

pfSense - Bug #2401

Mounting read-only after mounting read-write can be very slow on NanoBSD

04/27/2012 02:44 PM - Jim Pingle

Status:	Resolved	Start date:	04/27/2012
Priority:	High	Due date:	
Assignee:		% Done:	100%
Category:	Operating System	Estimated time:	0.00 hour
Target version:	2.1	Affected Architecture:	
Affected Version:	2.1		

Description

Mounting read-only after mounting read-write can be very slow on recent NanoBSD images on 2.1, based on FreeBSD 8.3

Does not seem to affect all CF cards, but most, especially slower ones.

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(1): time /etc/rc.conf_mount_rw  
0.457u 0.077s 0:01.04 50.0% 3184+477k 0+27io 0pf+0w
```

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(2): time /etc/rc.conf_mount_ro  
0.441u 0.112s 0:37.01 1.4% 3101+450k 0+597io 0pf+0w
```

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(4): time mount -o rw /  
0.005u 0.000s 0:00.00 0.0% 0+0k 0+3io 0pf+0w
```

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(5): time mount -o ro /  
0.000u 0.021s 0:35.56 0.0% 160+704k 0+542io 0pf+0w
```

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(2): time /sbin/mount -u -w -o sync,noatime /  
0.000u 0.006s 0:00.01 0.0% 0+0k 0+5io 0pf+0w
```

```
[2.1-DEVELOPMENT][root@pfSense.localdomain]/root(3): time /sbin/mount -u -r -f -o sync,noatime /  
0.000u 0.018s 0:28.37 0.0% 64+2680k 0+407io 0pf+0w
```

Normally the read-only mount step should only take a few seconds, 5 or so.

As a consequence, many operations on the firewall take a very long time to work as it will pause for >30 seconds every time something re-mounts ro during the process. Reboot takes a very long time as well.

There is a lot more data in this forum thread:

<http://forum.pfsense.org/index.php/topic.48256.msg257717.html#msg257717>

A few more data points:

- Enabling TRIM on the slice - no effect, still slow
- Enabling softupdates on the slice - no effect, still slow
- Running sync;sync;sync; before mounting ro - no effect, still slow
- The mount process sticks in biowr state, which is disk i/o
- While mount is stuck, the system is somewhat unresponsive - can't run a process or connect with ssh, but does appear to continue routing.
- If you mount ro twice in a row, the second is fast, so it's only a problem when transitioning from read-write to read-only.

Truss on the slow mount process:

<http://pastebin.com/?page=p&id=4f9ace216b5cb>

top -SH while mount is happening:

<http://pastebin.com/?page=p&id=4f9ae2a9bc080>

systat -iostat :

<http://pastebin.com/?page=p&id=4f9ae3017b4ef>

ststat -vmstat during and after :

<http://pastebin.com/?page=p&id=4f9ae3f4c746a>

Because this is NanoBSD, DMA is already disabled and so is write caching.

History

#1 - 04/27/2012 02:47 PM - Jim Pingle

- Description updated

#2 - 08/20/2012 06:00 AM - Marcus Brown

Hi,

I've been having this problem with NanoBSD too. However, I just noticed that when I updated Safari to v6.0 the problem disappeared. Previously I was using Chrome, and the problem was apparent. I never did any timing tests on the command line, but I was noticing very long >30 seconds to save any new setting using the web GUI.

#3 - 08/28/2012 02:04 PM - Ermal Luçi

- Status changed from New to Feedback

This should have gone away with the fixes that went in for reference counting fixes.

#4 - 08/28/2012 03:12 PM - Jim Pingle

- Status changed from Feedback to New

It may be improved a little but it's still slow even just running mount, no refcounts involved. Still only affects certain CF cards. One I have is fast, another is extremely slow.

```
: time /sbin/mount -u -w -o sync,noatime /
0.006u 0.000s 0:00.01 0.0%    0+0k 0+5io 0pf+0w
: time /sbin/mount -u -r -f -o sync,noatime /
0.000u 0.018s 0:45.66 0.0%    0+0k 0+398io 0pf+0w
```

#5 - 09/29/2012 08:10 AM - Bipin Chandra

i have a 2GB SLC based CF card and i installed the 1GB nanobsd, when using without much changes its fast but as soon as u start configuring more stuff and i also installed freeradius2 and after that it went dead slow, more than 45seconds for each edit, i also noticed the more the disk space is used, the more slower it gets

#6 - 10/12/2012 09:05 AM - David Jarosch

I've updated to 2.1-BETA0 (i386) built on Thu Oct 11 yesterday and having the same issues like described by you all. The GUI is very fast but when I have to do some changes or for example creating a new OpenVPN account, I have to wait like 1-2 minutes until one step finishes. I'm running on nanobsd (4g). I don't believe it's the CF card since it is, I hope I'm not wrong, a Transcend Ultimate 600x CF card.

Cheers,
Szop

#7 - 12/10/2012 02:51 PM - Jim Pingle

Two more points of data:

- Enabling write caching has no effect on speed.
- Switching from UFS to UFS2 has no effect on speed.

#8 - 12/11/2012 04:30 PM - Ermal Luçi

- Status changed from *New* to *Feedback*

- % Done changed from 0 to 100

Applied in changeset pfsense-tools:commit:008742971cb44d8c0f81929504ab7330442c4ba4.

#9 - 12/12/2012 10:56 PM - Phillip Davis

Fantastic, I have upgraded 7 Alix nanobsd systems to 2.1-BETA1. They all do conf_mount_ro in under 1 second now. Well done to all those who spent time on this. I can give all my friends a pfSense 2.1-BETA1 for Christmas:)

#10 - 12/12/2012 11:36 PM - Chris Buechler

- Status changed from *Feedback* to *Resolved*