Life after Xserve

Life after Xserve

Part I: Recreating netboot

Why netboot?

- Image standard image without local media
- Update in one place: available everywhere
- Recovery without building recovery media

Boot Service Discovery Protocol

- DHCP
- TFTP
- NFS / HTTP / AFP

Phase 1: Get kernel

- Standard DHCP request to get IP address
  - *Any DHCP server can do this.*
- Custom request for a *basic kernel*
  - *Understand and handle BSDP requests. (Darwin: bootpd, ISC DHCP server)*
- Load basic kernel from TFTP

Phase II: Get filesystem

- After loading basic kernel, custom DHCP request to get a root filesystem
  - disk image can be mounted via HTTP, NFS, (AFP)
  - Client does not speak DNS yet!

How about a picture?

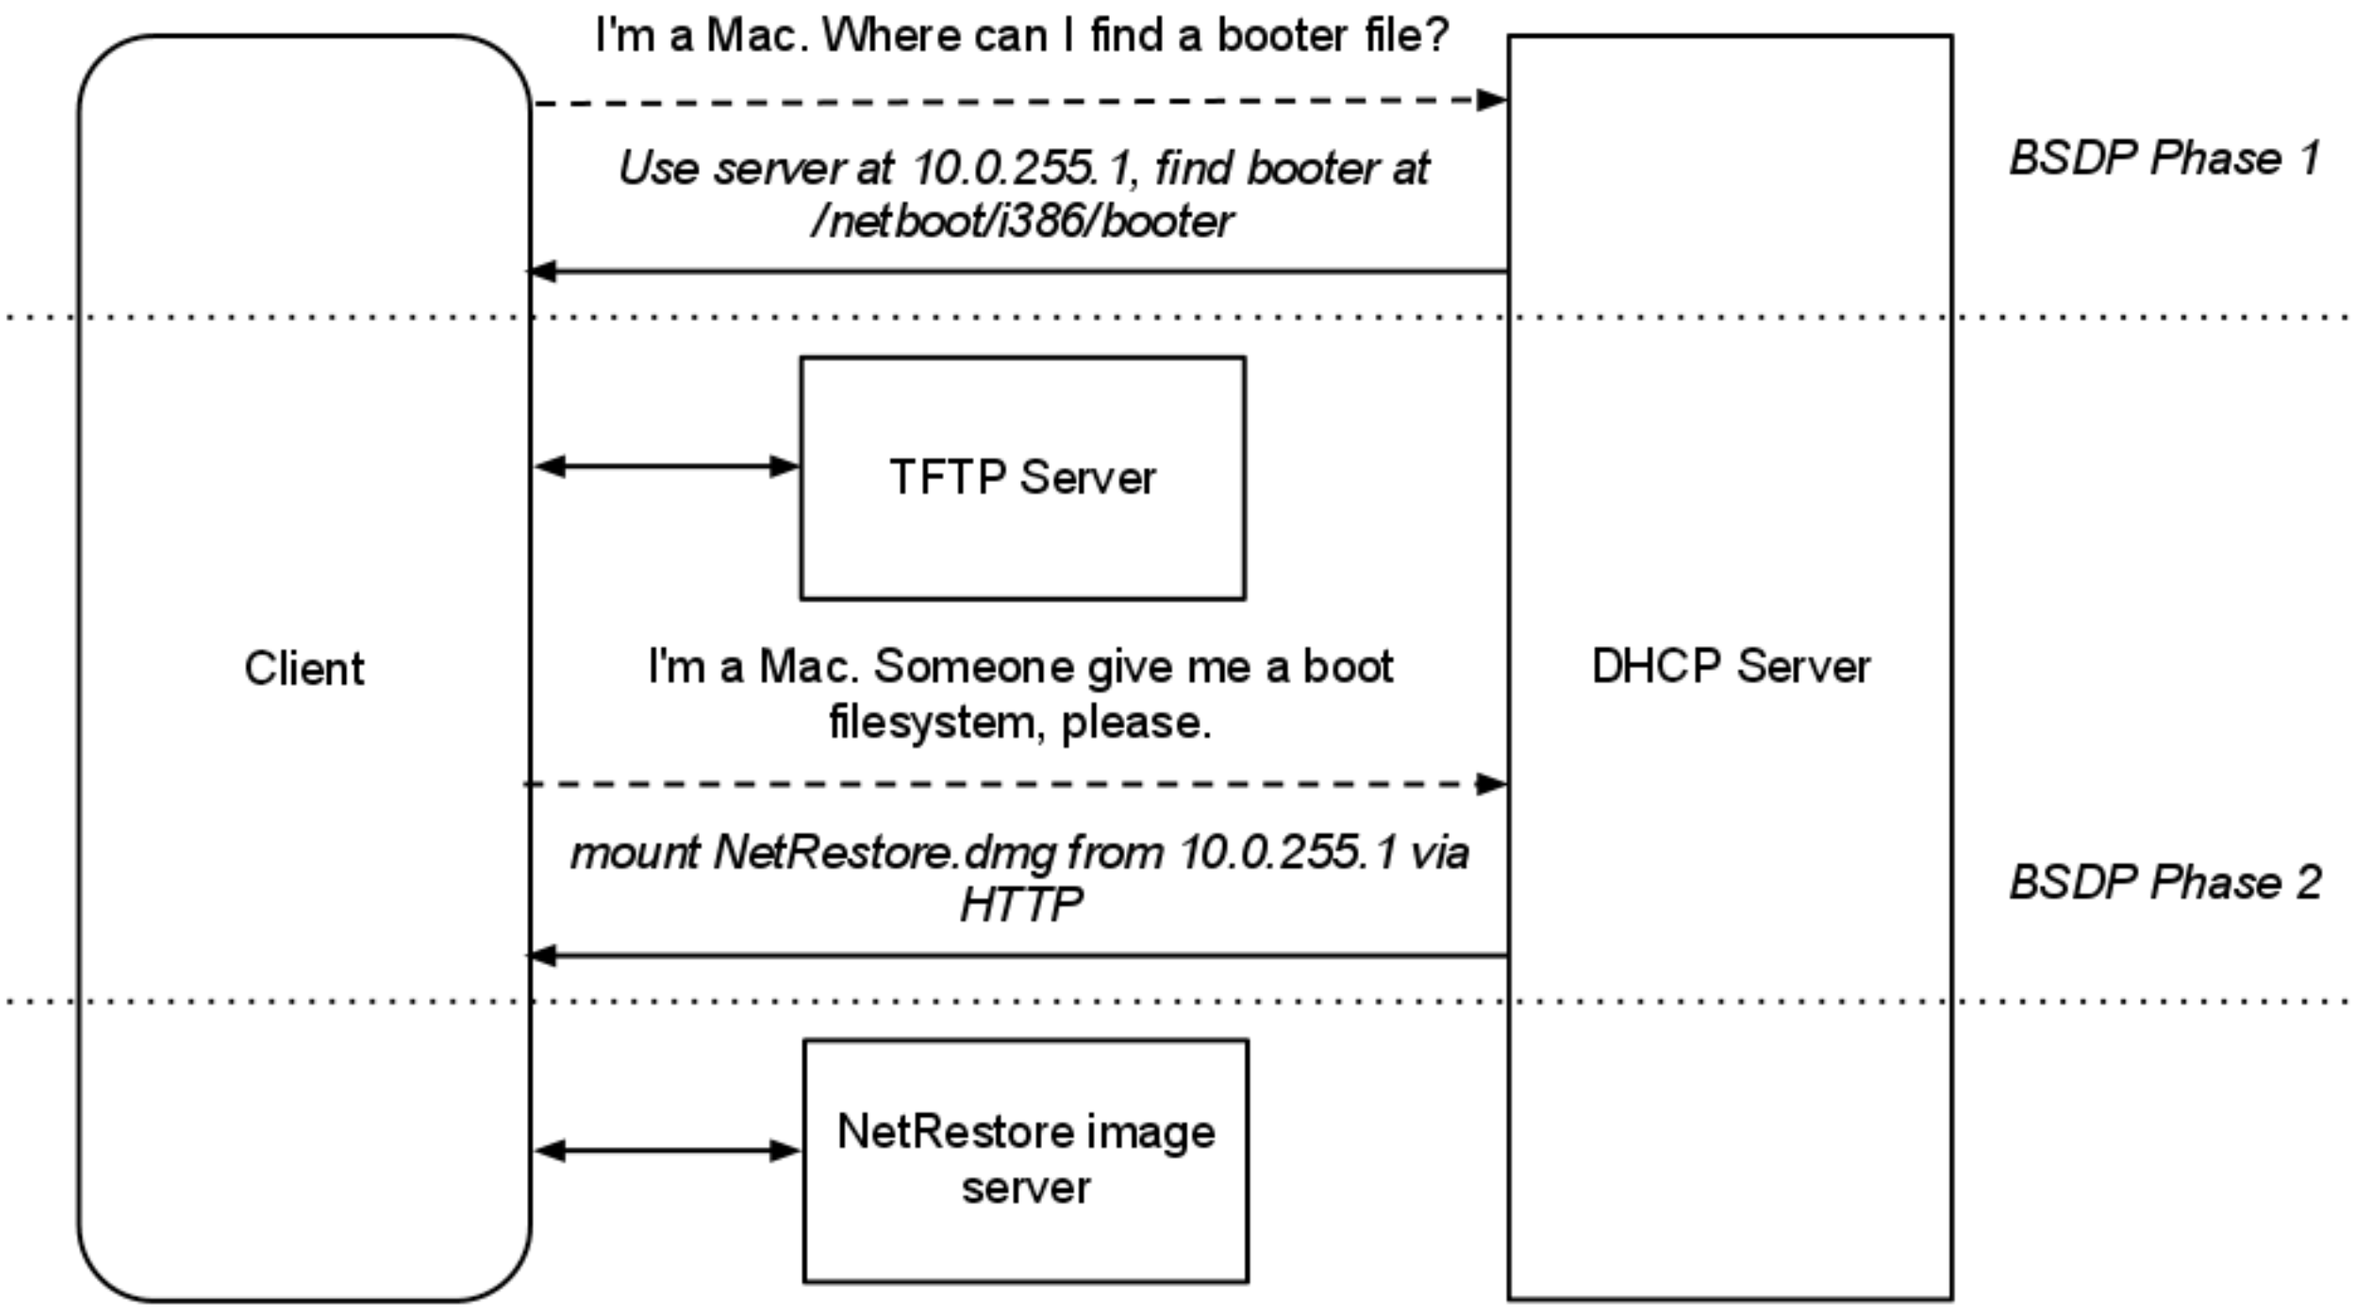
I'm a Mac. Where can I find a booter file?
Use server at 10.0.255.1, find booter at /netboot/i386/booter
BSDP Phase 1
TFTP Server
Client
I'm a Mac. Someone give me a boot filesystem, please.
DHCP Server
mount NetRestore.dmg from 10.0.255.1 via HTTP
BSDP Phase 2
NetRestore image server

Co-opting ISC DHCP Server

magic numbers and trivial servers

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {

}
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {
    match if substring (option vendor-class-identifier, 0, 14) = "AAPLBSDPC/i386";
}
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {
    match if substring (option vendor-class-identifier, 0, 14) = "AAPLBSDPC/i386";
    vendor-option-space AAPL; # this is set by Apple
    option vendor-class-identifier "AAPLBSDPC"; # this is set by Apple
    option dhcp-parameter-request-list 1,3,11,17,43,60; # dhcp options to set



}
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {
    match if substring (option vendor-class-identifier, 0, 14) = "AAPLBSDPC/i386";
    vendor-option-space AAPL; # this is set by bootpd
    option vendor-class-identifier "AAPLBSDPC/i386"; # this is set by bootpd
    option dhcp-parameter-request-list 1,3,11,17,43,60; # dhcp options to set

    if (option dhcp-message-type = 1) {
        option vendor-encapsulated-options 08:04:81:00:00:67; # magic number
    }


}
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {
    match if substring (option vendor-class-identifier, 0, 14) = "AAPLBSDPC/i386";
    vendor-option-space AAPL; # this is set by bootpd
    option vendor-class-identifier "AAPLBSDPC/i386"; # this is set by bootpd
    option dhcp-parameter-request-list 1,3,11,17,43,60; # dhcp options to set

    if (option dhcp-message-type = 1) {
        option vendor-encapsulated-options 08:04:81:00:00:67; # magic number
    }

    filename "netboot/i386/booter"; # downloaded from “next-server” via TFTP
    option root-path "http://192.168.254.1/path/to/NetInstall-Restore.dmg";
}
```

```
# private netboot network # sample private network

subnet 10.0.255.0 netmask 255.255.255.0 {
    range 10.0.255.2 10.0.255.20;
    option routers 10.0.255.1; # default gateway
}

# standard DHCP options

allow bootp; # for backwards compatibility with simple clients
default-lease-time 600; # leases will be reserved for 10min
max-lease-time 7200; # reserve for no longer than 2h
option domain-name-servers 10.0.255.1; # DNS server
next-server 10.0.255.1; # IP address of TFTP server
# netboot specific options

class "AppleNBI-i386" {
    match if substring (option vendor-class-identifier, 0, 14) = "AAPLBSDPC/i386";
    vendor-option-space AAPL; # this is set by bootpd
    option vendor-class-identifier "AAPLBSDPC/i386"; # this is set by bootpd
    option dhcp-parameter-request-list 1,3,11,17,43,60; # dhcp options to set

    if (option dhcp-message-type = 1) {
        option vendor-encapsulated-options 08:04:81:00:00:67; # magic number
    }

    filename "netboot/i386/booter"; # downloaded from “next-server” via TFTP
    option root-path "http://192.168.254.1/path/to/NetInstall-Restore.dmg";
}
```

TFTP server piftalls

- 5 minute timeout (in Apple's firmware)
- anonymous read access (no security)
- hardware specific booter files / kernel extensions
- *Lion:*
  SupportedPlatforms.plist
  com.apple.Boot.plist

DHCP server pitfalls

- finicky configuration syntax
- error logging leaves room for improvement
- other hardware

tcpdump

- network debugging swiss army knife
- tcpdump -vvv -n -s 0 -i eth0 -w dhcp.cap \
  udp port 67 or udp port 68 or broadcast
- tcpdump -vvv -n -s 0 -i eth0 -w tftp.cap \
  udp port 69
- tcpdump -vvv -n -s 0 -i eth0 -w http.cap \
  http port 80

NetRestore image

- build with System Image Utility or similar
- split up .nbi hierarchy between servers:
  - i386 directory served by TFTP server
  - NetRestore.dmg served by web server *(Install image served by web server)*

Life after XServe

Life after XServe

Part II: Hardware specific images

You're in a maze of twisty little images, all looking alike

OS X: Branches

- New hardware often runs own branch
- Standard OS releases missing new drivers
- Branches sometimes missing mainline fixes
- To avoid nasty surprises, build hardware specific images

Where to branch?

- OS X Server branches during BSDP
  - Possible with ISC DHCP
    - how often do you want to change your DHCP server configuration?
    - Maybe it's sufficient to branch during image selection?

Custom image catalogs

- add to startup using /etc/rc.cdrom.local
- check hardware model *(sysctl)*
- download hardware specific InstallPreferences.plist
- Success!

Further pitfalls

- NetRestore mostly read-only
- Batteries not included / some tools missing
- booter files still need support for new hardware (graphics & network drivers)

Demo time

- Ubuntu Linux
- tftp-hpa, ISC DHCP server, apache
- MacBook Pro 13” (late 2011)

tftp configuration

```
# mkdir -p /var/lib/tftroot/netboot

# cp -rav /path/to/your.nbi/i386 \
  /var/lib/tftproot/netboot/i386

# chmod -R a+r \
  /var/lib/tftproot/netboot/i386

# /etc/init.d/tftp-hpa start
```