SunSoft
A Sun Microsystems Company

FORCED UNMOUNT

Devesh Shah

03/11/99

Forced Unmount

- **Problem Description:**

Currenly, if NFS gets in certain states or a disk fails permanently, an umount command will fail with EIO. System administrators have to use “fuser -k” to find and kill all the processes and try to unmount again.

The only guaranteed method is to reboot the client !!

What is Force Unmount?

- Allows to forcibly unmount the file systems

  Example: umount -f /mnt
- It must be used when reboot is the only alternative
- Currently supports NFS and UFS file systems
- The feature will be available in Solaris 2.8 release

Why need force unmount?

- Customer requirement
- Cluster technology group needs the feature
- Other vendors have the feature available
- Failover mechanism works for read only file systems

Force unmount Design

- File system being forcibly unmounted simply disappears from the namespace
- The existing processes using the file system return EIO
- Prevents the system to consume the new resources and clean up the old resources as much as possible
- NFS client releases the locks of a forced unmounted file system
- File system resources are cleaned up in a lazy fashion

Force unmount Design

- **UMOUNT(1M)**
  generic and FS specific unmount commands
- **/usr/include/sys/mount.h**
  #define MS_FORCE 0x0400
- **New system call - umount2(char *mntpnt, int flag)**
- **VFS_UNMOUNT(vfs_t *vfsp, int flag, cred_t *cr)**
- **VFS reference count is required**
- **File systems that do not support the feature return ENOTSUP**
- **In case of a nested file systems, top level file system is forced unmounted**

Implementation (NFS file system)

- **nfs3_unmount(vfs_t *, int, cred *)**
- **Check for MS_FORCE flag**
- **rfscall() will not send the RPC request and returns EIO**
- **Dirty pages are cleaned up - B_INVAL flag is set**
- **NFS nodes (rnodes) are not recycled**
- **Removal of mi_rootvp - a pointer to root vnode**
- **nfs[3]_freevfs - free up the resources associated with mntinfo structure**

Implementation(NFS file system)

- The lockfs protocol can not be used to lock the file systems due to different semantics
- Blocked processes are not woken up
- Vfs referece count is balanced
- fuser -k can be run after file system is forced unmounted

Chages in other file systems

- **UFS file system**
  - **Implementaion for force unmount already exist**
- **Cachefs and Autofs**
  - **Need to evaluate if support is necessary**
- **Other file systems**
  - **ENOTSUP error**
- **Other changes**
  - **truss command is changed to reflect umount2() syscall**
  - **kadb macros is changed to see the vfs referece count**

IMPLICATIONS

- **When file operation is attempted on mmap'ed file, an application will get SIGSEGV or SIGBUS**
- **Any data prior to force unmount was not committed would be lost**
- **The processes may not receive any indication that a file system has been forcibly unmountd**