

# Automounter Improvements

Ashok Advani  
Sunsoft

*advani@eng.sun.com*

---

Connectathon '97  
Feb 28, 1997



## Objective

- To evaluate and discuss improvements made to the Solaris automounter in the upcoming Solaris 2.6 release, and in the ongoing Fileset Migration/Replication project

---

Connectathon '97  
Feb 28, 1997

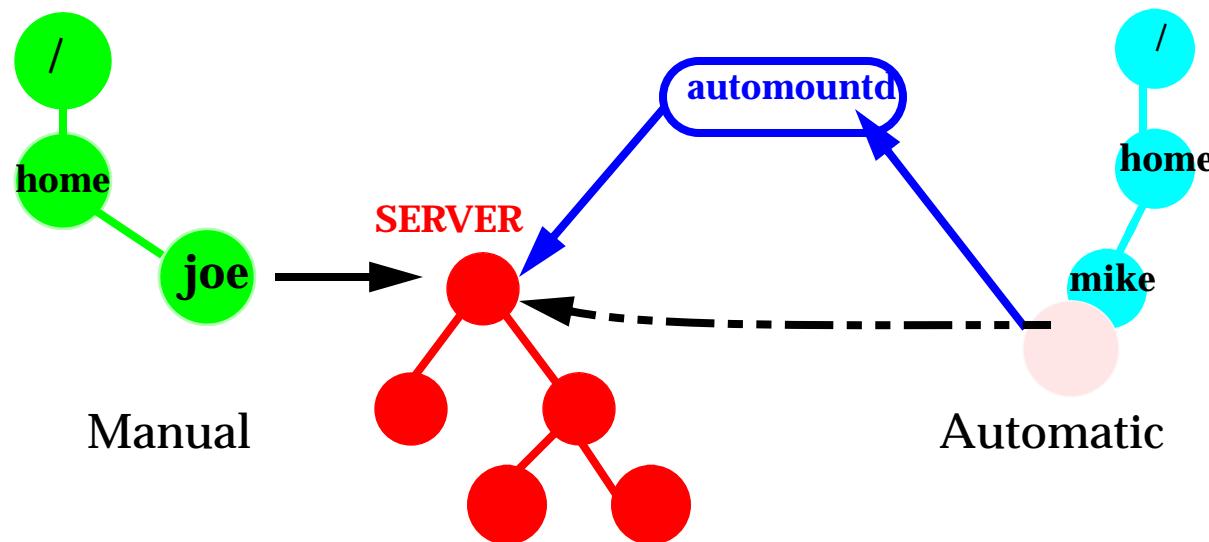


## SCOPE

- **Solaris Autofs Automounter**
- **Solaris 2.6 Automounter Features**
- **Atomic Unmounting and Mounting Hierarchies**
- **Automounter wish-list**
- **Summary**

## Solaris Autoofs Automounter

- automatic, on-demand mounting and unmounting of filesystems



## Solaris 2.6 Automounter Features

- Fully multithreaded automounter daemon
  - 2.5.1 Daemon is not completely threaded
  - daemon hangs => system hangs
    - mount from dead server =>daemon hang
    - recursive calls to daemon => daemon hang
  - Performance

## Solaris 2.6 Automounter Features

- **Browsability**

- Only mounted entries are visible in Solaris 2.5.1
- Needed for global name space construction (xfn)

% ls /home                    Solaris 2.6  
advani (+ 7000 more!!!)

- Performance

/ws      80 entries      instantaneous  
/home    7000 entries    15 seconds on cold cache  
/net     30,000 entries => you don't want to know!!

- Applications that mount the world!

/net   -hosts   -nobrowse

## Solaris 2.6 Automounter Features

- **Hierarchical Mounting of Filesystems**

- Server exporting a hierarchy of filesystems can cause mount storms in Solaris 2.5.1

Example: NFS server x exports 27 filesystems /, /a, /b,...,/z  
client % cd /net/x

- ALL 27 NFS mounts (/net/x, /net/x/a,/net/x/b,....,/net/x/z) take place up front!

- Solaris 2.6 does better on-demand mounting

Example: Only /net/x is NFS mounted.

- Place AUTOFS trigger nodes on /net/x/b, ...., /net/x/z

- Performance

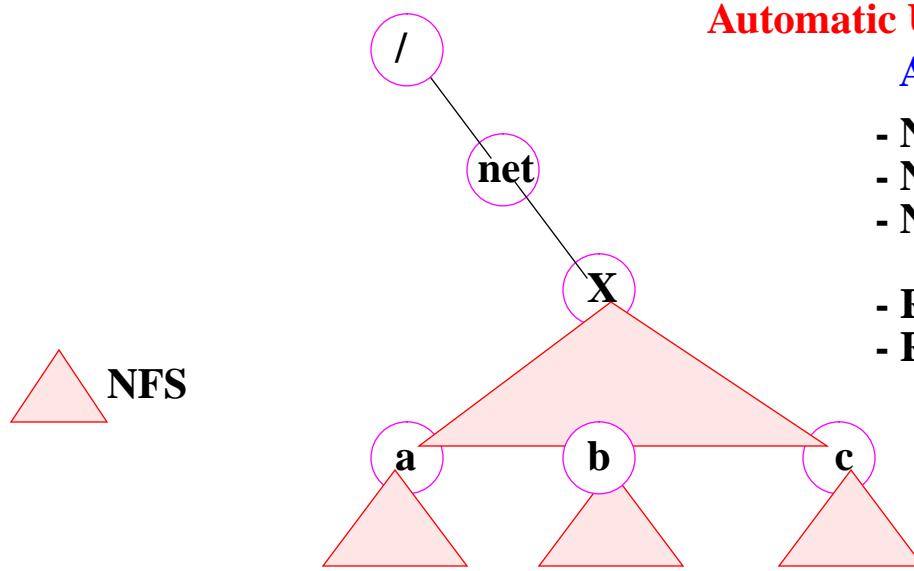
## Solaris 2.6 Automounter Features

- **Duplicate Request Cache**

- Recurrence of non-idempotent operations, eg., mount
- Solaris 2.5.1 automounter supports 16 entry duplicate request cache based on xid's
- Solaris 2.6 supports a time-based cache
  - hash table based on xid
  - time-based removal of entries in the cache (LRU)
  - current timer set at 900 seconds
  - built into the RPC library

# Atomic Unmounting/Mounting Hierarchies

- The Unmount/Remount Problem in Solaris 2.5.1



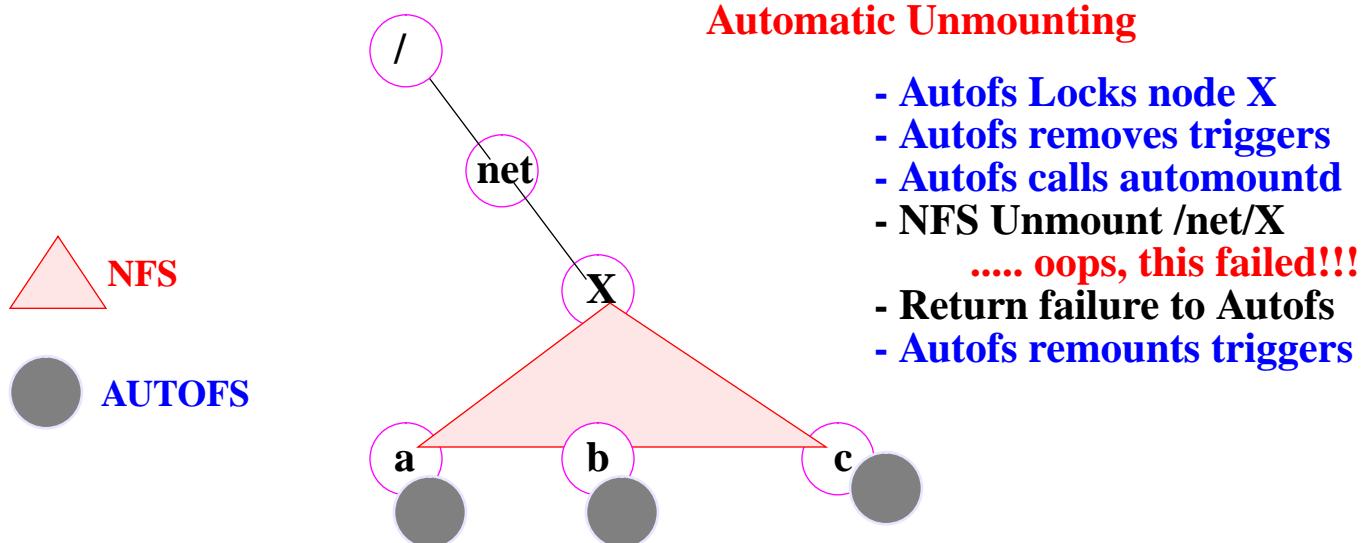
## Automatic Unmounting

Autofs calls automountd

- NFS Unmount /net/X/a
- NFS Unmount /net/X/b
- NFS Unmount /net/X/c
  - ..... oops, this failed!!!
- Remount /net/X/b
- Remount /net/X/a

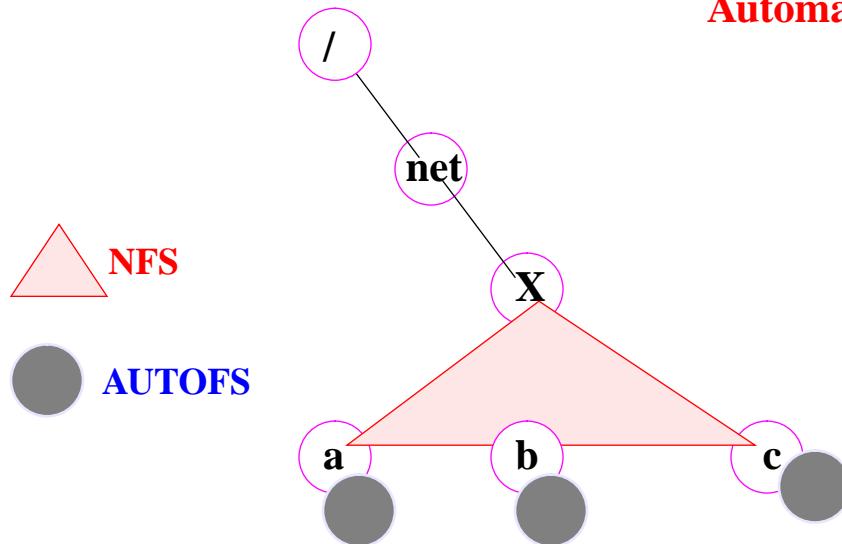
# Atomic Unmounting/Mounting Hierarchies

- The Unmount/Remount Problem in Solaris 2.6



# Atomic Unmounting/Mounting Hierarchies

## •Atomic Unmounting Design (Fileset Project)



### Automatic Unmounting

- Autoofs Locks node X
- Autoofs calls VFS\_BUSY(vfsp, expected\_ref)
- Autoofs removes triggers
- Autoofs calls Automountd
- NFS Unmount /net/X  
..... oops, this failed!!!
- Return failure to Autoofs
- Autoofs remounts triggers

# Atomic Unmounting/Mounting Hierarchies

- **Hierarchical Mounting problem in Solaris 2.5.1**
  - Autofs calls automountd
  - Automountd reads list of hierarchical mounts from map
    - Example: Server X exports /, /a, /b, /c, /d
    - for (i = 0; i < number\_mounts; i++) {
      - build\_nfs\_args();
      - nfs\_mount();
  - What if automountd dies in the middle? Autofs is confused!

# Atomic Unmounting/Mounting Hierarchies

- **Hierarchical Mounting Problem in Solaris 2.6**
  - Autofs calls automountd
  - Automountd reads list of hierarchical mounts from map
    - Example: Server X exports /, /a, /b/, /c, /d
    - nfs\_mount(/net/X);
    - for (i = 0; i < trigger\_mounts; i++)
    - build\_automountd\_args\_action\_list();
  - Automountd returns action list to Autofs
  - Autofs mounts trigger nodes
  - What if automountd dies in the middle? Autofs is still confused!

# Atomic Unmounting/Mounting Hierarchies

- **Atomic Mounting Design (Fileset Project)**
  - Autofs calls automountd with cookie vnode
  - Automountd reads list of hierarchical mounts from map
    - Example: Server X exports /, /a, /b/, /c, /d
    - mounta\_list = build\_nfs\_mounta(/net/X);
    - for (i = 0; i < trigger\_mounts; i++)
    - mounta\_list->next = build\_autofs\_mounta();
    - \_autofssys(AUTOFSSYS\_MOUNT, cookie, mounta\_list);
  - What if automountd dies in the middle? Autofs simply checks the vnode to see if its mounted or has trigger nodes installed

## Automounter Wish-list (Fileset Project)

- Rearchitect User-level mounting to use shared libraries
- Autofs Name Space Synchronization
  - Server X exports /a, /b and its mounted
  - Sysadmin now decides to export /c
  - Must wait for unmount until /c is seen by the automounter
- Mount option concatenation
  - mntpnt -nosuid server-x:/  
/a -rw server-x:/a
  - override the options vs. concatenate them
- in-kernel /etc/mnttab

## **Summary**

- Continuing to work to enhance the automounter
- Solaris 2.6 Automounter Features
- Atomic Unmounting and Mounting Hierarchies
- Automounter wish-list