



October 12-14, 2004

# Autofs NG – The new Linux automounter

Mike Waychison

Linux Software Engineering

Sun Microsystems, Inc

[michael.waychison@sun.com](mailto:michael.waychison@sun.com)



October 12-14, 2004

# Agenda

- History of Autofs
- Summary of functionality
- Comparisons
- Autofs NG



October 12-14, 2004

# History of Linux Autofs

- Autofs 3 - created by H. Peter Anvin ~ 1997
- Autofs 4 - created by Jeff Fitzhardinge
- Autofs 4.1 – maintained by Ian Kent
- AMD - BSD4.4



October 12-14, 2004

# Autofs 3

- Supported basic autofs functionality
  - No direct maps
  - No browsing
  - No multimounts
  - No nsswitch.conf support
    - Resulted in Linuxisms (yp:auto\_home)



October 12-14, 2004

# Autofs v4

- Based off v3
- Added rudimentary support for multimounts
- Yay! /net access (sorta..)
  - Used custom program maps to simulate the special -hosts map.



October 12-14, 2004

# Autofs v4 (cont)

- Multimounts are mounted all at the same time.
  - Originally 'all or nothing' approach
  - 'nostrict' option later added to deal with errors.
- Still no direct mount support.
- Still no browsing support.



October 12-14, 2004

# Autofs v4.1

- Ghosting (like browsing)
  - Indirect map entries picked up by the wildcard key stick around
- Direct(ish) map support
  - Direct maps are translated into indirect maps internally -> leads to 'covering' directories



October 12-14, 2004

# AMD

- Different syntax
  - more powerful
  - more confusing (IMHO)
- Portable
  - Implemented in userspace
  - Ported to Tru64, BSDi, Linux, Solaris, Irix and AIX





October 12-14, 2004

# AMD (cont)

- Uses a userspace NFS server
  - Traps on lookups of non-existent files
  - Processes map entry
  - Mounts the filesystem in some other location
  - Returns a newly created symlink that points to the mounted filesystem
- Leads to the classic `/bin/pwd` problem



October 12-14, 2004

# Design problems(v3->4.1)

- Kernel -> Userspace:
  - Communication done through pipe passed in at mount time.
  - Broken pipe -> broken mount
- Userspace -> Kernel
  - loctls on the base directory of the indirect mount
  - Does not work well with direct mounts (directory is covered!)



October 12-14, 2004

# Feature Comparison

	v3	v4	v4.1
Indirect Maps	Y	Y	Y
Direct Maps	N	N	Sort of
Browsing	N	N	Sort of
Native -hosts	N	N	N
Multimounts	N	Sort of	Soft of



October 12-14, 2004

# Goals

- Support:
  - indirect mounts
  - direct mounts
  - multimounts
  - browsing
- Allow transparent use of maps by Commercial Unix (Solaris in particular)



October 12-14, 2004

# Secondary Goals

- Support Linuxisms
  - mapname nameservice prefixes (program:, yp:, nisplus: ldap:)
  - Hesiod support
  - CD changer support
- Iterable program maps



October 12-14, 2004

# Design – The Situation

- Current design had flaws
  - Communications channels did not allow direct mounting nor multimounts
  - Linux per-process mount table namespaces would break autofs mounts
  - Current design would get 'stuck' if deamon killed



October 12-14, 2004

# Design – Requirements

- New means of communication
  - Must be able to communicate with direct mounts
- Remove duplicated state between kernel and userspace
- Be able to deal with userspace updates and brokenness
  - w/o hanging the machine :)



October 12-14, 2004

# Design – Trapping

- Provide magic filesystem
  - Two modes of operation
    - Direct – Traps on entry into root directory
    - Indirect – Traps on entry into any child directory
  - Trapping performed by calling out to userspace (/sbin/autofs)
    - Done in namespace context of the trapping application (but as root)
    - Assumes administrator trusts /sbin/autofs





October 12-14, 2004

# Design – Expiry

- Handled in the kernel
  - Provided as a generic facility
- Atomic expiry of trees
  - Trees can be tagged to expire as part of an expiring subtree
- Flaw – Userspace has no notification that the expiry occurred.



October 12-14, 2004

# Example – Indirect

```
/etc/auto_master:
```

```
/home      /etc/auto_home
```

```
/etc/auto_home:
```

```
mikew      bignfs.mydomain.com: /export/mikew
```

```
stews      bignfs.mydomain.com: /export/stews
```



October 12-14, 2004

# Example – Indirect (cont)

- When automount(8) is run:

```
mount -o indirect,mapname=/etc/auto_home -t  
autofs none /home
```

- When /home/mikew is accessed,  
kernel calls:

```
MAPNAME=/etc/auto_home \  
MAPKEY=mikew \  
MOUNTFD=3 \  
/sbin/autofs mount
```



October 12-14, 2004

# Example – Indirect (cont)

- /sbin/autofs:
  - Looks up \$MAPKEY in \$MAPNAME
  - Mounts found filesystem
  - Moves filesystem to directory specified by \$MOUNTFD
  - Returns 0
- Kernel lets userspace walk into the newly available filesystem



October 12-14, 2004

# Example – Direct

- Direct mounts are done in the same way:
- ```
mount -o direct,mapname=/etc/auto_direct,mapkey=/full/path -t  
autofs none /full/path
```
- **MAPNAME=/etc/auto\_direct**
- **MAPKEY=/full/path**
  - (specified at autofs mount time, no runtime resolution done)



October 12-14, 2004

# Lazy Multimounts

- Handle by installing direct mounts within mounted filesystems
- Introduces the concept of entry offsets
- Example:


```
/usr/src      /      bignfs:/export/sources \  
              linux    othernfs:/export/linux
```




October 12-14, 2004

# Multimounts (cont)

```
 /usr/src          /          bignfs:/export/sources  \  
                   linux    othernfs:/export/linux
```



← -t ext3 /dev/hda1



← -t autofs  
-o mapname=/etc/auto\_direct  
-o mapkey=/usr/src




October 12-14, 2004

# Multimounts (cont)



```

/usr/src          /          bignfs:/export/sources  \
                  linux    othernfs:/export/linux
  
```




```
-t ext3 /dev/hda1
```

```
-t autofs
-o mapname=/etc/auto_direct
-o mapkey=/usr/src
```



```
-t nfs bignfs:/export/sources
```



```
-t autofs
-o mapname=/etc/auto_direct
-o mapkey=/usr/src
-o mapoffset=linux/
```






October 12-14, 2004

# Multimounts (cont)

```

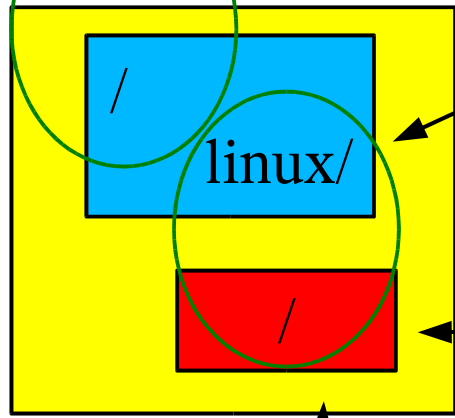

/usr/src          /          bignfs:/export/sources  \
                  linux      othernfs:/export/linux

```



```
-t ext3 /dev/hda1
```

```
-t autofs
-o mapname=/etc/auto_direct
-o mapkey=/usr/src
```



```
-t nfs bignfs:/export/linux
```

```
-t autofs
-o mapname=/etc/auto_direct
-o mapkey=/usr/src
-o mapoffset=linux/
```

**Mounted Atomically / Expires Atomically**



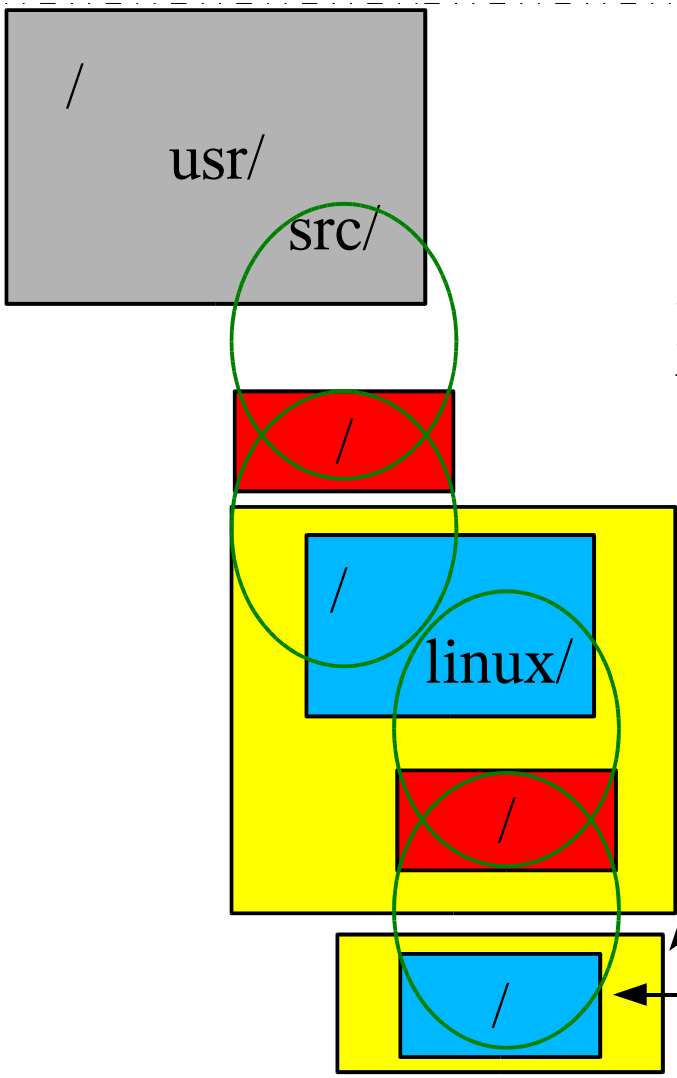
October 12-14, 2004

# Multimounts (cont)

```

/usr/src          /          bignfs:/export/sources  \
                 linux    othernfs:/export/linux

```



Each level expires atomically but independent of each other

`-t nfs othernfs:/export/linux`



October 12-14, 2004

# Project Status

- Alpha
- Need to work on Linuxisms
- Pass most of autofs connectathon testsuite
- Kernel + Userspace GPL  
<http://autofsng.bkbits.net>  
<ftp://ftp-eng.cobalt.com/pub/users/ssmith>