

## FreeNAS - Bug #42803

### Unblock speculative prefetcher on pool creation

08/23/2018 07:13 PM - Alexander Motin

<b>Status:</b> Done	
<b>Priority:</b> No priority	
<b>Assignee:</b> Alexander Motin	
<b>Category:</b> OS	
<b>Target version:</b> 11.2-BETA3	
<b>Seen in:</b> 11.2-BETA2	<b>Needs Merging:</b> No
<b>Severity:</b> Medium	<b>Needs Automation:</b> No
<b>Reason for Closing:</b>	<b>Support Suite Ticket:</b> n/a
<b>Reason for Blocked:</b>	<b>Hardware Configuration:</b>
<b>Needs QA:</b> No	<b>ChangeLog Required:</b> No
<b>Needs Doc:</b> No	
<b>Description</b> Fix at r331950 appeared to be incomplete, fixing only case of pool import, but not pool creation, leaving prefetcher still blocked for newly created pools. It dramatically reduces sequential read speed between pool creation and first reboot or reimport. It can be diagnosed by not growing numbers in `sysctl kstat.zfs.misc.zfetchstats` during sequential read.	

#### Associated revisions

##### Revision fb319eab - 08/23/2018 07:17 PM - Alexander Motin

Unblock speculative prefetcher also on pool creation.

Fix at r331950 appeared to be incomplete, fixing only case of pool import, but not pool creation, leaving prefetcher still blocked for newly created pools.

Ticket: #42803

(cherry picked from commit 6dafcae1dddb8858d56d3337f4ae49071dff752a)

#### History

##### #1 - 08/23/2018 07:19 PM - Alexander Motin

- Status changed from Unscreened to In Progress

PR: <https://github.com/freenas/os/pull/142>

##### #2 - 08/24/2018 03:59 AM - Dru Lavigne

- Target version changed from 11.2-RC1 to 11.2-BETA3

##### #3 - 08/24/2018 05:34 AM - Dru Lavigne

- Needs Doc changed from Yes to No

- Needs Merging changed from No to Yes

#### #4 - 08/27/2018 07:13 AM - Alexander Motin

- Status changed from In Progress to Ready for Testing

#### #5 - 08/27/2018 07:16 AM - Dru Lavigne

- Subject changed from Speculative prefetcher blocked on pool creation to Unblock speculative prefetcher on pool creation

- Needs Merging changed from Yes to No

#### #7 - 08/30/2018 09:50 AM - Michael Reynolds

Grabbing a CentOS ISO as it is a 4.16G ISO

##### **kstat.zfs.misc.zfetchstats values before starting the wget**

```
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 18586419
kstat.zfs.misc.zfetchstats.misses: 18599666
kstat.zfs.misc.zfetchstats.hits: 4974
```

wget http://mirror.fileplanet.com/centos/7/isos/x86\_64/CentOS-7-x86\_64-DVD-1804.iso

##### **Values after starting the wget**

```
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 18591686
kstat.zfs.misc.zfetchstats.misses: 18604975
kstat.zfs.misc.zfetchstats.hits: 5906
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 18797352
kstat.zfs.misc.zfetchstats.misses: 18810746
kstat.zfs.misc.zfetchstats.hits: 14960
```

##### **A random time during the wget**

```
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 18664013
kstat.zfs.misc.zfetchstats.misses: 18677391
kstat.zfs.misc.zfetchstats.hits: 12631
```

##### **After wget finished the download**

```
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 19210404
kstat.zfs.misc.zfetchstats.misses: 19224005
kstat.zfs.misc.zfetchstats.hits: 39209
```

##### **Started a cp of the ISO from 1 pool to another after the wget finished**

```
cp CentOS-7-x86_64-DVD-1804.iso ../poolio
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 19612637
kstat.zfs.misc.zfetchstats.misses: 19626509
kstat.zfs.misc.zfetchstats.hits: 41607
```

##### **When the cp of the ISO finished**

```
root@freenas[./mnt/pool1]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 19643510
kstat.zfs.misc.zfetchstats.misses: 19657394
kstat.zfs.misc.zfetchstats.hits: 43590
```

Change between start and finish.

**kstat.zfs.misc.zfetchstats.max\_streams** Difference between start and finish: **1057091**

**kstat.zfs.misc.zfetchstats.misses** Difference between start and finish: **1057728**

**kstat.zfs.misc.zfetchstats.hits** Difference between start and finish: **38616**

**#9 - 08/30/2018 11:46 AM - Michael Reynolds**

Using dd to test this change more specifically.  
Since I already fetched the CentOS ISO may as well use it.

**before dd**

```
root@freenas[~]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 30477903
kstat.zfs.misc.zfetchstats.misses: 30499898
kstat.zfs.misc.zfetchstats.hits: 57077
```

**dd if=CentOS-7-x86\_64-DVD-1804.iso of=/dev/null bs=128k**

**Right after the dd**

```
root@freenas[~]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 30800414
kstat.zfs.misc.zfetchstats.misses: 30822993
kstat.zfs.misc.zfetchstats.hits: 93811
```

**#11 - 09/07/2018 02:36 PM - Michael Reynolds**

- Status changed from Ready for Testing to Passed Testing  
- Needs QA changed from Yes to No

Turned off compression

```
root@freenas[~]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 117022
kstat.zfs.misc.zfetchstats.misses: 117599
kstat.zfs.misc.zfetchstats.hits: 338
```

ran dd if=/dev/zero of=BUG-42803 bs=1024k count=32768

```
root@freenas[~]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 124182
kstat.zfs.misc.zfetchstats.misses: 124794
kstat.zfs.misc.zfetchstats.hits: 394
```

```
root@freenas[~]# dd if=BUG-42803 of=/dev/null bs=1024k count=32768
5363+0 records in
5363+0 records out
5623513088 bytes transferred in 0.683319 secs (8229698419 bytes/sec)
root@freenas[~]# sysctl kstat.zfs.misc.zfetchstats
kstat.zfs.misc.zfetchstats.max_streams: 151175
kstat.zfs.misc.zfetchstats.misses: 151933
kstat.zfs.misc.zfetchstats.hits: 5819
```

max\_streams went up by 26381  
misses went up by 27139  
hits went up by 5425

Checked with mav to see if this is acceptable.

He said the number of hits looks right, the number of misses seems high, but that's a question for a different ticket.

#12 - 09/08/2018 05:50 AM - Dru Lavigne

- Status changed from *Passed Testing* to *Done*