

FreeNAS - Bug #27422

FreeNAS 11.1 possible memory leak

12/26/2017 01:33 AM - Lukasz Cepowski

Status:	Closed: Duplicate	
Priority:	Critical	
Assignee:	Alexander Motin	
Category:	OS	
Target version:	N/A	
Seen in:	11.1	Needs Merging: Yes
Severity:	New	Needs Automation: No
Reason for Closing:		Support Suite Ticket: n/a
Reason for Blocked:		Hardware Configuration: Hardware: - Intel S3420GP mobo - Xeon X3420 - 16GB DDR3 ECC - 2x onboard Intel NICs used for management only - 4x Intel 82575GB 1GBe, bonded in LACP, no jumbo currently - onboard 6x SATA controller - 2x Marvell 88SE9235 controllers, two additional separate 4x SATA controllers, on dedicated to two HDD, one to SSD
Needs QA:	Yes	ChangeLog Required: No
Needs Doc:	Yes	

Description

Hi All,

first of all, I would like to thank iXsystems guys and FreeBSD team for amazing piece of software FreeNAS is. I've been using it for more than 5 years and always been happy with stability and feature set FreeNAS offers.

Recently, decided to restore my server lab used for development and devops muscle training and rebuilt the primary storage server that exposes NFS to four VM hypervisors. Details are listed below, to keep it short, it is a 8x HDD Stripe-Mirror with single SSD divided into ZIL and L2ARC and Swap, Xeon, 16GB ECC, networking over bonded 4x 1Gbe using LACP.

Soon after setting it up, I've noticed the system hangs unexpectedly without printing any output neither in console nor in any logs, here's what happened:

1. Initial setup, 8x HDD stripe mirror + Intel 520 SSD, 8GB ZIL, 196GB L2ARC, no swap, LACP + jumbo frames, default tunables - system froze in couple of hours, initially responsive to ping, but impossible to log into SSH, seemed like stuck on IO. Before it happened, noticed relatively high CPU usage on System and IRQ.
2. Thought it might be SSD, so replaced it with Intel 320 SSD, left 8GB, 128GB L2ARC, no swap, LACP + jumbo frames, default tunables - again system froze without crashing in couple of hours.
3. Suspected Jumbo Frames to be the culprit, so reverted back to 1500 MTU - system froze again but after a longer time.
4. Had an L2ARC hunch, set it in tunables to 12GB and increased l2arc_write_max and l2arc_write_boost to 128MB - system froze again but this time worked even longer, more than 24 hours, although I wasn't doing any load tests in meanwhile.
5. Finally noticed there is no swap so I added 16GB partition on the SSD and swapon'ed it - system froze but again after more than a day, noticed FreeNAS started swapping soon before crashing.
6. Turned on autotune, got 13GB for L2ARC and started doing load tests - froze again in around 12 hours or so, noticed swapping again soon before system froze
7. Bless you guys for the "Remote Graphite Server Hostname" option to upload carbon stats, done that, decreased L2ARC to 10GB, started load test to see what's happening - froze again and rebooted itself around an hour later early morning.
8. Retried load tests and watched Grafana - same result, froze but this time got a pattern in Grafana dashboard.

So, it seems, FreeNAS, kernel or some process, is running out of memory, regardless of the L2ARC size, at some point it starts to

swap heavily and soon after that, poof... freezed.

Snapshots of the Grafana dashboards:

- 3 hours before the freeze: <https://snapshot.raintank.io/dashboard/snapshot/9xM3Z6ropc6eVBZIUqtC0737Yvw3Y9X0?orgId=2>
- 24 hours before the freeze: <https://snapshot.raintank.io/dashboard/snapshot/zlYgJkmanglGNhKV6XbeoqNrwW9GhuUp?orgId=2>

Have a look on the ZFS Hit Ratio, Memory, Swap, and SSD IOPS.

Attached: lspci, ps aux, sysctl -a, vmstat, zfs-stat, dmesg taken soon before freeze.

// Copied from the original post: <https://forums.freenas.org/index.php?threads/freenas-11-1-possible-memory-leak.60109/>

Related issues:

Related to FreeNAS - Bug #27270: Fix memory leak

Resolved

12/15/2017

History

#1 - 12/26/2017 06:22 AM - Kris Moore

- Assignee changed from Release Council to William Grzybowski
- Priority changed from No priority to Critical
- Target version set to 11.1-U1

#2 - 12/26/2017 07:25 AM - William Grzybowski

- Assignee changed from William Grzybowski to Alexander Motin

I dont see anything on userland consuming too much memory.

Any ideas?

#3 - 12/26/2017 09:17 AM - Lukasz Cepowski

FYI, freezed again after around 8 hours. I've double checked are there any cron jobs or tasks running every couple of hours and had only a smartctl test running every 8 hours, but changed that to 1 hour interval and hasn't changed the behaviour of the server.

This is snapshot from last 48 hours, again same pattern with swapping.

<https://snapshot.raintank.io/dashboard/snapshot/OSBozFc3csRBpBXLWJNJTyLXN5WWYITN?orgId=2>

#4 - 12/26/2017 09:34 AM - Alexander Motin

- Related to Bug #27270: Fix memory leak added

#5 - 12/26/2017 09:38 AM - Alexander Motin

- Status changed from Unscreened to Closed: Duplicate
- Seen in changed from TrueNAS 11.1-U1 to 11.1

To say for sure I'd need `vmstat -z` output, but I suspect this ticket is a duplicate of [#27270](#). At least it does look like a memory leak. Lukasz, can you show that output, especially the line for g_bio, if it happen again?

#6 - 12/26/2017 09:54 AM - Lukasz Cepowski

Hi Alexander,

sure, I'll dump output from `vmstat -z` captured right before the freeze, according to the pattern observed it should happen in next 8 or so hours.

BTW, [#27270](#) seems to be fixed, is it possible to get either a patch or nightly .iso build with fixed version? Or should I downgrade to some other major release (10.x, 9.10.x)?

#7 - 12/26/2017 10:40 AM - Alexander Motin

The fix (if it is the same problem) is applied to nightly and will be part of 11.1-U1. Though nightly already include new set of changes from the upstream, so may introduce any new problems meanwhile. The most safe way for you would be revert to the latest 11.0-Ux and wait there until 11.1-U1 is released.

#8 - 12/26/2017 11:55 AM - Lee Marzke

This bug may have triggered for me, also in 11-1-Release, on system with 32GB ECC RAM and it ran fine for 8 days then just died. Running mostly ESXi VM's over NFS. I saw some logs referencing killing processes, no swap available.

I didn't have time to debug much as I'm running production loads, reverted back to 11.0-U4

#9 - 12/26/2017 10:11 PM - Lukasz Cepowski

- File *vmstat-1514349900.out* added

System locked up again, as expected, I've dumped vmstat -z every minute, this is the last one before the hang up:

```
g_bio:                376,          0,35194941,          239,314447189,          0,          0
```

#10 - 12/27/2017 05:56 AM - Alexander Motin

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

#11 - 12/31/2017 03:25 PM - Dru Lavigne

- Target version changed from 11.1-U1 to N/A

#12 - 01/01/2018 01:02 PM - Tobias Müllauer

think i have this bug to. my swap is full and all my 24g memory.

#13 - 01/01/2018 02:00 PM - Lee Marzke

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

Which bug is that ?

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

What is the bug number please? I can't find any open bugs against memory leaks and I wan't to be sure it is really fixed in 11.1-U1 before I try this branch again.

#14 - 01/01/2018 02:23 PM - Tobias Müllauer

Lee Marzke wrote:

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

Which bug is that ?

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

What is the bug number please? I can't find any open bugs against memory leaks and I wan't to be sure it is really fixed in 11.1-U1 before I try this branch again.

This one?

Relaterar till FreeNAS - Bug [#27270](#): ARC and g_bio zone memory leak

#15 - 01/01/2018 07:51 PM - Lee Marzke

Tobias Müllauer wrote:

Lee Marzke wrote:

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

Which bug is that ?

Alexander Motin wrote:

Yes, that is a leak we saw and that should be fixed in 11.1-U1.

What is the bug number please? I can't find any open bugs against memory leaks and I wan't to be sure it is really fixed in 11.1-U1 before I try this branch again.

This one?

Relaterar till FreeNAS - Bug [#27270](#): ARC and g_bio zone memory leak

Thank you for the help Tobias, and thank you Alexander for getting this patched quickly.
-Lee

#16 - 02/28/2018 04:18 PM - Dru Lavigne

- File deleted (*lspci.txt*)

#17 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*dmesg.txt*)

#18 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*psaux.txt*)

#19 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*vmstat.txt*)

#20 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*zfs-stat.txt*)

#21 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*sysctl.txt*)

#22 - 02/28/2018 04:19 PM - Dru Lavigne

- File deleted (*vmstat-1514349900.out*)