

DragonFlyBSD - Bug #2499

DRAGONFLY_3_2 lockd not responding correctly

01/22/2013 12:41 AM - Nerzhul

Status:	In Progress	Start date:	01/22/2013
Priority:	Urgent	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
Hello, i must use lockd for concurrent access on a webserver with nfs extended storage. There is some concurrent access and lockd isn't responding correctly.			
On the NFSv3 client, timeout appears and console logs: nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd not responding nfs server A.B.C.65:/nfs/fbsd_pkg: lockd is alive again			
After "netstat -an -f inet" i see there is a queue on rpc socket			
netstat -an -f inet			
Active Internet connections (including servers)			
Proto Recv-Q Send-Q Local Address Foreign Address (state)			
tcp4 0 0 A.B.C.65.nfsd WebCluster1.977 ESTABLISHED			
tcp4 0 0 A.B.C.65.nfsd WebCluster1.611 ESTABLISHED			
tcp4 0 0 localhost.smtp *.* LISTEN			
tcp4 0 0 *.ssh *.* LISTEN			
tcp4 0 0 *.1017 *.* CLOSED			
tcp4 0 0 *.1020 *.* LISTEN			
tcp4 0 0 *.nfsd *.* LISTEN			
tcp4 0 0 *.1023 *.* LISTEN			
tcp4 0 0 *.1022 *.* LISTEN			
tcp4 0 0 *.sunrpc *.* LISTEN			
tcp4 0 0 A.B.C.65.nfsd A.B.C.96.811 ESTABLISHED			
tcp4 0 0 A.B.C.65.nfsd WebCluster1.972 ESTABLISHED			
tcp4 0 48 A.B.C.65.ssh 129.175.196.190.60067 ESTABLISHED			
udp4 0 0 *.918 *.*			
udp4 0 0 A.B.C.65.1028 ntp.u-psud.fr.ntp			
udp4 456 0 *.1017 *.*			
udp4 18656 0 *.1018 *.*			
udp4 0 0 *.nfsd *.*			
udp4 0 0 *.1021 *.*			
udp4 0 0 *.1020 *.*			
udp4 0 0 *.1022 *.*			
udp4 0 0 *.sunrpc *.*			

When i see that, i make tcpdump -nni em0 to see what's happening:

```
22:12:42.781597 IP 10.117.100.95.961 > 10.117.100.65.1017: UDP, length 212
22:12:48.801935 IP 10.117.100.95.961 > 10.117.100.65.1017: UDP, length 212
22:12:54.669917 IP 10.117.100.95.961 > 10.117.100.65.1017: UDP, length 212
22:13:00.148965 IP 10.117.100.95.961 > 10.117.100.65.1017: UDP, length 212
```

After a little time, lockd respond to all request, but many failed because of timeout

On the dragonflyBSD server i can see this in /var/log/messages

```
Jan 21 22:14:19 webfiler1 rpc.lockd: duplicate lock from WebCluster1.srv.
Jan 21 22:14:19 webfiler1 last message repeated 3 times
Jan 21 22:14:19 webfiler1 rpc.lockd: no matching entry for WebCluster1.srv.
Jan 21 22:14:29 webfiler1 dntpd[571]: issuing offset adjustment: 0.026637s
Jan 21 22:14:44 webfiler1 rpc.lockd: rpc to statd failed: RPC: Timed out
Jan 21 22:14:44 webfiler1 rpc.lockd: duplicate lock from WebCluster1.srv.
Jan 21 22:14:44 webfiler1 last message repeated 3 times
Jan 21 22:14:44 webfiler1 rpc.lockd: no matching entry for WebCluster1.srv.
```

I think there is a problem on DragonFlyBSD which queue many lockd requests.

History

#1 - 01/22/2013 12:39 PM - tuxillo

- Status changed from New to In Progress

Hi,

It's been told several times in the IRC channel that rpc.lockd wasn't working but it hasn't been up until now when we find someone using it and reporting it doesn't work.

If I recall correctly, there was a short discussion about it and syncing ours with FreeBSD's was considered, but I've seen no work towards it. I'll add this to our Projects Page (<http://www.dragonflybsd.org/docs/developer/ProjectsPage/>) for better visibility.

What is exactly your setup, if I may ask?

Cheers,
Antonio Huete

#2 - 01/22/2013 12:47 PM - Nerzhul

Hi Antonio,

I use a FreeBSD cluster which mount it's 3rd party FS via NFS share under DFly (/usr/local).
The system also mount /usr/ports/distfiles and /var/db/pkg.

/var/db/pkg contain a Sqlite DB for pkgng.

in /usr/local we have all 3rd party compiled ports for clustered web servers, with all web user datas in /usr/local/www/* (except pg and my DB which are on another server) and some different websites (owncloud, dokuwiki...)

The goal is to have a main freebsd which only mount required 3rd party folders, machine extensible and exactly same datas. (the goal is 3 great Apache).

As Francois Tigeot adviced me, i have disabled lockd via /etc/fstab nolockd option for /usr/local and /usr/ports/distfiles which are static, but not for /var/db/pkg which contains a sqlite db.

Then lockd problem is partially resolved. But it slows my "pkg info" command and other "pkg xx" commands (like portmaster uses).