

FreeNAS - Bug #41910

Fix system crash/freeze when deleting many files

08/16/2018 09:28 AM - Michael Johnson

Status:	Done	
Priority:	No priority	
Assignee:	Alexander Motin	
Category:	OS	
Target version:	11.2-BETA3	
Seen in:	11.2-BETA2	Needs Merging: No
Severity:	Med High	Needs Automation: No
Reason for Closing:		Support Suite Ticket: n/a
Reason for Blocked:		Hardware Configuration:
Needs QA:	No	ChangeLog Required: No
Needs Doc:	No	

Description

I have a plex database backup with 550,000 files and 255,000 directories. If I try and delete the top directory of the database backup, the system freezes and never becomes responsive until a hard reset.

I am able to repeat the crash on different pools and even multiple systems (i7-3770 and xeon 1231v3). FWIW, I was unable to reproduce it on a Ryzen 1700 in an esxi VM.

Booting up 11.1, I am able to successfully remove the files.

Related issues:

Related to FreeNAS - Bug #40820: 11.2 Beta2 - system lock up when deleting files	Closed
Related to FreeNAS - Bug #41292: System crashes when deleting files	Closed
Related to FreeNAS - Bug #38923: #27514 fix causes panic on dataset quota ove...	Closed
Related to FreeNAS - Bug #42299: Rm file will cause freenas crash restart	Closed
Related to FreeNAS - Bug #40600: FreeNAS crashes when modifying files shared ...	Closed
Has duplicate FreeNAS - Bug #42605: watchdog timeout / panic / reboot when de...	Closed
Has duplicate FreeNAS - Bug #42972: Deleting large sparsebundle crashes freenas	Closed
Has duplicate FreeNAS - Bug #44914: Likely panic deleting specific ZFS path	Closed
Has duplicate FreeNAS - Bug #44383: When i open services it says (red exclama...	Closed

Associated revisions

Revision 14c9fd68 - 08/22/2018 10:14 AM - Alexander Motin

Create separate taskqueue to call `zfs_unlinked_drain()`.

r334810 introduced `zfs_unlinked_drain()` dispatch to taskqueue on every deletion of a file with extended attributes. Using `system_taskq` for that with its multiple threads in case of multiple files deletion caused all available CPU threads to uselessly spin on busy locks, completely blocking the system.

Use of single dedicated taskqueue is the only easy solution I've found, while it would be great if we could specify that some task should be executed only once at a time, but never in parallel, while many tasks could use different threads same time.

Sponsored by: iXsystems, Inc.
Ticket: #41910

(cherry picked from commit c1c7ce6e54560a6cfd7340d6bd6d85a5dc47798e)

Revision 5592caf9 - 08/22/2018 10:21 AM - Alexander Motin

Create separate taskqueue to call `zfs_unlinked_drain()`. (#141)

r334810 introduced `zfs_unlinked_drain()` dispatch to taskqueue on every deletion of a file with extended attributes. Using `system_taskq` for that with its multiple threads in case of multiple files deletion caused all available CPU threads to uselessly spin on busy locks, completely blocking the system.

Use of single dedicated taskqueue is the only easy solution I've found, while it would be great if we could specify that some task should be executed only once at a time, but never in parallel, while many tasks could use different threads same time.

Sponsored by: iXsystems, Inc.
Ticket: #41910

(cherry picked from commit `c1c7ce6e54560a6cfd7340d6bd6d85a5dc47798e`)

History

#1 - 08/16/2018 09:37 AM - Michael Johnson

Other things I tried:

- Removed all snapshots before attempting to remove the directory.
- Booted in single user and tried to remove

#2 - 08/16/2018 11:01 AM - Dru Lavigne

- *Private changed from No to Yes*
- *Reason for Blocked set to Need additional information from Author*

Michael: please attach a debug (System -> Advanced -> Save debug) to this ticket.

#3 - 08/16/2018 01:42 PM - Michael Johnson

- *File debug.tgz added*

Here is the debug logs.

#4 - 08/17/2018 01:22 AM - Dru Lavigne

- *Assignee changed from Release Council to Alexander Motin*

#5 - 08/17/2018 12:20 PM - Alexander Motin

- *Status changed from Unscreened to In Progress*
- *Target version changed from Backlog to 11.2-RC1*
- *Severity changed from New to Med High*
- *Reason for Blocked deleted (Need additional information from Author)*

Among the few reports about hangs on file deletion this is the first one that looks promising, thanks to one kernel dump, triggered by software watchdog, I found in the provided debug. As I see, system may not completely hang, but being very busy (inefficiently). As I see, at the moment of the panic out of 8 CPU cores 7 were waiting on lock

```
lock_delay() at lock_delay+0x42/frame 0xfffffe0a91dfe6b0
_sx_xlock_hard() at _sx_xlock_hard+0x178/frame 0xfffffe0a91dfe760
zfs_zget() at zfs_zget+0x1f8/frame 0xfffffe0a91dfe810
zfs_unlinked_drain() at zfs_unlinked_drain+0x99/frame 0xfffffe0a91dfe9e0
taskqueue_run_locked() at taskqueue_run_locked+0x154/frame 0xfffffe0a91dfea40
taskqueue_thread_loop() at taskqueue_thread_loop+0x98/frame 0xfffffe0a91dfea70
fork_exit() at fork_exit+0x83/frame 0xfffffe0a91dfeab0
fork_trampoline() at fork_trampoline+0xe/frame 0xfffffe0a91dfeab0
```

, while one remaining could be still working, or at least looked different:

```
arc_buf_access() at arc_buf_access+0xfa/frame 0xfffffe0a91e1c620
dbuf_hold_impl() at dbuf_hold_impl+0x79/frame 0xfffffe0a91e1c680
dbuf_hold() at dbuf_hold+0x25/frame 0xfffffe0a91e1c6b0
dnode_hold_impl() at dnode_hold_impl+0x130/frame 0xfffffe0a91e1c730
dmu_bonus_hold() at dmu_bonus_hold+0x1d/frame 0xfffffe0a91e1c760
zfs_zget() at zfs_zget+0xb4/frame 0xfffffe0a91e1c810
zfs_unlinked_drain() at zfs_unlinked_drain+0x99/frame 0xfffffe0a91e1c9e0
taskqueue_run_locked() at taskqueue_run_locked+0x154/frame 0xfffffe0a91e1ca40
taskqueue_thread_loop() at taskqueue_thread_loop+0x98/frame 0xfffffe0a91e1ca70
fork_exit() at fork_exit+0x83/frame 0xfffffe0a91e1cab0
fork_trampoline() at fork_trampoline+0xe/frame 0xfffffe0a91e1cab0
```

. Behind all that "activity" system had no time to react on anything else.

#6 - 08/17/2018 12:27 PM - Alexander Motin

- Related to Bug #40820: 11.2 Beta2 - system lock up when deleting files added

#7 - 08/17/2018 12:27 PM - Alexander Motin

- Related to Bug #41292: System crashes when deleting files added

#8 - 08/20/2018 07:10 AM - Alexander Motin

- Related to Bug #38923: #27514 fix causes panic on dataset quota overflow added

#9 - 08/20/2018 02:34 PM - Alexander Motin

- Related to Bug #42299: Rm file will cause freenas crash restart added

#10 - 08/22/2018 10:20 AM - Alexander Motin

- Needs Doc changed from Yes to No

- Needs Merging changed from Yes to No

PR: <https://github.com/freenas/os/pull/141>

Test case: in empty directory create many files, add some extended attributes to each and delete all of them:

```
jot 10000 | xargs -n 1 touch
setextattr system zzz 1 *
setextattr system zzy 1 *
setextattr system zzx 1 *
rm *
```

System without the patch may either hang or at least get very busy for longer then needed. System with the patch should handle it using only one CPU core and for much smaller time.

#11 - 08/22/2018 10:21 AM - Alexander Motin

- Status changed from In Progress to Ready for Testing

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

#12 - 08/22/2018 12:48 PM - Dru Lavigne

- Subject changed from System crash (freeze) when deleting hundreds of thousands of files. to Fix system crash/freeze when deleting many files

#13 - 08/22/2018 12:49 PM - Dru Lavigne

- File deleted (debug.tgz)

#14 - 08/22/2018 12:49 PM - Dru Lavigne

- Private changed from Yes to No

#15 - 08/24/2018 08:16 AM - Dru Lavigne

- Has duplicate Bug #42605: watchdog timeout / panic / reboot when deleting files/folders over SMB added

#16 - 08/24/2018 06:12 PM - Liviu Sas

Not sure if it's helpful or not, I have the same issue, but it's only happening on encrypted pools.
I can't reproduce the issue on non-encrypted pools.

#17 - 08/26/2018 04:58 PM - Dru Lavigne

- Has duplicate Bug #42972: Deleting large sparsebundle crashes freenas added

#18 - 08/26/2018 10:08 PM - Disk Didler

The quicker we can roll out Beta 3 (with this?) the better.
I kinda don't like watching my system fall over.

I thought this wasn't an issue for me, alas, I was wrong.

Not a good one!

#19 - 08/27/2018 08:45 AM - Dru Lavigne

- Related to Bug #40600: FreeNAS crashes when modifying files shared with a jail added

#21 - 09/06/2018 08:21 AM - Dru Lavigne

- Has duplicate Bug #44914: Likely panic deleting specific ZFS path added

#22 - 09/06/2018 09:59 AM - Dru Lavigne

- Has duplicate Bug #44383: When i open services it says (red exclamation) Sorry an error has occurred added

#23 - 09/06/2018 11:11 AM - Bonnie Follweiler

- Status changed from Ready for Testing to Passed Testing

- Needs QA changed from Yes to No

Test Passed in FreeNAS-11.2-MASTER-201809060900

#24 - 09/06/2018 11:31 AM - Dru Lavigne

- Status changed from Passed Testing to Done