jail.allow_raw_sockets: 0
security.jail.chflags_allowed: 0
security.jail.jailed: 1
$
```



# stop and start jail

exit the jail, (ssh)



```
Terminal — ssh — 80x24
$ hostname
chick.diversaform.net
$ exit
Connection to 10.0.1.200 closed.
[root@chicken /usr/local/jails]# hostname
chicken.diversaform.net
[root@chicken /usr/local/jails]#
```

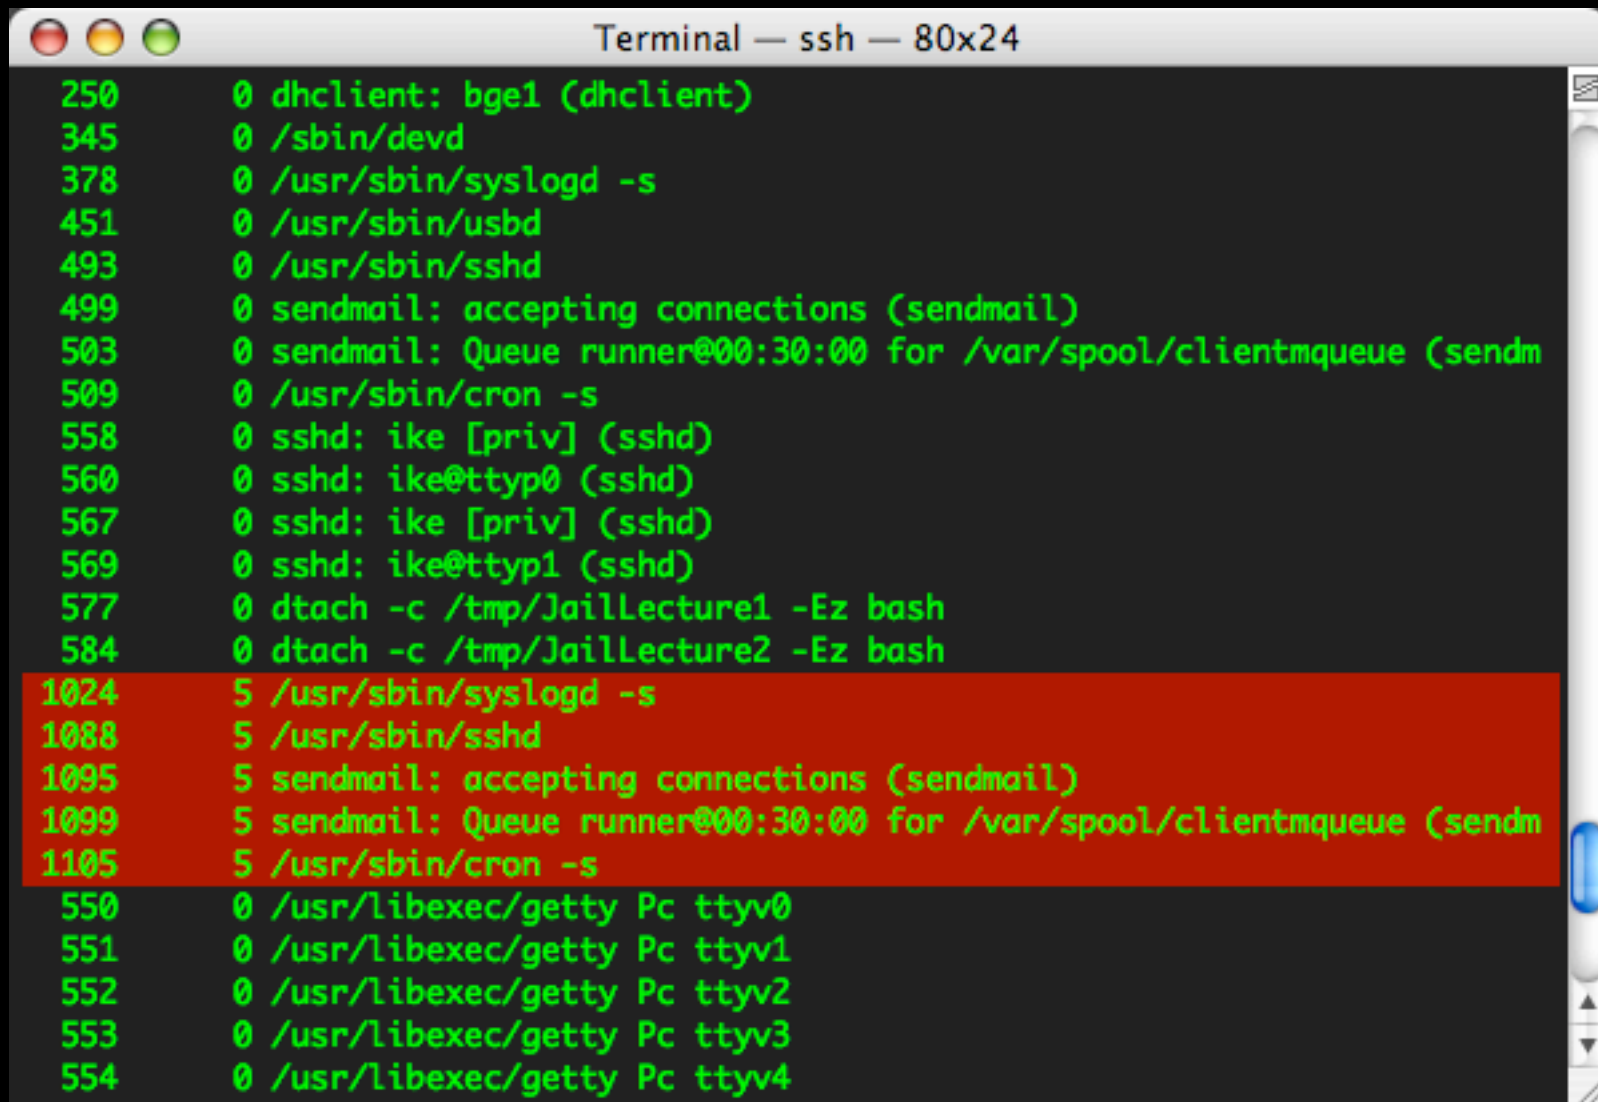
# stop and start jail

look at jailed processes (man page goodies)

```
Terminal — ssh — 80x24
5 10.0.1.200 chick.diversaform.net /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# ps ax -o pid,jid,args
PID    JID  COMMAND
  0      0  [swapper]
  1      0  /sbin/init --
  2      0  [g_event]
  3      0  [g_up]
  4      0  [g_down]
  5      0  [kqueue taskq]
  6      0  [thread taskq]
  7      0  [acpi_task0]
  8      0  [acpi_task1]
  9      0  [acpi_task2]
 10      0  [ktrace]
 11      0  [idle: cpu1]
 12      0  [idle: cpu0]
 13      0  [swi1: net]
 14      0  [swi4: clock sio]
 15      0  [swi3: vm]
 16      0  [yarrow]
 17      0  [swi2: cambio]
 18      0  [swi6: task queue]
 19      0  [swi6: +]
```

# stop and start jail

look at jailed processes (man page goodies)



```
Terminal — ssh — 80x24
250    0 dhclient: bge1 (dhclient)
345    0 /sbin/devd
378    0 /usr/sbin/syslogd -s
451    0 /usr/sbin/usbd
493    0 /usr/sbin/sshd
499    0 sendmail: accepting connections (sendmail)
503    0 sendmail: Queue runner@00:30:00 for /var/spool/clientmqueue (sendm
509    0 /usr/sbin/cron -s
558    0 sshd: ike [priv] (sshd)
560    0 sshd: ike@tty0 (sshd)
567    0 sshd: ike [priv] (sshd)
569    0 sshd: ike@tty1 (sshd)
577    0 dtach -c /tmp/JailLecture1 -Ez bash
584    0 dtach -c /tmp/JailLecture2 -Ez bash
1024   5 /usr/sbin/syslogd -s
1088   5 /usr/sbin/sshd
1095   5 sendmail: accepting connections (sendmail)
1099   5 sendmail: Queue runner@00:30:00 for /var/spool/clientmqueue (sendm
1105   5 /usr/sbin/cron -s
550    0 /usr/libexec/getty Pc ttyv0
551    0 /usr/libexec/getty Pc ttyv1
552    0 /usr/libexec/getty Pc ttyv2
553    0 /usr/libexec/getty Pc ttyv3
554    0 /usr/libexec/getty Pc ttyv4
```

# stop and start jail

use killall with -j flag

```
Terminal — ssh — 80x24
555    0 /usr/libexec/getty Pc ttyv5
556    0 /usr/libexec/getty Pc ttyv6
557    0 /usr/libexec/getty Pc ttyv7
199    0 dhclient: bge0 [priv] (dhclient)
230    0 dhclient: bge1 [priv] (dhclient)
561    0 -bash (bash)
576    0 dtach -c /tmp/JailLecture1 -Ez bash
570    0 -bash (bash)
583    0 dtach -c /tmp/JailLecture2 -Ez bash
578    0 bash
757    0 bash
1471   0 ps ax -o pid,jid,args
585    0 bash
587    0 bash
[root@chicken /usr/local/jails]# jls
      JID  IP Address      Hostname                Path
      5   10.0.1.200      chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# killall -j 5
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/p
roc
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/d
ev
[root@chicken /usr/local/jails]#
```

# stop and start jail

watch out for stacking mount points!

```
Terminal — ssh — 80x24
555    0 /usr/libexec/getty Pc ttyv5
556    0 /usr/libexec/getty Pc ttyv6
557    0 /usr/libexec/getty Pc ttyv7
199    0 dhclient: bge0 [priv] (dhclient)
230    0 dhclient: bge1 [priv] (dhclient)
561    0 -bash (bash)
576    0 dtach -c /tmp/JailLecture1 -Ez bash
570    0 -bash (bash)
583    0 dtach -c /tmp/JailLecture2 -Ez bash
578    0 bash
757    0 bash
1471   0 ps ax -o pid,jid,args
585    0 bash
587    0 bash
[root@chicken /usr/local/jails]# jls
      JID  IP Address      Hostname                Path
      5   10.0.1.200      chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# killall -j 5
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/p
roc
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/d
ev
[root@chicken /usr/local/jails]#
```

# stop and start jail

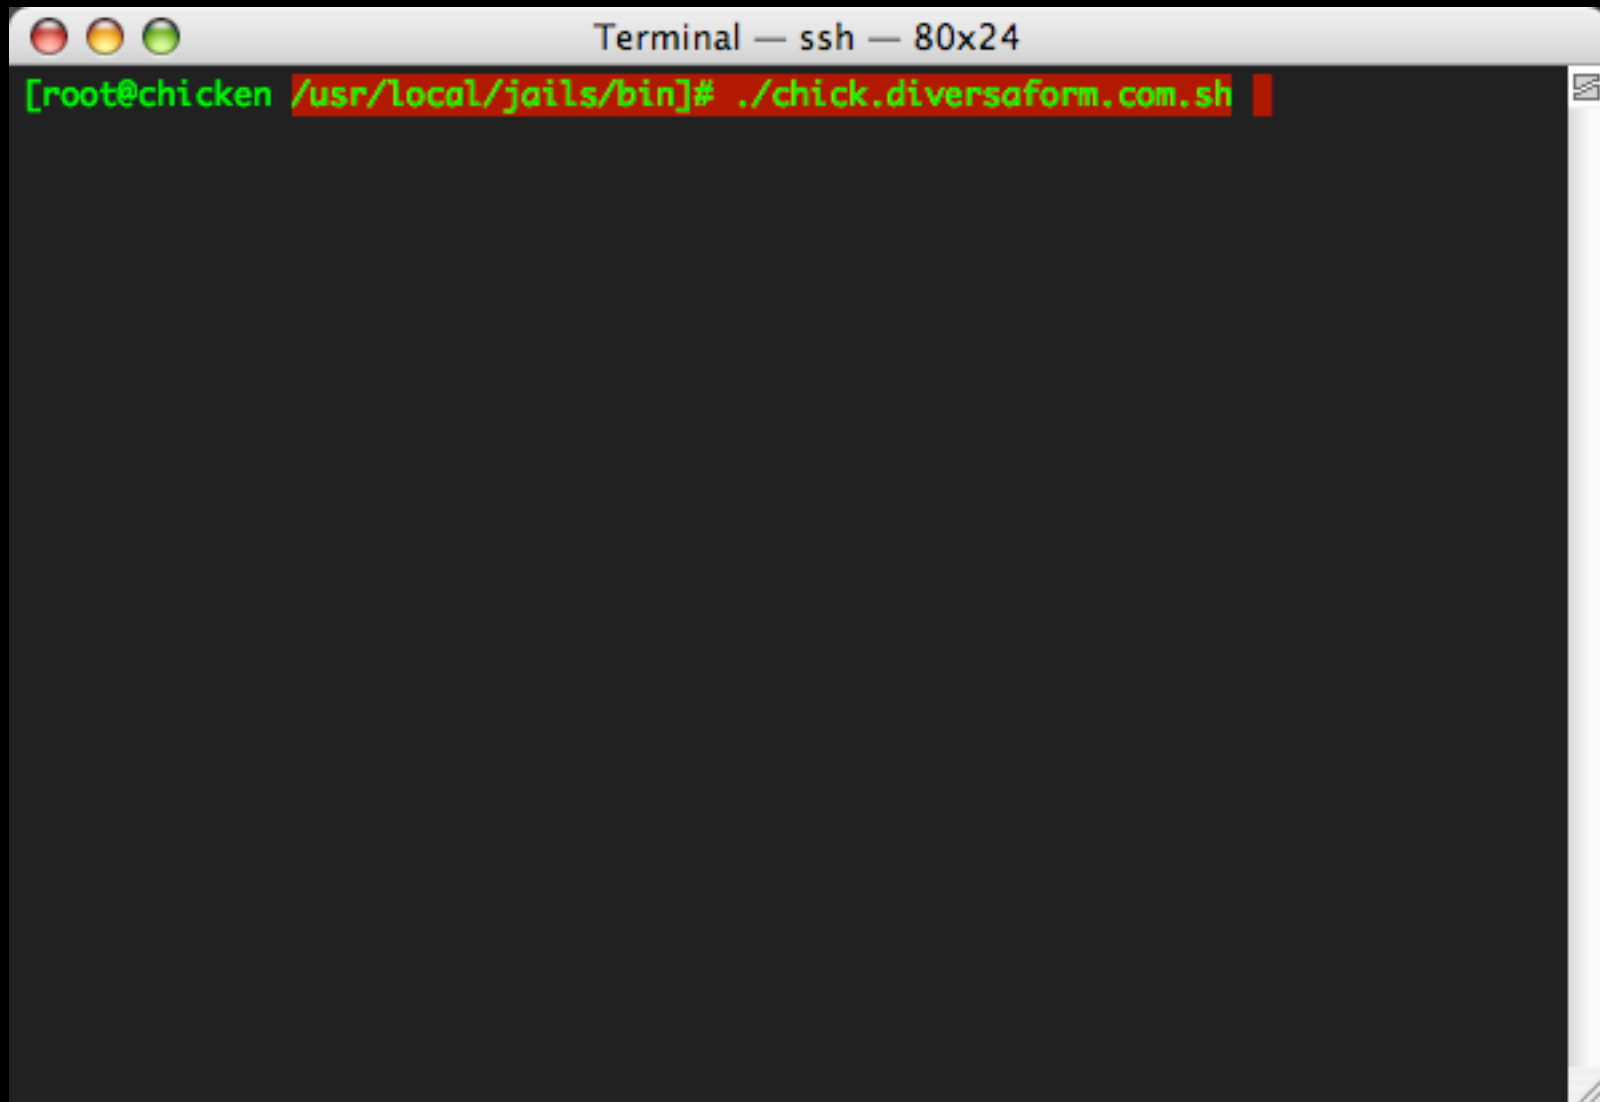
watch out for stacking mount points!

```
Terminal — ssh — 80x24
555    0 /usr/libexec/getty Pc ttyv5
556    0 /usr/libexec/getty Pc ttyv6
557    0 /usr/libexec/getty Pc ttyv7
199    0 dhclient: bge0 [priv] (dhclient)
230    0 dhclient: bge1 [priv] (dhclient)
561    0 -bash (bash)
576    0 dtach -c /tmp/JailLecture1 -Ez bash
570    0 -bash (bash)
583    0 dtach -c /tmp/JailLecture2 -Ez bash
578    0 bash
757    0 bash
1471   0 ps ax -o pid,jid,args
585    0 bash
587    0 bash
[root@chicken /usr/local/jails]# jls
      JID  IP Address      Hostname                Path
      5   10.0.1.200      chick.diversaform.net  /usr/local/jails/chick.div
ersaform.net
[root@chicken /usr/local/jails]# killall -j 5
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/p
roc
[root@chicken /usr/local/jails]# umount /usr/local/jails/chick.diversaform.net/d
ev
[root@chicken /usr/local/jails]#
```



# stop and start jail

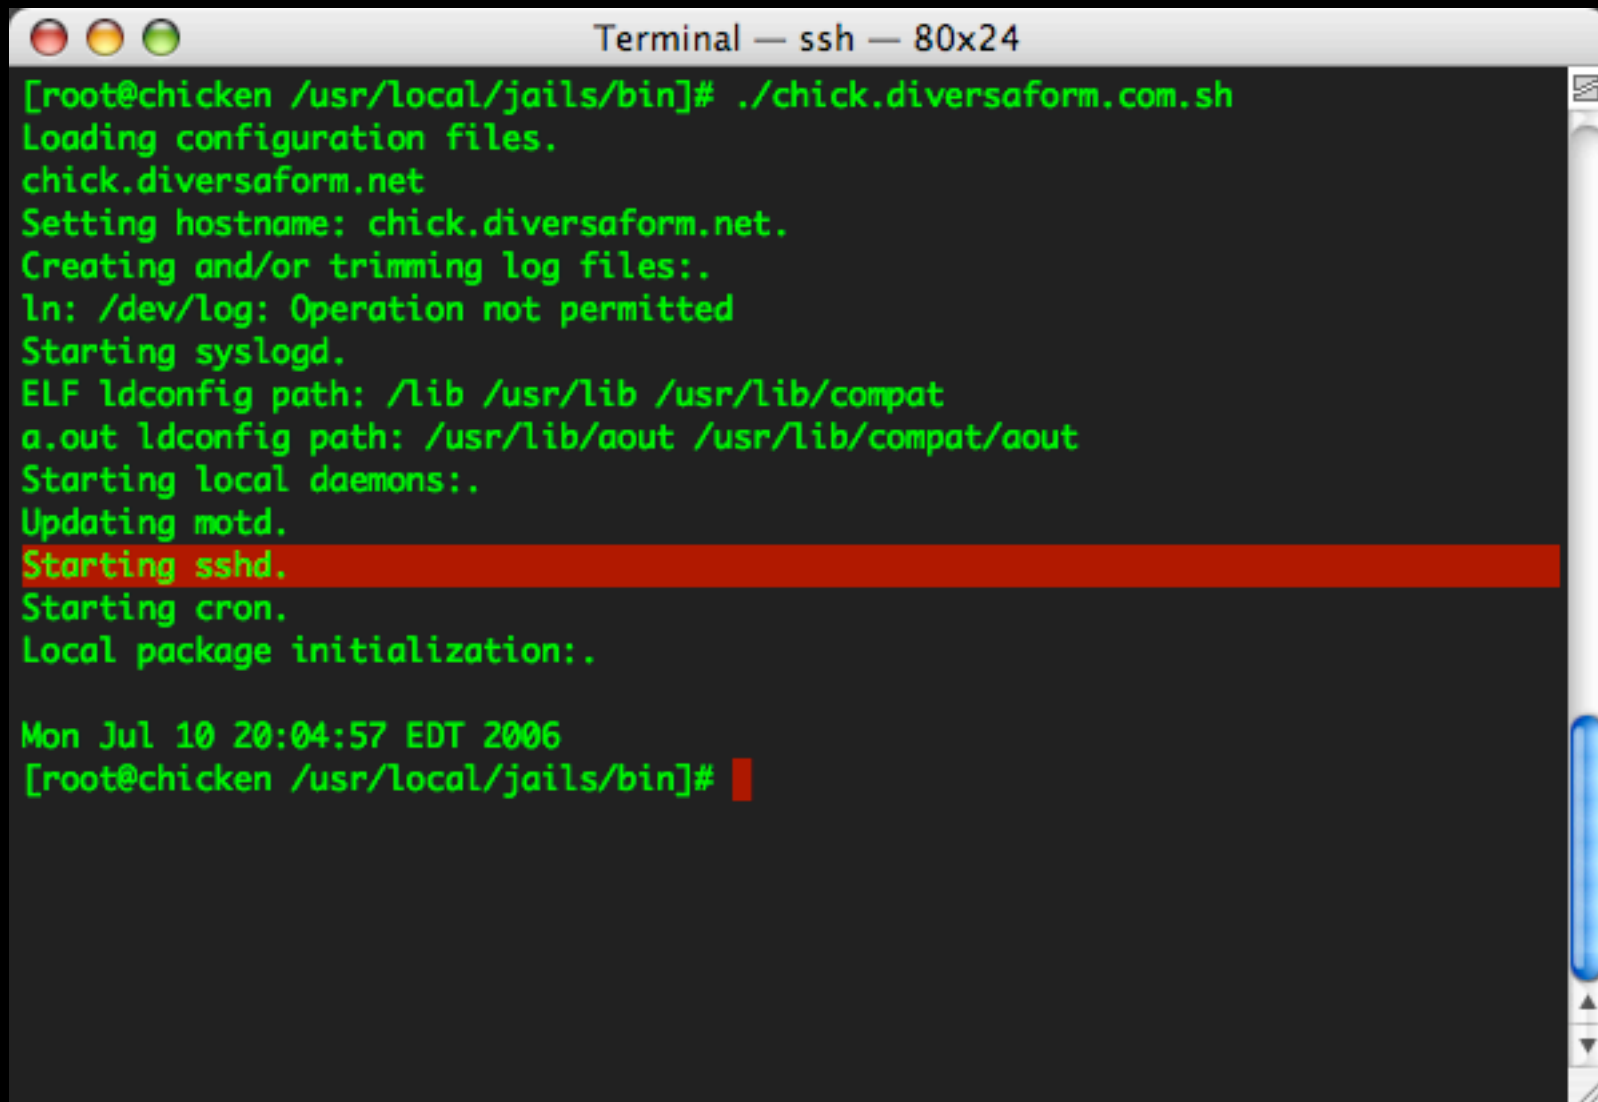
restarting with the script this time,

A terminal window titled "Terminal — ssh — 80x24" with three window control buttons (red, yellow, green) in the top-left corner. The terminal prompt is "[root@chicken /usr/local/jails/bin]#". The command being executed is "./chick.diversaform.com.sh", which is highlighted in red. The rest of the terminal area is empty and black.

```
Terminal — ssh — 80x24
[root@chicken /usr/local/jails/bin]# ./chick.diversaform.com.sh
```

# stop and start jail

restarting with the script this time,

A terminal window titled "Terminal — ssh — 80x24" with a dark background and green text. The window shows the execution of a script named "chick.diversaform.com.sh" from the directory "/usr/local/jails/bin". The script performs several system configuration tasks: loading configuration files, setting the hostname to "chick.diversaform.net", creating or trimming log files (with a warning for /dev/log), starting syslogd, setting ldconfig paths for ELF and a.out, starting local daemons, updating motd, starting sshd (highlighted with a red bar), starting cron, and local package initialization. The terminal shows the date "Mon Jul 10 20:04:57 EDT 2006" and the prompt "[root@chicken /usr/local/jails/bin]#".

```
Terminal — ssh — 80x24
[root@chicken /usr/local/jails/bin]# ./chick.diversaform.com.sh
Loading configuration files.
chick.diversaform.net
Setting hostname: chick.diversaform.net.
Creating and/or trimming log files:.
In: /dev/log: Operation not permitted
Starting syslogd.
ELF ldconfig path: /lib /usr/lib /usr/lib/compat
a.out ldconfig path: /usr/lib/aout /usr/lib/compat/aout
Starting local daemons:.
Updating motd.
Starting sshd.
Starting cron.
Local package initialization:.

Mon Jul 10 20:04:57 EDT 2006
[root@chicken /usr/local/jails/bin]#
```

# stop and start jail

now the jid has incremented once, to 6

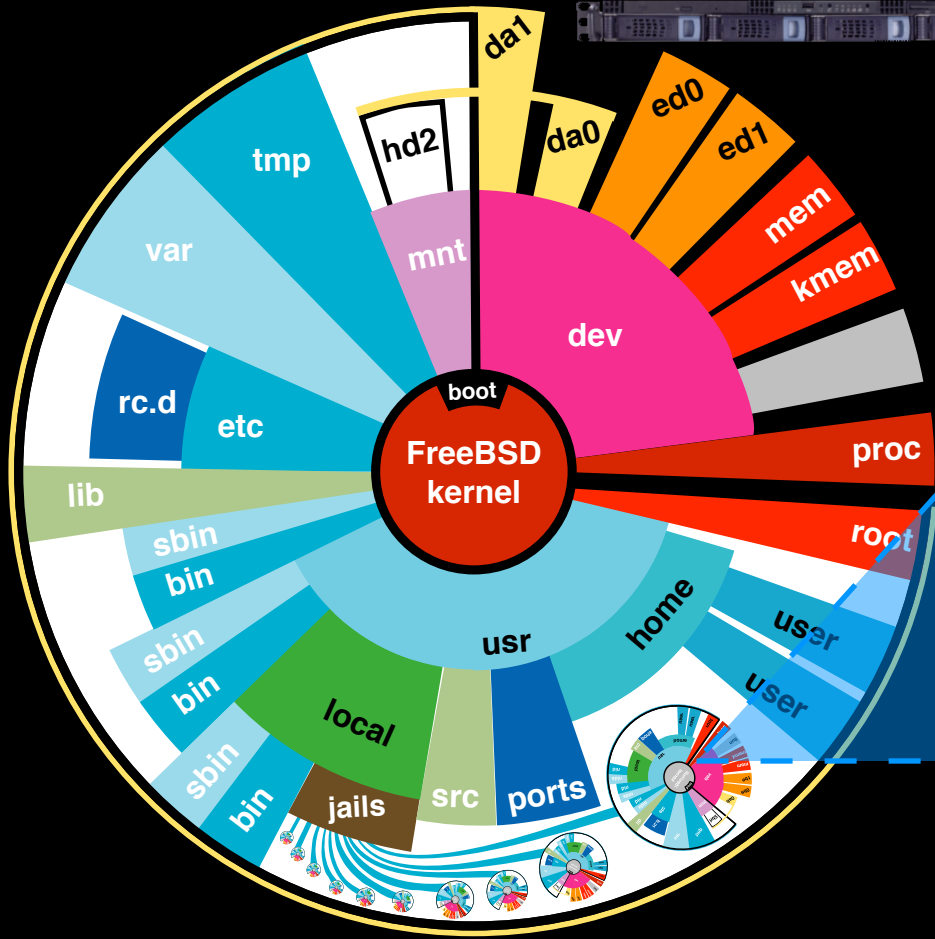
```
Terminal — ssh — 80x24
378 0 /usr/sbin/syslogd -s
451 0 /usr/sbin/usbd
493 0 /usr/sbin/sshd
499 0 sendmail: accepting connections (sendmail)
503 0 sendmail: Queue runner@00:30:00 for /var/spool/clientmqueue (sendm
509 0 /usr/sbin/cron -s
558 0 sshd: ike [priv] (sshd)
560 0 sshd: ike@tty0 (sshd)
567 0 sshd: ike [priv] (sshd)
569 0 sshd: ike@tty1 (sshd)
577 0 dtach -c /tmp/JailLecture1 -Ez bash
584 0 dtach -c /tmp/JailLecture2 -Ez bash
1571 6 /usr/sbin/syslogd -s
1623 6 /usr/sbin/sshd
1629 6 sendmail: accepting connections (sendmail)
1633 6 sendmail: Queue runner@00:30:00 for /var/spool/clientmqueue (sendm
1639 6 /usr/sbin/cron -s
550 0 /usr/libexec/getty Pc ttyv0
551 0 /usr/libexec/getty Pc ttyv1
552 0 /usr/libexec/getty Pc ttyv2
553 0 /usr/libexec/getty Pc ttyv3
554 0 /usr/libexec/getty Pc ttyv4
555 0 /usr/libexec/getty Pc ttyv5
556 0 /usr/libexec/getty Pc ttyv6
```

# running processes


jexec to check processes (bad idea, in practice)

```
Terminal — ssh — 80x24
557    0 /usr/libexec/getty Pc ttyv7
199    0 dhclient: bge0 [priv] (dhclient)
230    0 dhclient: bge1 [priv] (dhclient)
561    0 -bash (bash)
576    0 dtach -c /tmp/JailLecture1 -Ez bash
570    0 -bash (bash)
583    0 dtach -c /tmp/JailLecture2 -Ez bash
578    0 bash
757    0 bash
1659   0 ps ax -o pid,jid,args
585    0 bash
587    0 bash
[root@chicken /usr/local/jails/bin]#
[root@chicken /usr/local/jails/bin]# jexec 6 ps auxc
USER      PID %CPU %MEM    VSZ   RSS  TT  STAT  STARTED      TIME  COMMAND
root    1571  0.0  0.0  1300   948  ??   IsJ    8:04PM    0:00.00  syslogd
root    1623  0.0  0.1  3356  2876  ??   IsJ    8:04PM    0:00.00  sshd
root    1629  0.0  0.1  3420  2840  ??   SsJ    8:04PM    0:00.01  sendmail
smmsp   1633  0.0  0.1  3300  2720  ??   IsJ    8:04PM    0:00.00  sendmail
root    1639  0.0  0.0  1312  1044  ??   IsJ    8:04PM    0:00.00  cron
root    1718  0.0  0.0  1344   864  p2    R+J    8:21PM    0:00.00  ps
[root@chicken /usr/local/jails/bin]# jexec 6 whoami
root
[root@chicken /usr/local/jails/bin]#
```

# Practical Comparison



host



jail

# Process Tree:

JailingServer

\\_init

- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...

\\_jail (Jail 1)

- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...

\\_jail (Jail 2)

- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...


\\_jail (Jail 3)

- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...


\\_jail (Jail 4)

- \\_daemon/process etc...
- \\_daemon/process etc...
- \\_daemon/process etc...

host



jail







```
Terminal — ssh — 80x24
chickenhawk:/home/ike ike$ sudo cat /proc/35721/status
sshd 35721 1 35721 35721 - sldr 1138578391,270482 4,956102 27,766581 select 0 0
0,0 chick.diversaform.net
chickenhawk:/home/ike ike$ sudo cat /proc/93030/status
sshd 93030 93025 93025 93025 noflags 1152553319,908960 0,61426 0,40950 select
1001 1001 1001,1001,1001,0 -
chickenhawk:/home/ike ike$
```



host




jail

```
Terminal — ssh — 80x24
chickenhawk:/home/ike ike$ sudo cat /proc/35721/status
sshd 35721 1 35721 35721 - sldr 1138578391,270482 4,956102 27,766581 select 0 0
0,0 chick.diversaform.net
chickenhawk:/home/ike ike$ sudo cat /proc/93030/status
sshd 93030 93025 93025 93025 - noflags 1152553319,908960 0,61426 0,40950 select
1001 1001 1001,1001,1001,0
chickenhawk:/home/ike ike$
```



host



jail




```

Terminal — ssh — 80x24
chickenhawk:/home/ike ike$ sudo cat /proc/35721/status
sshd 35721 1 35721 35721 - sldr 1138578391,270482 4,956102 27,766581 select 0 0
0,0 chick.diversaform.net
chickenhawk:/home/ike ike$ sudo cat /proc/93030/status
sshd 93030 93025 93025 93025 - noflags 1152553319,908960 0,73969 0,52749 select
1001 1001 1001,1001,1001,0 -
chickenhawk:/home/ike ike$ ps auxc | grep VSZ; ps auxc | grep 35721; ps auxc | g
rep 93030
USER          PID %CPU %MEM  VSZ   RSS TT  STAT  STARTED    TIME COMMAND
root          35721 0.0  0.1  3352  1152 ??  IsJ   29Jan06    0:32.72 sshd
ike           93030 0.0  0.1  6092  2492 ??   S     1:41PM    0:00.13 sshd
chickenhawk:/home/ike ike$

```



host



jail



you may also want to run `/etc/rc.shutdown` from within the jail. To kill processes from outside the jail, use the `jexec(8)` utility in conjunction with the one of the `kill(1)` commands above.

The `/proc/pid/status` file contains, as its last field, the hostname of the jail in which the process runs, or ```-''` to indicate that the process is not running within a jail. The `ps(1)` command also shows a ``J'` flag for processes in a jail. However, the hostname for a jail may be, by default, modified from within the jail, so the `/proc` status entry is unreliable by default. To disable the setting of the hostname from within a jail, set the `security.jail.set_hostname_allowed` sysctl variable in the host environment to `0`, which will affect all jails. You can have this sysctl set on each boot using `sysctl.conf(5)`. Just add the following line to `/etc/sysctl.conf`:

```
security.jail.set_hostname_allowed=0
```


You can also list/kill processes based on their jail ID. To show processes and their jail ID, use the following command:

```
ps ax -o pid,jid,args
```



To show and then kill processes in jail number 3 use the following com-

:|





# jail(8) best practices





and opportunities...



# break out of jail?



**Poul-Henning Kamp (PHK)** wrote the jail feature for R&D Associates <http://www.rndassociates.com/> who contributed it to FreeBSD around 1998.

To my knowledge, nobody has broken out of a jail directly, ever. It is however assumed that nobody has tried that hard yet, as it is still considered 'esoteric'.

If someone breaks jail, PHK wrote that he would love to know about it.

# best practices

- ssh into jails to manage their processes!!!!
- You always can see the jailed filesystem/userland from host server, be careful.
- Design your jailing system carefully, be creative with core UNIX utilities.
- Use your highest secure practices for host server...

# great utilities

- 4.x, jps, jkill, jtop
- 5.x, 6.x, onward builtin ps, kill
  - !plus jls(8), jexec(8) jattach(2), sysctl features for jailing
- Design your jailing system carefully, be creative (note about nullfs, devfs)
- additionally, handy: pstree, xtail, disk images via mdconfig

# common weak points

- lost jail?
  - [hostname lockdown]
- resource attacks
  - disks full
    - [partitions, disk images]
  - fork bombs, memory hogs
    - [securelevels, login.conf]
  - process control
- direct driver access
  - [flags to mount devfs, procfs]


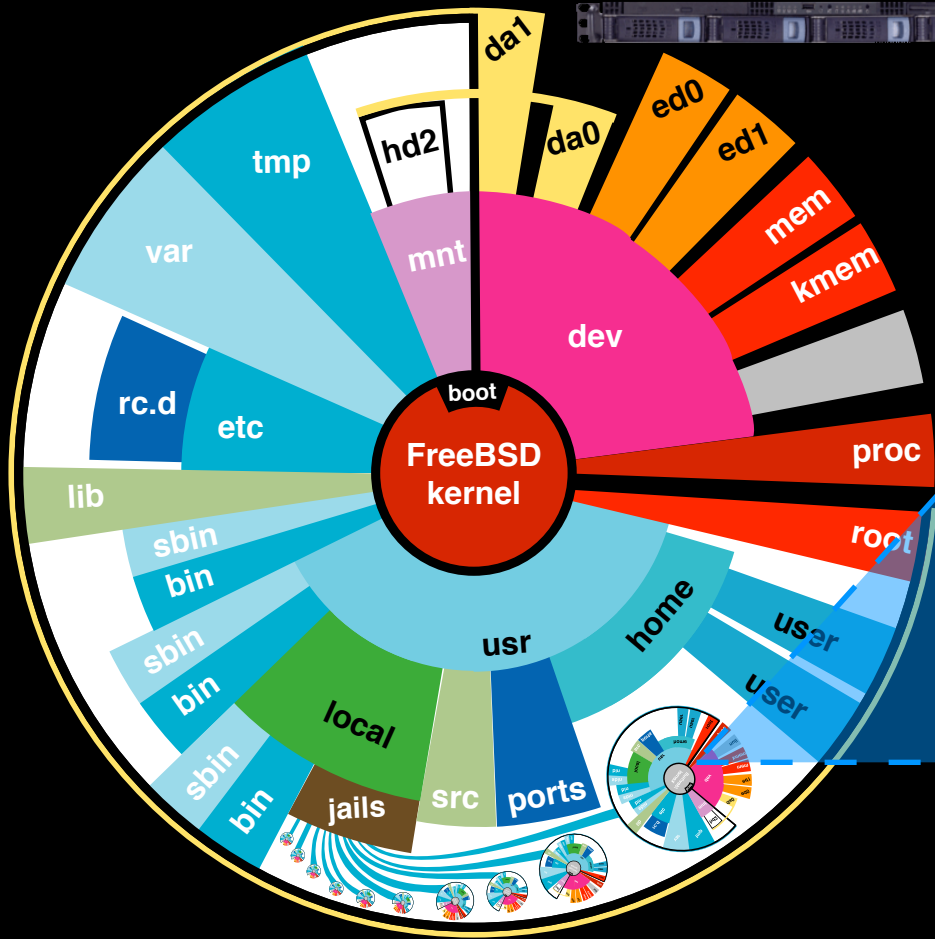
# common weak points

- lost jail?
  - [hostname lockdown]
- resource attacks
  - disks full
    - [partitions, disk images]
  - fork bombs, memory hogs
    - [securelevels, login.conf]
  - process control
- direct driver access
  - [flags to mount devfs, procfs]





# Comments on Isolation



# memory/process attacks



<http://www.samag.com/documents/s=1151/sam0105d/0105d.htm>

[http://www.freebsd.org/doc/en\\_US.ISO8859-1/books/arch-handbook/jail.html](http://www.freebsd.org/doc/en_US.ISO8859-1/books/arch-handbook/jail.html)

OpenRoot Project, fork-bombs, FreeBSD  
SecureLevels/maxproc, reality, and process control

# memory/process attacks

(check the Defcon 14 CD)

```
# hog.c, a small utility to hog system memory  
# written by Brian Redman (BER) sometime around 1986  
# Basic Instructions, Compile this code to a binary:
```

```
cc hog.c -o hog
```

```
# then run something like:
```

```
hog 10
```

```
# and the hog will do just that- sit and hog 10mb of ram.  
# To run a hog stampede, (a fork bomb):
```

```
while (1)  
hog 99m&  
end
```

# memory/process attacks

(check the Defcon 14 CD)

```
# STEP 1)
# jailed /etc/login.conf file, example of restricted values:
:maxproc=30:\
:memoryuse=25M:\

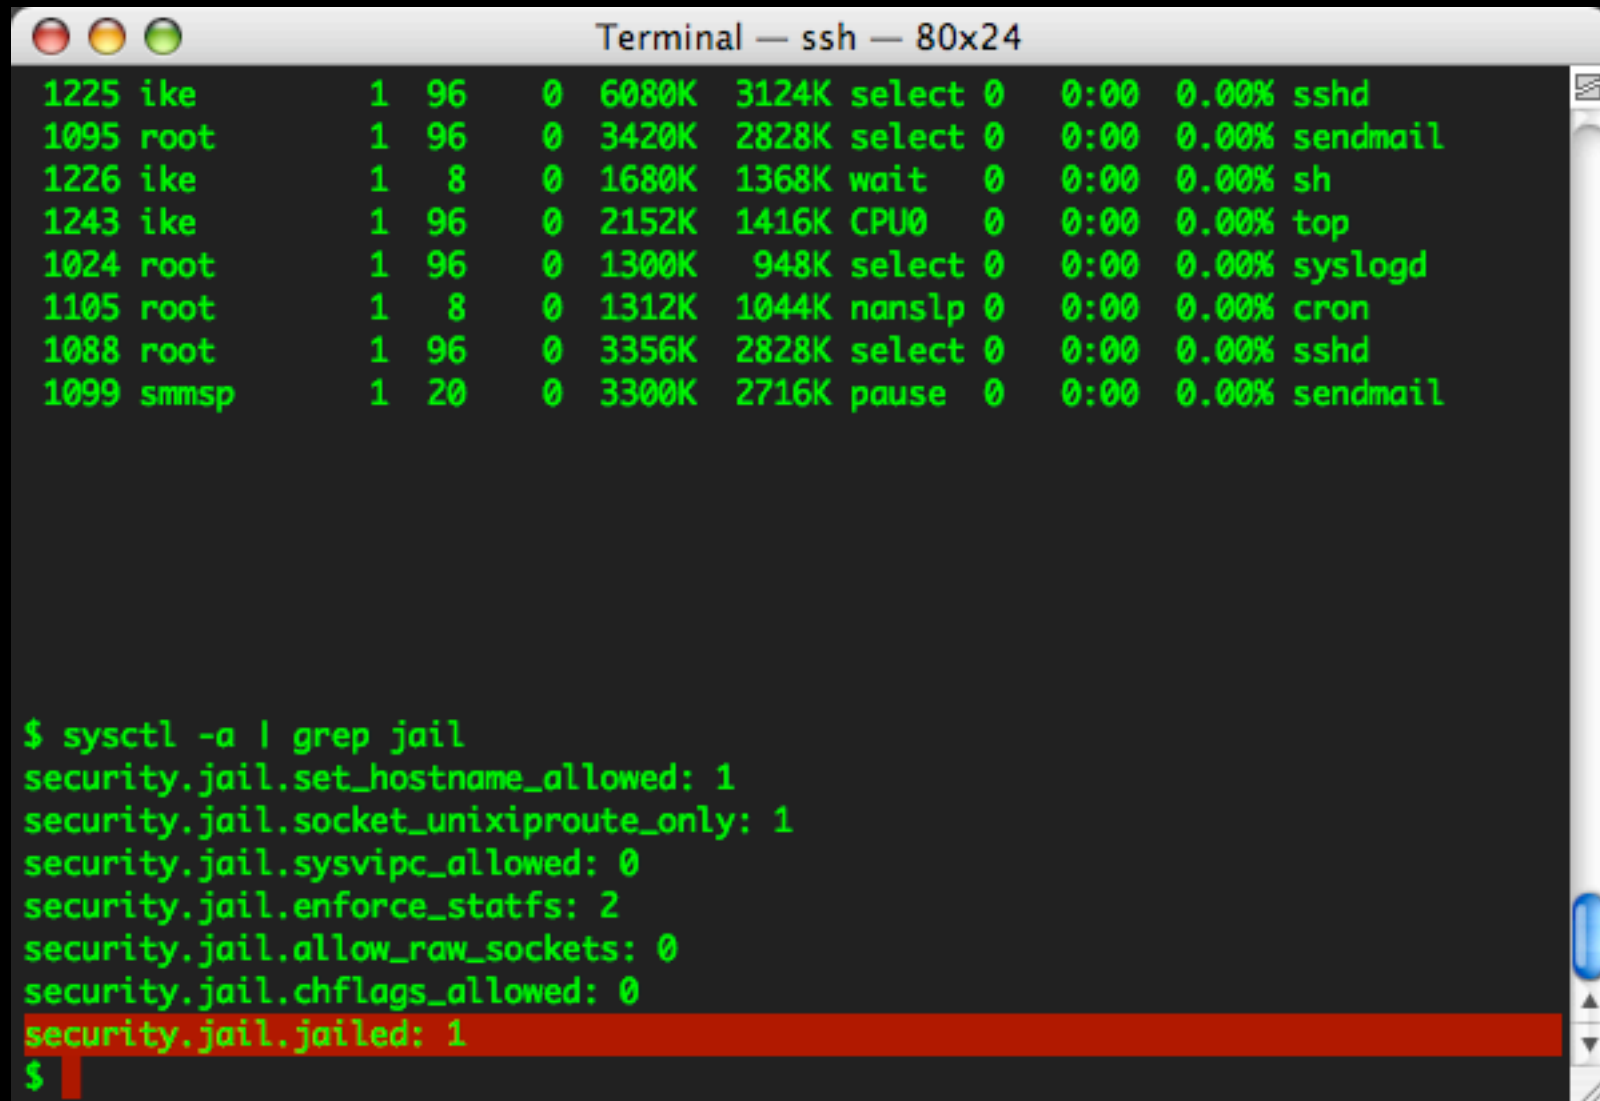
# STEP 2)
# Set immutable flags on jailed /etc/login.conf, example:
chflags schg $D/etc/login.conf

# STEP 3)
# Set a higher securelevel on a per-jail basis
# (5.x onward, 4.x jailing only securlevels for entire host)
# add the following line to the jailed /etc/sysctl.conf:

kern.securelevel=2
# securelevel 1 is minimum, read the man page for securelevel
```

# honeypot?

compile and give the jail a kernel, fix sysctl:



```
Terminal — ssh — 80x24
1225 ike      1  96    0 6080K 3124K select 0   0:00  0.00% sshd
1095 root     1  96    0 3420K 2828K select 0   0:00  0.00% sendmail
1226 ike      1   8    0 1680K 1368K wait   0   0:00  0.00% sh
1243 ike      1  96    0 2152K 1416K CPU0   0   0:00  0.00% top
1024 root     1  96    0 1300K  948K select 0   0:00  0.00% syslogd
1105 root     1   8    0 1312K 1044K nanslp 0   0:00  0.00% cron
1088 root     1  96    0 3356K 2828K select 0   0:00  0.00% sshd
1099 smmsp    1  20    0 3300K 2716K pause  0   0:00  0.00% sendmail

$ sysctl -a | grep jail
security.jail.set_hostname_allowed: 1
security.jail.socket_unixiproute_only: 1
security.jail.sysvipc_allowed: 0
security.jail.enforce_statfs: 2
security.jail.allow_raw_sockets: 0
security.jail.chflags_allowed: 0
security.jail.jailed: 1
$
```

# disk resource control

- Put at least your jailed systems on a separate partition, or perhaps each jail (rigid in practice)
- File-Backed Disk Images (mdconfig, in handbook)- insanely flexible, but take extra memory (usually negligible)



# file-backed disks (.dmg)

- WOW, they're convenient.
- watch out for device numbering (or things get lost), here's where Jailing strategies from 4.x come in handy... unless someone has a better way of mangling device nodes
- speed is getting excellent for file-backed memory disks, but will always introduce some overhead in file I/O

# file-backed disks (.dmg)

FreeBSD handbook has tons more information!

```
# writing 1gb blank file, (analagous to creating an  
unformatted harddrive)
```

```
dd if=/dev/zero of=1gb.img bs=1k count=1024k
```

```
# attaching the file (analagous to attaching a harddrive)...
```

```
mdconfig -a -t vnode -f 1gb.img -u 1101
```

```
# formatting the disk...
```

```
disklabel -r -w md1101 auto
```

```
# detaching the disk (analagous to ejecting a harddrive)...
```

```
mdconfig -d -u 1101
```

# file-backed disks (.dmg)

mount disks when starting jails,  
<snip - jail start script>

```
mdconfig -a -t vnode -f /path/to/jaildisk_file.dmg -u 200  
mount /dev/md200c /path/to/jail_userland_mount_dir
```

```
# regarding '-u 200' above, it can be handy to use some  
# variant of a jail's respective IP address for it's disk  
# image devide node id, so it's easy to track down on host  
# system with many jailed servers.
```

```
# later in script,  
jail /path/to/jail_userland_mount_dir \  
hostname.fqdn.com \  
10.0.1.200 \  
/bin/sh /etc/rc
```

</snip>

# automation

- Tarball packaging is your friend.  
clean, simple, reliable.  
    be aware of dev/proc mounts  
    be aware of symlinks
- use FreeBSD Ports Mechanism!  
    (**not** for the ports collection, that's insanely presumptuous, [borderline irresponsible])
- CVS/SVN anyone?

# upgrading jailed systems

- Simply use buildworld, (FROM HOST SYSTEM),
- toss buildworld DESTDIR flag, with a jail's userland path
- follow the handbook: [http://www.freebsd.org/doc/en\\_US.ISO8859-1/books/handbook/makeworld.html](http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/makeworld.html)

# /etc/sysctl.conf (host)

(check the Defcon 14 CD)

```
# $FreeBSD: src/etc/sysctl.conf,v 1.8 2003/03/13 18:43:50 mux Exp $
#
# This file is read when going to multi-user and its contents piped thru
# ``sysctl'' to adjust kernel values.  ``man 5 sysctl.conf'' for details.
#

# Uncomment this to prevent users from seeing information about processes that
# are being run under another UID.
#security.bsd.see_other_uids=0

# ikenote jailing additives
security.jail.set_hostname_allowed=0      # default = 1 # jailed resetting hostname.
security.jail.enforce_statfs=2           # default = 2 # mount point info.
security.jail.allow_raw_sockets=0        # default = 0 # for ping, etc...
security.jail.socket_unixiproute_only=1  # default = 1 # access to routing sockets.
security.jail.sysvipc_allowed=0          # default = 0 # SysV shared mem? Ha!
security.jail.chflags_allowed=0         # default = 0 # root less than root...
```



# sysctl (stock values)

(check the Defcon 14 CD)

```
$ sysctl -a | grep jail
security.jail.set_hostname_allowed: 0
security.jail.socket_unixiproute_only: 1
security.jail.sysvipc_allowed: 0
security.jail.enforce_statfs: 2
security.jail.allow_raw_sockets: 0
security.jail.chflags_allowed: 0
security.jail.jailed: 0
```

# firewalls (quick comment)

- context:
  - *why jail in the first place again?*
- threats affect an entire host server
- firewall at a higher level (mental shift to treat the host like a network gateway!)
  - global system firewalling, throttling
  - different boxes? different rules?

# Start Script w/ Disk Image

(check the Defcon 14 CD)

```
#!/bin/sh

# simple, complete script to start a jail.

# define the absolute path to the jail,
J=/usr/local/jails/jailed.userland.directory

# define the ip address for the jail,
I=10.0.1.192

# define a hostname,
H=fqdn.com

ifconfig en0 inet alias $I/32

mount -t procfs proc $J/proc
mount_devfs devfs $J/dev
## add additonal flags to mount_devfs, to hide unnecessary devices!!!
## check the man page for mount_devfs

jail $J $H $I /bin/sh /etc/rc
```

# jail crontab misc...

(check the Defcon 14 CD)

```
# comment out the following, just to keep syslog quiet for irrelevant items.  
  
# Save some entropy so that /dev/random can re-seed on boot.  
# */11 * * * * operator /usr/libexec/save-entropy  
  
# Adjust the time zone if the CMOS clock keeps local time, as opposed to  
# UTC time. See adjkerntz(8) for details.  
# 1,31 0-5 * * * root adjkerntz -a
```

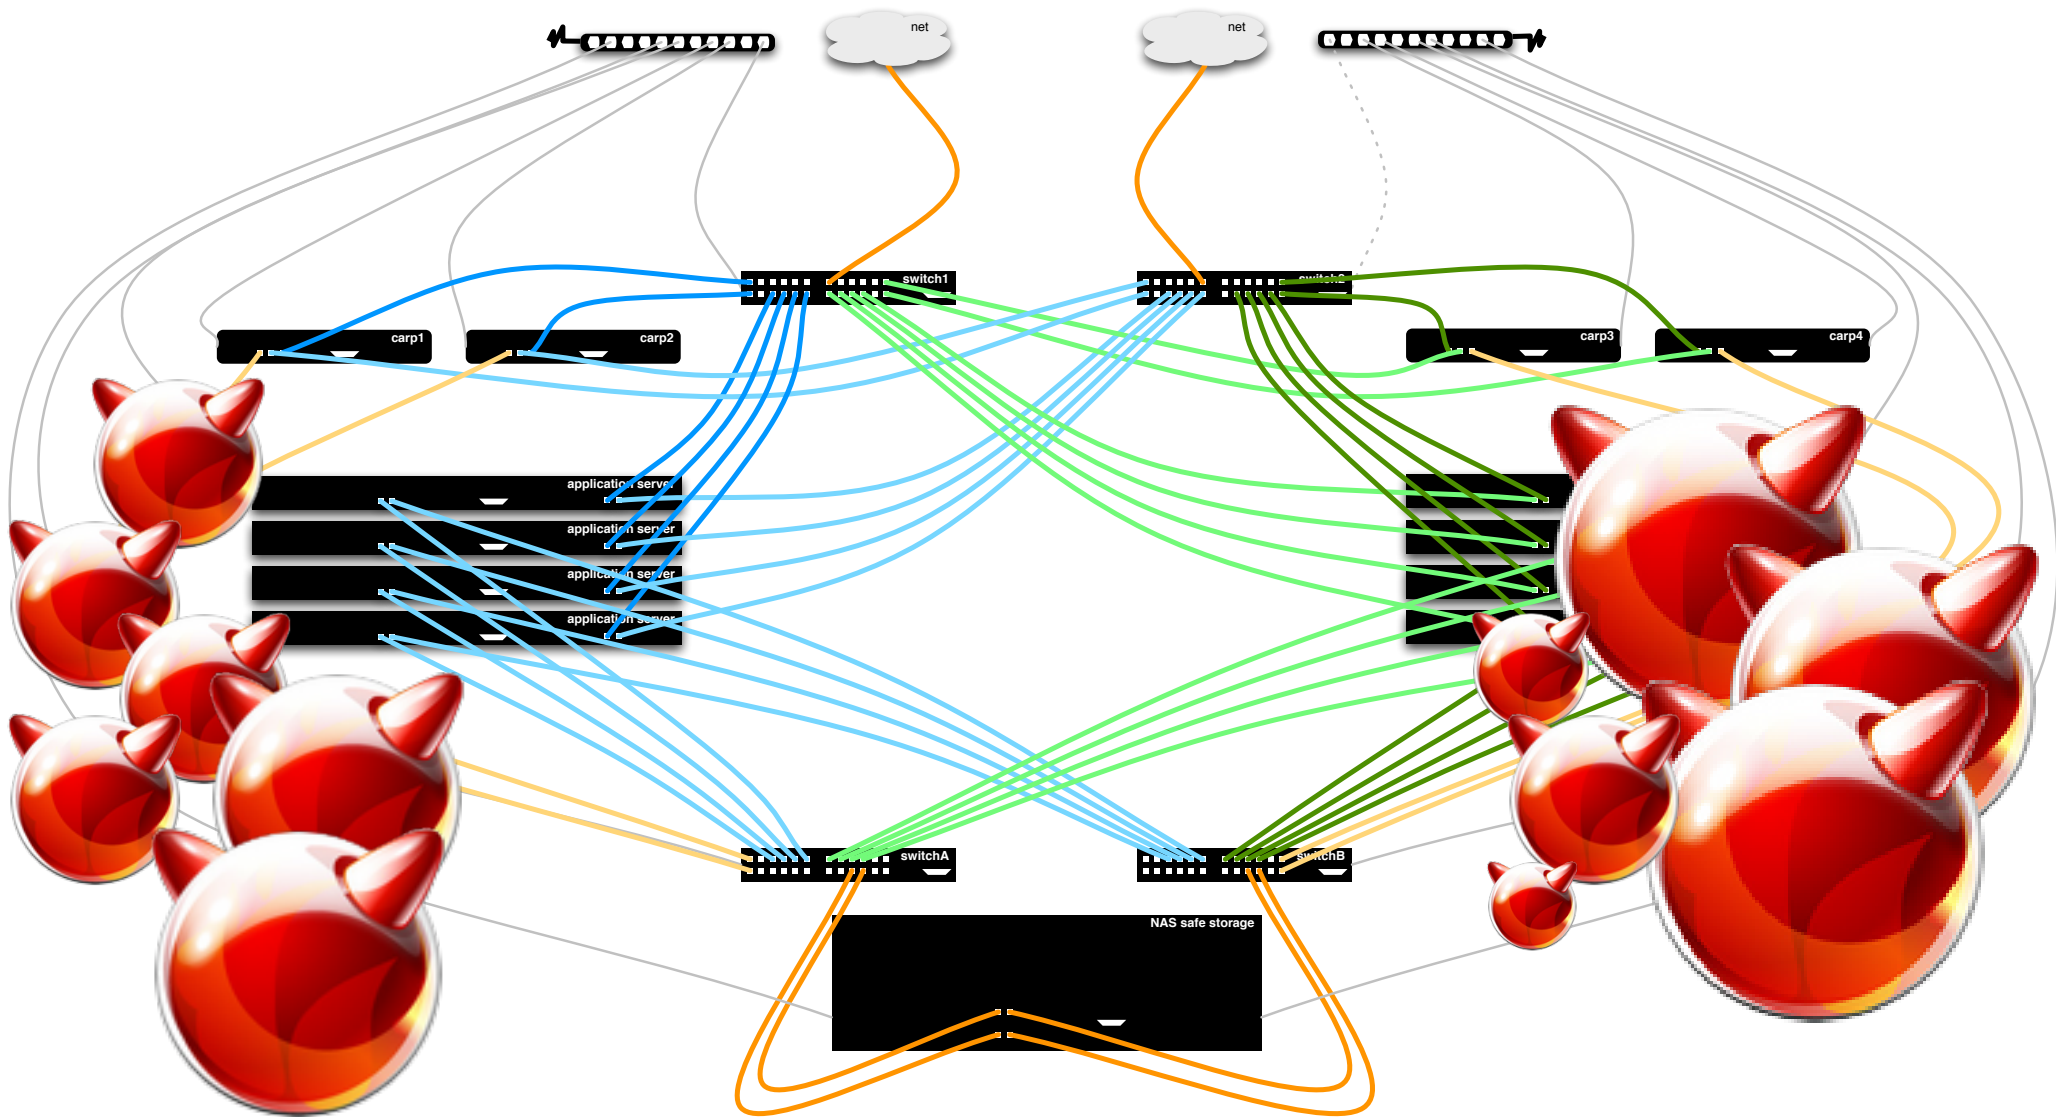
# future directions...

## important fun:

- CARP, from PF/OpenBSD
- GEOM
- NFS Improvements
- more NAS/SAN support (GEOM, ggated)
- FreeBSD 4.x, 5.x, 6.x, (7x!)

# sick possibilities...

GEOM Gate, CARP, fun with failover jails...



# misc

- Compile md(4) into the kernel for File-Backed Disks, for better performance
- GOTCHA: rm a jail directory? `chflags -R noschg jaildir`





# jailing.net

 Search

news members

you are not logged in log in join

## Go to Jail(8) Free Site

**jailing.net is a website dedicated to jail(8), a unix system to imprison a process and its descendants, used for virtualizing machine services by isolating an operating system.**

jail(8) utility, and jail(2) system call, first appeared in FreeBSD 4.0 and currently only exists in

# Stillborn.

For knowledge, this is the only jail(8) website which exists to date.

The mission of jailing is as follows:

- 1) To centralize and keep track resources on the topic of jailing which exist on the internet
- 2) To provide a platform for new documentation and information about jailing
- 3) To provide a platform for individuals and organizations involved with jailing, providing membership for content authorship

As we are currently under construction a stock home journal, for more information:

[one website](#)

Last modified 2005-03-22 03:16

### navigation

Home

### log in

Name

Password

Log in

Forgot your

password?

Register?

### news

[BSDCan in Canada!](#)  
2005-03-22

[More...](#)

### upcoming events

[BSDCan FreeBSD Developer Summit](#)  
Ottawa, Canada,  
2005-05-11

[BSDCan in Canada](#)  
Ottawa, Canada,  
2005-05-12

March 2005						
Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

# Suggestions?

# **NMCBUG**

NEW YORK CITY \*BSD USER GROUP

**Jailing Party - jail(8)  
Friday-Monday, Sept. 3-6 2004  
Next time we get more resilient hardware...  
Thanks everyone for all the fun!**





# Special Thanks:



**wintermute** (of iMeme), taught me to jail(8).  
He's here somewhere- buy him a drink.



**reality** schooled me more BSD than he knows...



**Poul-Henning Kamp** wrote the jail feature for R&D Associates <http://www.rndassociates.com/> who contributed it to FreeBSD around 1998.



**Robert Watson** wrote the extended documentation, found a few bugs, added a few new features, and cleaned up the userland jail environment.





[isaac@diversaform.com](mailto:isaac@diversaform.com)

Ike is proud to be a part of the *New York City \*BSD Users Group*,  
and the *Lower East Side Mac Unix Users Group*



**NYCBUG**  
NEW YORK CITY \*BSD USER GROUP