



October 12-14, 2004

Managing NFS with Autofs

Brent Callaghan

Sun Microsystems, Inc.

`brent.callaghan@sun.com`



October 12-14, 2004

What's Coming...

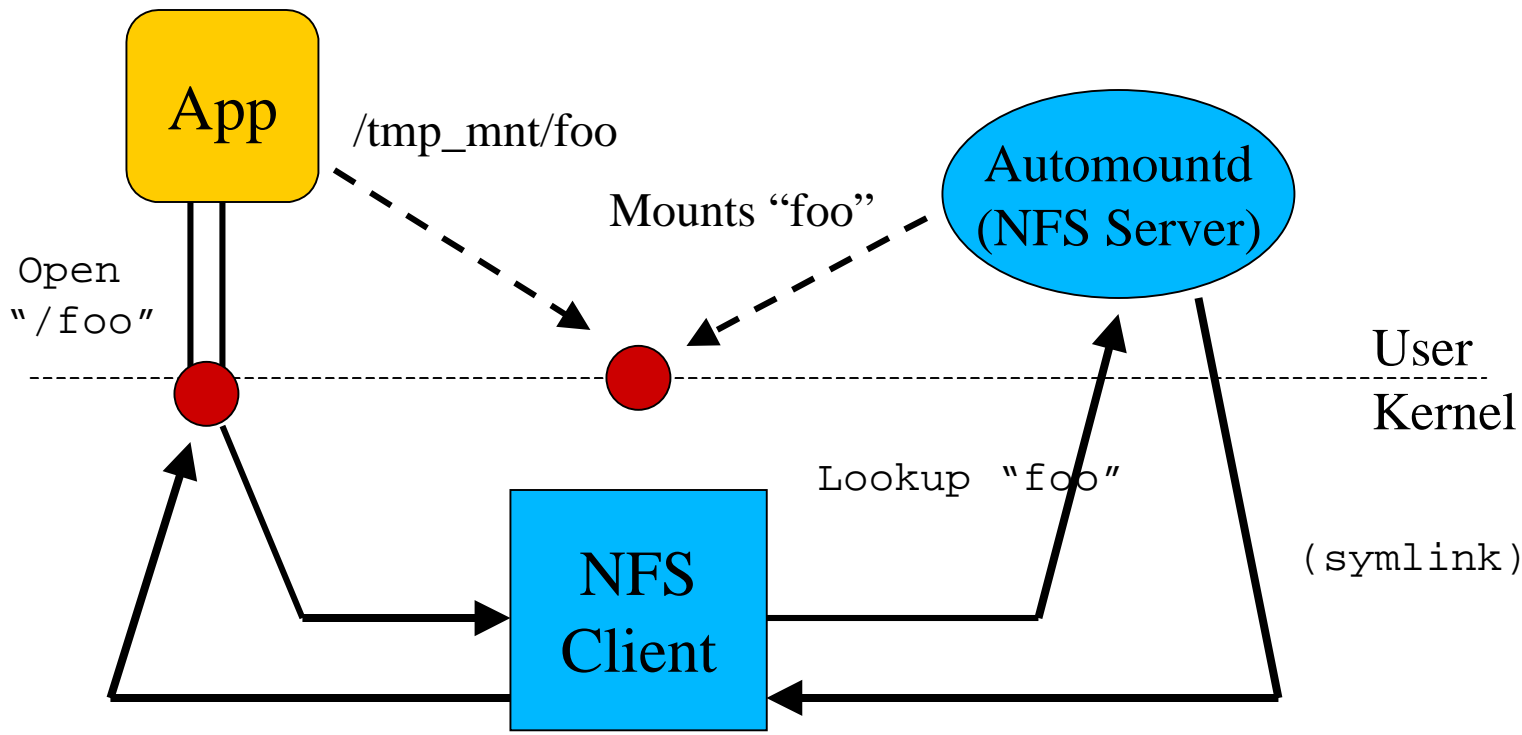
- Some history
- Direct Maps
- Indirect Maps
- Master Map
- Executable Maps
- Map Sources
- Offset Mounts
- Mount Options
- Multiple Mounts
- Replica Servers
- Map Variables
- Key Substitution
- Auto-unmounting
- Autofs Mounts
- Visible Directories
- and finally some
- Recommendations



October 12-14, 2004

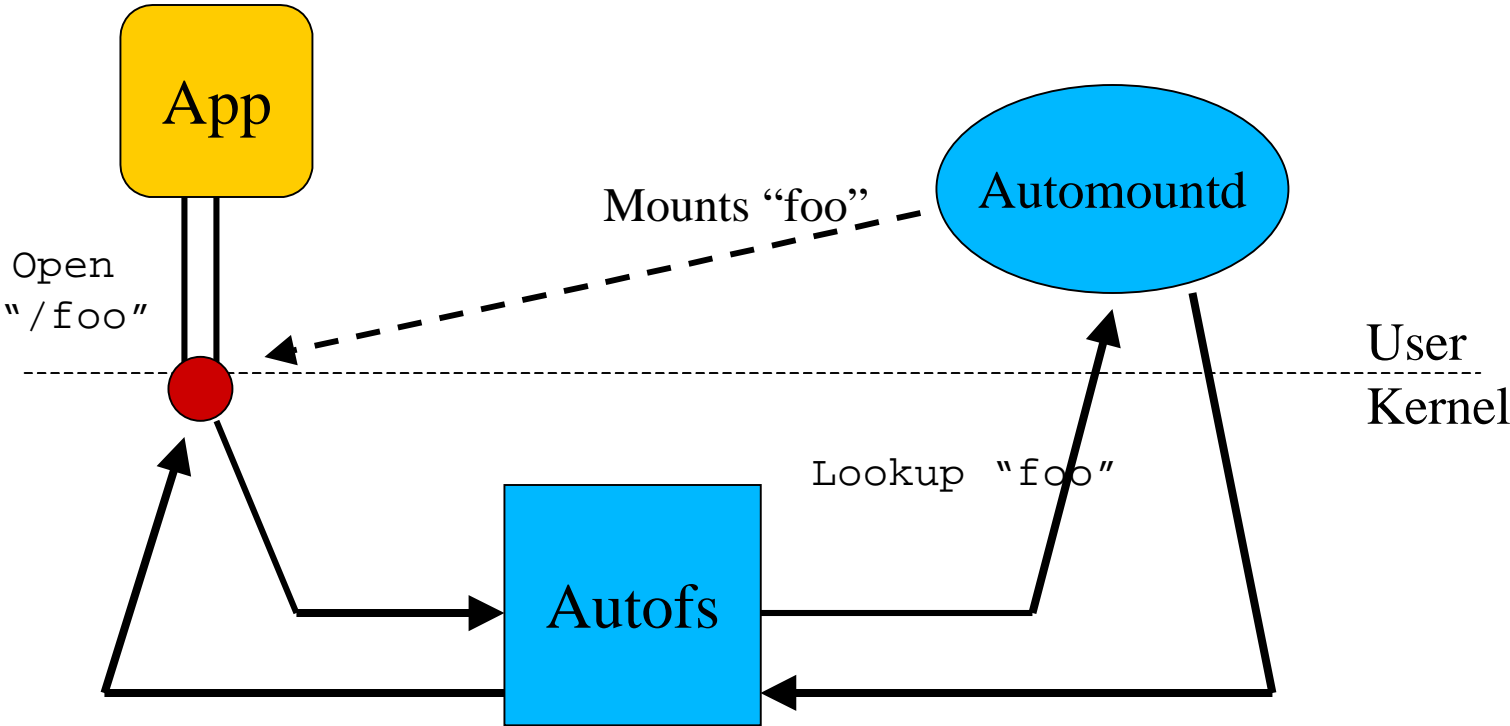
In the beginning - The Automounter

No kernel changes - loopback NFS



Enter Autofs

- A kernel filesystem
- Eliminates symbolic link





October 12-14, 2004

Automounter Maps

Map a name to a location

Name

Location

```
blorf      rtzblf:/export/fnarg
fratzl     nord:/export/frib
gloib      zygorf:/export/snitz
...
```



October 12-14, 2004

Direct Maps

Directly map autofs mountpoints to locations
(not very scalable)

Autofs Mountpoints

Locations

```
/blorf          rtzblf:/export/fnarg  
/usr/fratzl    nord:/export/frib  
/opt/glolb     zygorf:/export/snitz  
...
```



October 12-14, 2004

Indirect Maps

Map autofs directory entries to locations
(hugely scalable)

Autofs mountpoint

/home

```
alan      rtzblf:/export/alan
jane      nord:/export/jane
peter     zygorf:/export/peter
...
```



October 12-14, 2004

Master Map

Establishes initial autofs mountpoints

Autofs mountpoint

Map name

Default mount options

```
/blorf          auto_blorf  
/home         auto_home -noquota  
/-            auto_direct  
+auto_local
```





October 12-14, 2004

Executable Maps


Generate locations on the fly

godzilla



Executable program
or script

```
#!/bin/sh
SERVER=$1
dfshares -h $1 |
awk '{print $1 "\t'$SERVER':" $1 "\\\"}' |
sort
```



godzilla:/export/shares



October 12-14, 2004

Executable Maps (cont)

Useful for:

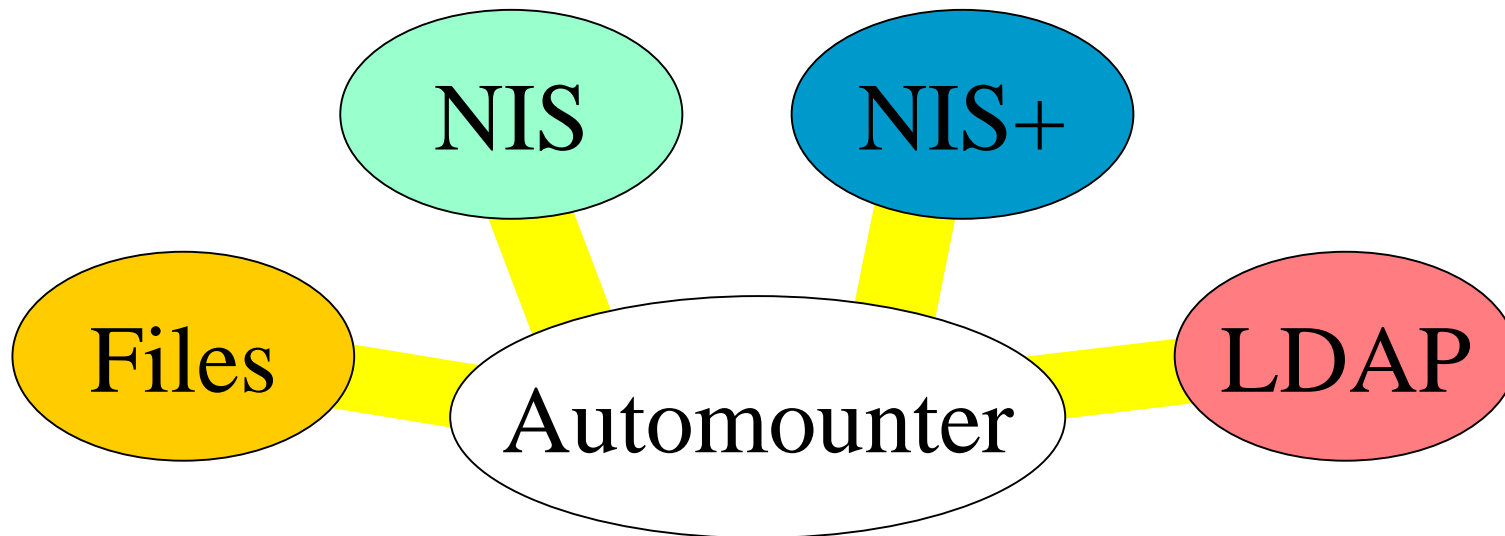
- Map sources that are not directly accessible by automounter.
- Implementing policy based on client, time of day, location, etc.
- Highly dynamic location data.



October 12-14, 2004

Map Sources

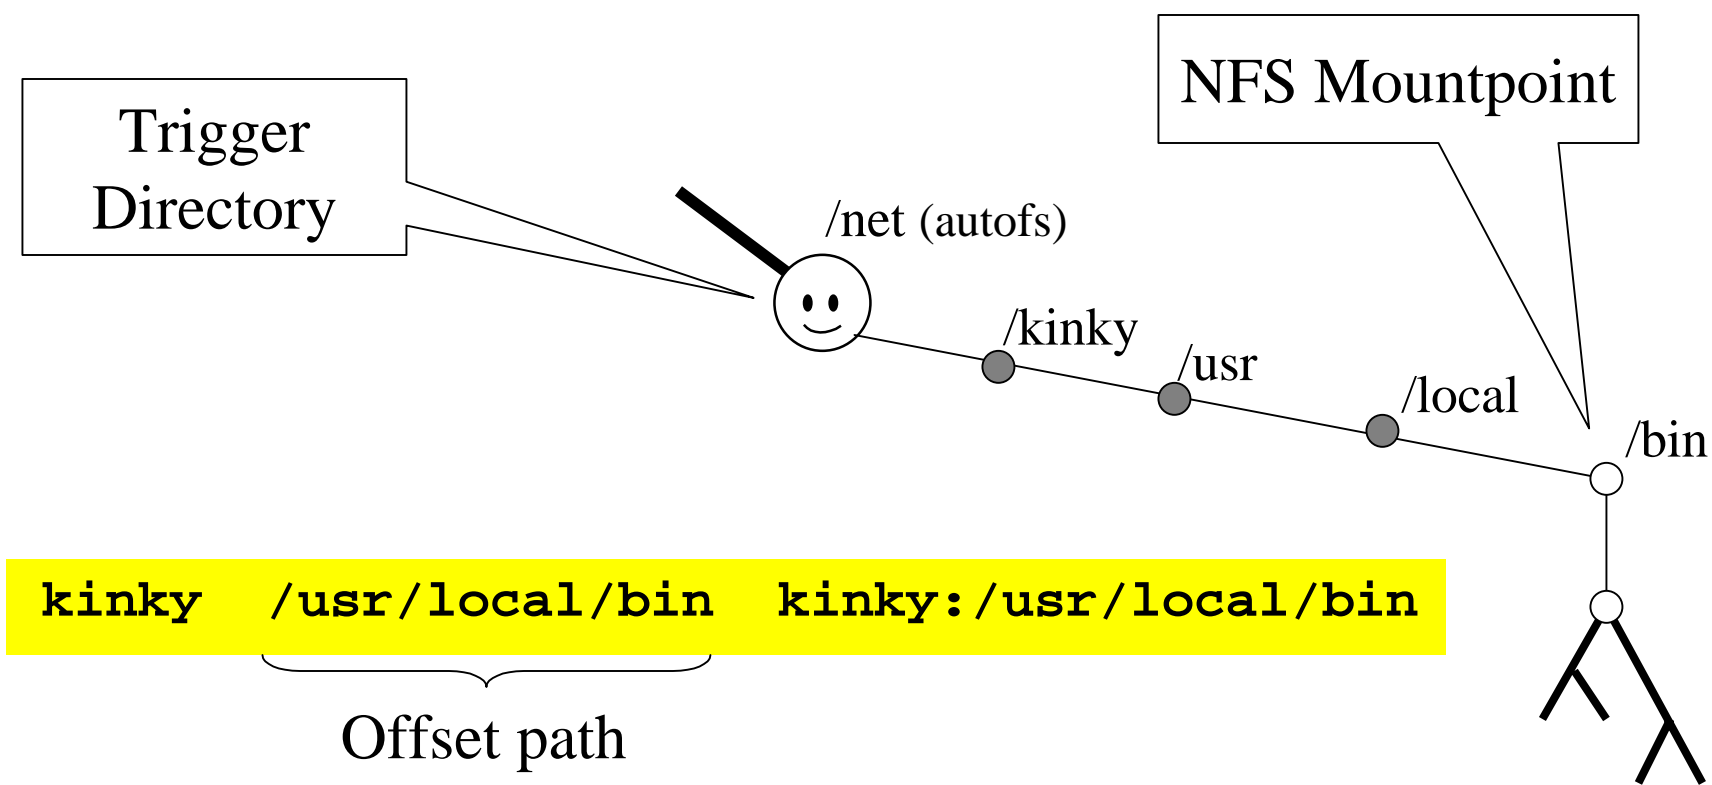
Maps managed in local files or
in a variety of name services.





October 12-14, 2004

Offset Mounts





October 12-14, 2004

Offset Mounts (cont)

Offsets useful for constructing
a mounted namespace

```
on10 \  
    /archives                stub:/export/archives \  
    /Codemgr_wsdata         stub:/export/gate/Codemgr_wsdata \  
    /deleted-files          stub:/export/gate/deleted-files \  
    /packages                stub:/export/packages \  
    /proto                   stub:/export/gate/proto \  
    /public                  stub:/export/gate/public \  
    /usr                     stub:/export/gate/usr
```

Offset paths



October 12-14, 2004

Mount Options (optional)

Map entries can have mount options for each mount that override the map default.

`bandersnatch`

`-ro,noquota`

`frumious:/jabber`

But keep them to a minimum!
Use mount defaults or map defaults



October 12-14, 2004

Multiple Mounts

A single map entry can initiate multiple mounts

```
key offset location { offset location } ...
```

```
on10 \  
    /archives          stub:/export/archives \  
    /Codemgr_wsdata   stub:/export/gate/Codemgr_wsdata \  
    /deleted-files     stub:/export/gate/deleted-files \  
    /packages          stub:/export/packages \  
    /proto             stub:/export/gate/proto \  
    /public            stub:/export/gate/public \  
    /usr               stub:/export/gate/usr
```



October 12-14, 2004

Replica Servers

A map entry can have multiple locations if the location is replicated.

```
key location {location} ...
```

```
pkg server1:/pkg server2:/pkg server3:/pkg
```

```
pkg server1,server2,server3:/pkg
```

```
pkg server1,server2(1),server3(2):/pkg
```


Weighting reduces preference



October 12-14, 2004

Replica Servers (cont)

- Replica list ordered by:
 - My host
 - Same server
 - My subnet (subject to weighting)
 - My network (subject to weighting)
- Solaris mounts entire list if read-only
- Otherwise, mount the best





October 12-14, 2004

Map Variables

Select a mount based on a client characteristic

```
/usr/local/bin server1,server2:/export/local/$OSNAME
```

using map variable text in exported path

```
share /export/local/linux  
share /export/local/solaris  
share /export/local/macosx  
share /export/local/netbsd  
share /export/local/hpux
```



October 12-14, 2004

Map Variables (cont)

Most useful map variables are predefined

Variable	Meaning	Derived From	Example
\$ARCH	Architecture type	uname -m	i86pc
\$CPU	Processor type	uname -p	i386
\$HOST	Hostname	uname -n	dinky
\$OSNAME	Operating system name	uname -s	solaris
\$OSREL	Operating system release	uname -r	5.7
\$OSVERS	Operating system version	uname -v	Generic
&	Insert key text	Entry key	



October 12-14, 2004

Key Substitution

The map key often reappears in the location

```
arnold      homeserve: /export/home/&  
brian      homeserve: /export/home/&  
charles    homeserve: /export/home/&  
jane       tinytots: /export/home/&  
kim        tinytots: /export/home/&  
oscar      homeserve: /export/home/&  
rumplestiltskin  tinytots: /export/home/&
```



October 12-14, 2004

Auto-unmounting

- Without auto-unmounting, mounts would “pile up”
- Every 5 min, automountd attempts to unmount
- Multiple mounts are unmounted in a bottom-up fashion.



October 12-14, 2004

Autofs Mounts

Map entries default to NFS, but other filesystems can be mounted, incl autofs!

```
src -fstype=autofs auto_src
```

Mount option

Location
is map name

Can use a hierarchy of autofs mounts to build a flexible, hierarchical namespace.



October 12-14, 2004

Autofs Mounts (cont)

Namespace using autofs hierarchy

auto_master

/shared auto_shared

auto_shared

```
home      -fstype=autofs auto_home
local     aa,bb,cc: /export/local/$ARCH
pkgs      -fstype=autofs auto_pkgs
src       -fstype=autofs auto_src
```

auto_home


```
arnold    homeserve: /export/home/&
brian     homeserve: /export/home/&
charles   homeserve: /export/home/&
jane      tinytots: /export/home/&
kim       tinytots: /export/home/&
oscar     homeserve: /export/home/&
rumplestiltskin tinytots: /export/home/&
```



October 12-14, 2004

Visible Directories

- Automounter directories used to show only mounted filesystems
 - To avoid “mount storms”






October 12-14, 2004

Visible Directories (cont)

- Solaris 2.6 provided fully visible directories via delayed mounts.
- Includes a -nobrowse option for huge directories, e.g. /home /net





October 12-14, 2004

Recommendations

- **Keep it simple!**
 - Avoid the use of mount options, map variables, multiple mounts, etc unless they provide real value.
 - Not all automounters are equal. Some cannot handle fancy map technology.
- **Avoid maps in files**
 - Yields non-uniform namespace
 - Difficult to distribute, maintain




October 12-14, 2004

Recommendations (cont)

- Avoid direct maps
 - use indirect maps.
- Plan your namespace
 - don't put everything in one map!
- Test your automounter implementation at Connectathon.



October 12-14, 2004



Questions & Answers