Subject: Re: NFS mounts disappearing from OAS servers
Posted by nfshelp on Mon, 22 Oct 2007 21:34:23 GMT

Here's some further information:

What mount option was used?

Mnt point	options	server	exported filesystem
common	-ro,soft,intr,retrans=9,tcp	celerra01_cge0:/Mid-Range_Export_DBA_Common	

Are you automounting? YES

Are the mounts clustered? no

What does "nfsstat –m" say?

```
[root@hsg05lt /]# nfsstat -m
nfsstat: invalid option -- m
Client rpc stats:
calls      retrans    authrefrsh
6119       0          0
Client nfs v3:
null     getattr   setattr   lookup    access     readlink
0    0%  1600  26% 299   4%  669   10% 347    5%  0      0%
read      write     create    mkdir     symlink    mknod
73    1%  2721  44% 21    0%  0     0%  0      0%  0      0%
remove    rmdir     rename    link      readdir    readdirplus
0    0%   0     0%  0     0%  0     0%  0      0%  32     0%
fsstat    fsinfo    pathconf  commit
104   1%  104   1%  0     0%  149   2%


[root@rac02lt /]# nfsstat -m
nfsstat: invalid option -- m
Server rpc stats:
calls      badcalls   badauth    badclnt    xdrcall
0          0          0          0          0
Client rpc stats:
calls      retrans    authrefrsh
14108      0          0
Client nfs v3:
null     getattr   setattr   lookup    access     readlink
0    0%  1210  8%  592   4%  147   1%  525    3%  0      0%
read      write     create    mkdir     symlink    mknod
13    0%  11139 78% 47    0%  0     0%  0      0%  0      0%
remove    rmdir     rename    link      readdir    readdirplus
```

0   0% 0   0% 0   0% 0   0% 0   0% 60   0%
fsstat   fsinfo   pathconf commit
203   1% 60   0% 0   0% 112   0%

What OS is this being run on?
Client side
Hsg05lt - Red Hat Enterprise Linux AS release 3 (Taroon Update 7)
        Rac02lt - Red Hat Enterprise Linux AS release 4 (Nahant Update 3)

What's the timeout?
The default timeout is .7 sec. with the retrans set to 9 I believe that we will get 7 minor timeouts (at which point we are at 60 secs.), this will generate a major timeout. We will get a mount failure when we have hit 3 major timeouts.