

Autofs Enhancements


Ricardo Labiaga
labiaga@Eng.Sun.Com
SunSoft

Overview

- **Lazy Mounting of Hierarchies**
- **Improved Concurrency**
- **Browsability of Indirect Maps**
- **XFN user-relative context binding**
- **Avoid use of the Space Protocol**

Existing *Autofs* Bottleneck

- **Concurrent lookups serialized by *autofs***
 - ❑ **Blocking done one level too high**
 - ❑ **No benefit from MT *automountd***



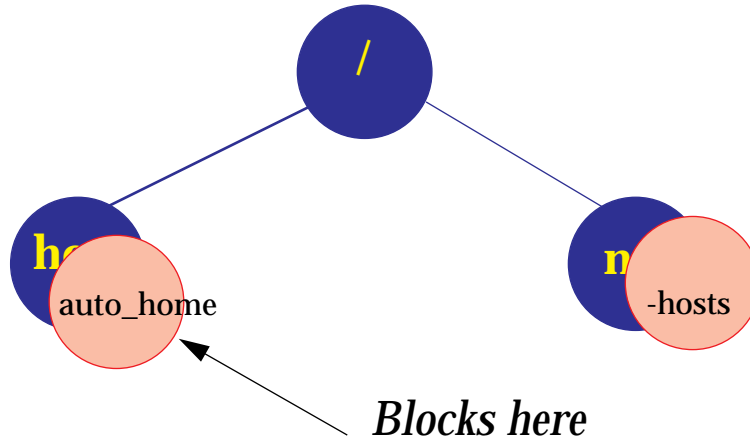
auto_home

-hosts

\$ /usr/lib/autofs/automount -v
automount: /home mounted
automount: /net mounted

- mount -F autofs auto_home /home
- mount -F autofs -hosts /net

```
stat("/home/joe", &buf);
```






jurassic:/export/home/joe

Automountd creates node and performs the mount

New Autofs

- **Mountpoint created by *autofs* module**
 - ❑ **Block on mountpoint node**
 - ❑ **Concurrent lookup of different keys**


```
stat("/home/joe", &buf);
```



Autofs creates node, then calls automountd to perform mount.



jurassic:/export/home/joe

Browsability

- **Needed for global name space construction**
 - ❑ **XFN**
 - ❑ **/shared**
- **Existing autofs limitation**
 - ❑ **Only mounted entries are visible**
- **Requirement**
 - ❑ **Avoid mount storms**

Solution: Modified mount policy


Browsability - Details

- ***Autofs* triggers mounts of real filesystem on**
 - ❑ **open**
 - ❑ **lookup of component within filesystem**
- ***Autofs/automountd* new functions**
 - ❑ **AUTOFS_READDIR**
 - ❑ **AUTOFS_LOOKUP**

auto_home:

aikman cowboys:/export/home/&
emmit cowboys:/export/home/&
sanders cowboys:/export/home/&
novacek cowboys:/export/home/&
moose cowboys:/export/home/&

Already Mounted



- 1) **getdents (/home, 30)**
returns: **emmit, brown, ps**
n_offset: **AUTOFS_READDIRCOOKIE**
- 2) **getdents (/home, 30)**
returns: **aikman, sanders, novacek**
n_offset: **AUTOFS_READDIRCOOKIE + 3**
- 3) **getdents (/home, 30)**
returns: **moose**
sets **EOF**

- ***automountd* caches map listing**
 - ❑ **entire map vs. entry at a time**
- **async thread in automountd frees cache periodically**
- **No autonodes are created on readdir()**
- **lookup() will create the autonode**
 - ❑ **freed when reference is released**

Browsability Performance

- **/home [~7,000 entries in NIS+]**
 - ❑ **15 seconds on cold cache**
 - ❑ **3 seconds on warm cache**
- **/ws [~ 80 entries]**
 - ❑ **instantaneous**
- **/net [~ 30,000 entries in NIS+]**
 - ❑ **You don't want to know :-) on cold cache**
 - ❑ **6 seconds on warm cache**
- **Room for improvement**

Other Features

- **Avoid use of Space Protocol**
- **XFN user-relative binding lookup support**
 - ❑ **Lookup of /xfn/thisuser resolves to different path depending of EUID.**
 - ❑ **/xfn/thisuser -> /xfn/labiaga**
 - ❑ **/xfn/thisuser -> /xfn/advani**