

The NFS  
Implementation  
for  
Solaris 2.1

John Corbin

corbin@sun.com

Network Resources

Sun Microsystems Computer Corp.

December 8, 1992

# Overview

- Overview of NFS
- What Has Changed In Solaris 2.1
- NFS Client Modifications
- NFS Server Modifications
- Administration Issues
- Switching Your NFS Server To Solaris 2.1

# Definitions

- Network File System (NFS)

A simple remote procedure-based protocol that allows remote access to file data and file attributes.

- File Data

Data contained within a file.

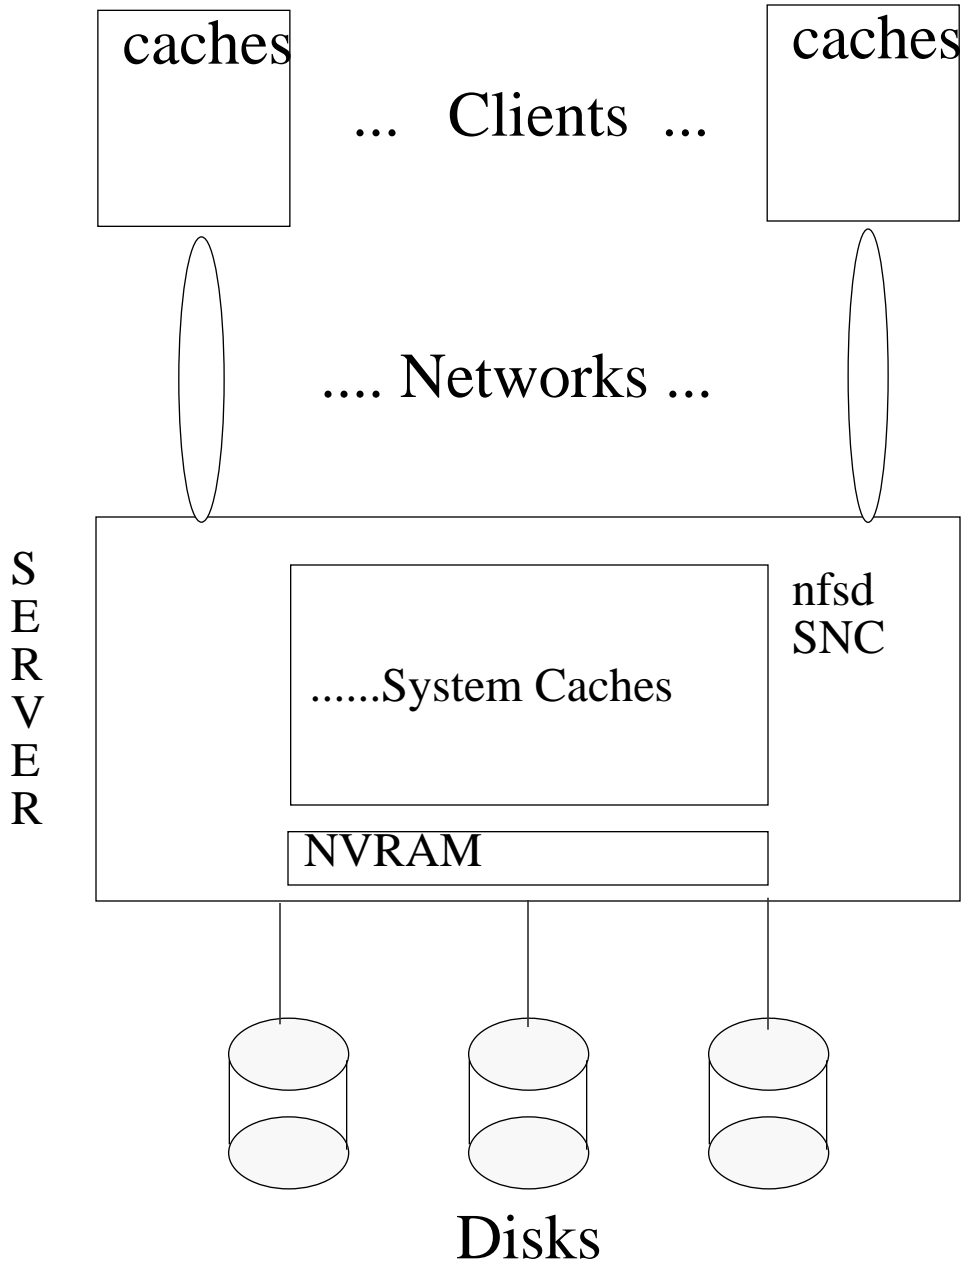
- File attributes

Also referred to as file meta-data.

EX:

file size, file owner, read/write/execute permissions

# NFS Components



## What's New

- Symmetric Multiprocessor Implementation

NFS Kernel Concurrency ~50%

It will continue to improve.

- NFS Running On Top Of STREAMS

- UDP Checksums On By Default

- *biod* Processes Converted

Are now dynamically allocated kernel threads.

- *nfsd* Processes Undergoing Conversion

There are the usual number of *nfsd* processes, which typically do nothing.

Dynamically allocated kernel threads do most the work.

- Support Added For Kerberos Authentication

- */tmp* mounted on *tmpfs* by default

Files put in */tmp* are guaranteed to not survive a reboot.

Do not store permanent files in */tmp*

# NFS Client Modifications

- Block I/O Handled By Kernel Threads

Default Number of Async I/O Threads Per Mount = 8

SunOS 4.X Had Default Of 4 biods Per Client

Override Example: set nfs:nfs\_async\_threads=5

Read Ahead Default = 1

Same as SunOS 4.X

Override Example: set nfs:nfs\_nra=2

Maximum Queued Async Requests = 100

- Dynamic Setting Of Retransmission Parameters

I Haven't Had A Chance To Experiment With It

Turned On By Default

Turn Off - set nfs:nfs\_dynamic=0

# NFS Server Modifications

- Performance Improvements

Streamlined the NFS server implementation.

*/kernel/drv/nfsfpmo*d Introduced

No more socket buffer size limitations - not using socket buffers.

Request queue is arbitrarily large, but not too large.

Dynamically allocated kernel threads

Maximum of 12 per CPU

- Kerberos Authentication Supported For NFS Mounts

- Export *tmpfs* File Systems

Great for diskless clients that need to remotely mount */tmp*

- NFS Loop Back Mounts

Neat for proof of concept on a clean NFS implementation

Not a good idea for end users

Why introduce NFS overhead to UFS and *tmpfs* mounts?

# Administration Issues

- The NFS administration conversion is fairly painless.

*/etc/exports* is now */etc/dfs/dfstab*

*exportfs* still there

share is SVR4 way

Same semantics as *exportfs*, different command format

*mount* and *umount* have modified option syntax

*mount -a* is now *mount\_all*

*umount -a* is now *umount\_all*

*rpc.mountd* is now *mountd*

*etherfind* replaced by *snoop*

*nfsstat*, *netstat*, are all still there

- New Book Out

“Solaris System Administrator’s Guide”, Janice Winsor,  
Ziff-Davis Press (ISBN 1-56276-080-7).

## Switching NFS Servers To Solaris 2.1

- Transition when it makes sense to do so.

Client doesn't care what OS the server is running, as long as the server is serving.

Keep Uni-Processor platforms running SunOS 4.1.3

Convert SS600s to Solaris 2.1 if possible

Keep NIS maps and administration tools on 4.X machines if desired

- Unbundled Product Ports Are Coming

Prestoserve

Sun Network Coprocessor