

FreeNAS - Bug #28539

Fix renamed function in zilstat MIB

02/15/2018 02:14 PM - Bill O'Hanlon

Status: Done	
Priority: No priority	
Assignee: Alexander Motin	
Category: OS	
Target version: 11.2-BETA1	
Seen in: Unspecified	Needs Merging: No
Severity: Medium	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 zilstat is failing because fbt::zil_lwb_write_start:entry no longer exists as an trace point for dtrace. /usr/local/bin/snmp-agent.py calls zilstat, so the FREENAS-MIB is returning 0 for a handful of variables.	
Related issues: Related to FreeNAS - Bug #40768: Revert recent zilstat commit as there is no ... Done	

Associated revisions

Revision 248de883 - 02/23/2018 06:39 AM - Bill O'Hanlon

fix(snmp/zilstat) dtrace was triggering off zil_lwb_write_start, which has been removed.

Ticket: #28539

Revision 3bda775e - 04/19/2018 03:47 PM - Bill O'Hanlon

fix(snmp/zilstat) dtrace was triggering off zil_lwb_write_start, which has been removed.

Ticket: #28539

History

#1 - 02/16/2018 11:43 AM - Dru Lavigne

- Assignee changed from Release Council to Alexander Motin

#2 - 02/16/2018 01:50 PM - Alexander Motin

Are you sure you saw the problem on 11.1-U2, not 11.2 (11-Nightly/11-stable)? Because that function was renamed only recently and that change is not in 11.1-stable branch.

#3 - 02/16/2018 03:09 PM - Bill O'Hanlon

- Seen in changed from TrueNAS 11.1-U2 to Unspecified

Actually, I'm not sure.

#4 - 02/22/2018 07:46 AM - Alexander Motin

- Assignee changed from Alexander Motin to Benno Rice

Benno, take a look at this please. It may be a trivial function rename in new ZFS code, though the ZIL code logic has changed also, so closer look may be needed.

#5 - 02/22/2018 10:51 AM - Benno Rice

zil_lwb_write_start was renamed to zil_lwb_write_issue in r324011. I'll do some further digging to see whether the semantics still match or not.

#6 - 02/22/2018 10:58 AM - Benno Rice

Looking at zilstat my read is that it's trying to measure (possibly among other things) the number of bytes being pushed through the ZIL, which is why it's hooking zil_lwb_write_start. If my read on this is correct I think the right fix is just to hook zil_lwb_write_issue instead.

#7 - 02/22/2018 11:09 AM - Benno Rice

- Status changed from Not Started to In Progress

#8 - 02/22/2018 12:01 PM - Bill O'Hanlon

FWIW, I did that on my system while trying to get past this issue to test another one, and it seemed to work well. So I agree that it seems like a good approach. :-)

#9 - 02/22/2018 02:25 PM - Benno Rice

Bill, did you want to take this over from here or did you want me to work up the patch for zilstat?

#10 - 02/23/2018 06:32 AM - Bill O'Hanlon

Sure, I can do that.

#11 - 02/23/2018 06:36 AM - Dru Lavigne

- Assignee changed from Benno Rice to Bill O'Hanlon

#12 - 02/26/2018 08:18 AM - Bill O'Hanlon

- Status changed from In Progress to Done

#13 - 02/26/2018 08:25 AM - Dru Lavigne

- Subject changed from zilstat is broken, which breaks FREENAS-MIB in SNMP to Fix renamed function in zilstat MIB

- Needs Doc changed from Yes to No

#14 - 02/26/2018 08:29 AM - Dru Lavigne

- Target version changed from 11.2-RC2 to 11.2-BETA1

- Needs Merging changed from Yes to No

#15 - 05/02/2018 05:41 AM - Dru Lavigne

- Status changed from Done to Ready for Testing

#16 - 06/21/2018 01:39 PM - Nick Wolff

- Status changed from Ready for Testing to Failed Testing

Zilstat is puking. Not sure if it's related to this or new issue

```
root@fncertified:/var/log # zilstat
```

```

dtrace: invalid probe specifier
#pragma D option quiet
inline int OPT_time = 0;
inline int OPT_txg = 0;
inline int OPT_pool = 0;
inline int OPT_mega = 0;
inline int INTERVAL = 1;
inline int LINES = -1;
inline int COUNTER = -1;
inline int FILTER = 0;
inline string POOL = "";
dtrace:::BEGIN
{
    /* starting values */
    MEGA = 1000000;
    counts = COUNTER;
    secs = INTERVAL;
    interval = INTERVAL;
    interval == 0 ? interval++ : 1;
    line = 0;
    last_event[""] = 0;
    nused=0;
    nused_max_per_sec=0;
    nused_per_sec=0;
    size=0;
    size_max_per_sec=0;
    size_per_sec=0;
    syncops=0;
    size_4k=0;
    size_4k_32k=0;
    size_32k=0;
    OPT_txg ? printf("waiting for txg commit...\n") : 1;
}

/*
 * collect info when zil_lwb_write_start fires
 */
fbt::zil_lwb_write_issue:entry
/OPT_pool == 0 || POOL == args[0]->zl_dmu_pool->dp_spa->spa_name/
{
    nused += args[1]->lwb_nused;
    nused_per_sec += args[1]->lwb_nused;
    size += args[1]->lwb_sz;
    size_per_sec += args[1]->lwb_sz;
    syncops++;
    args[1]->lwb_sz <= 4096 ? size_4k++ : 1;
    args[1]->lwb_sz > 4096 && args[1]->lwb_sz < 32768 ? size_4k_32k++ : 1;
    args[1]->lwb_sz >= 32768 ? size_32k++ : 1;
}

/*
 * Timer
 */
profile:::tick-1sec
{
    OPT_txg ? secs++ : secs--;
    nused_per_sec > nused_max_per_sec ? nused_max_per_sec = nused_per_sec : 1;
    nused_per_sec = 0;
    size_per_sec > size_max_per_sec ? size_max_per_sec = size_per_sec : 1;
    size_per_sec = 0;
}

/*
 * Print header
 */
profile:::tick-1sec
/OPT_txg == 0 && line == 0/
{
    /* print optional headers */
    OPT_time ? printf("%-20s ", "TIME") : 1;

    /* print header */
    OPT_mega ? printf("%10s %10s %10s %10s %10s %10s",
        "N-MB", "N-MB/s", "N-Max-Rate",
        "B-MB", "B-MB/s", "B-Max-Rate") :

```

```

        printf("%10s %10s %10s %10s %10s %10s",
            "N-Bytes", "N-Bytes/s", "N-Max-Rate",
            "B-Bytes", "B-Bytes/s", "B-Max-Rate");
    printf(" %6s %6s %6s %6s\n",
        "ops", "<=4kB", "4-32kB", ">=32kB");
    line = LINES;
}

fbt::txg_quiesce:entry
/OPT_txg == 1 && POOL == args[0]->dp_spa->spa_name && line == 0/
{
    OPT_time ? printf("%-20s ", "TIME") : 1;

    OPT_mega ? printf("%10s %10s %10s %10s %10s %10s %10s",
        "txg", "N-MB", "N-MB/s", "N-Max-Rate",
        "B-MB", "B-MB/s", "B-Max-Rate") :
        printf("%10s %10s %10s %10s %10s %10s %10s",
            "txg", "N-Bytes", "N-Bytes/s", "N-Max-Rate",
            "B-Bytes", "B-Bytes/s", "B-Max-Rate");
    printf(" %6s %6s %6s %6s\n",
        "ops", "<=4kB", "4-32kB", ">=32kB");
    line = LINES;
}

/*
 * Print Output
 */
profile::tick-1sec
/OPT_txg == 0 && secs == 0/
{
    OPT_time ? printf("%-20Y ", walltimestamp) : 1;
    OPT_mega ?
        printf("%10d %10d %10d %10d %10d %10d",
            nused/MEGA, nused/(interval*MEGA), nused_max_per_sec/MEGA,
            size/MEGA, size/(interval*MEGA), size_max_per_sec/MEGA) :
        printf("%10d %10d %10d %10d %10d %10d",
            nused, nused/interval, nused_max_per_sec,
            size, size/interval, size_max_per_sec);
    printf(" %6d %6d %6d %6d\n",
        syncops, size_4k, size_4k_32k, size_32k);
    nused = 0;
    nused_per_sec = 0;
    nused_max_per_sec = 0;
    size=0;
    size_max_per_sec=0;
    size_per_sec=0;
    syncops=0;
    size_4k=0;
    size_4k_32k=0;
    size_32k=0;
    secs = INTERVAL;
    counts--;
    line--;
}

fbt::txg_quiesce:entry
/OPT_txg == 1 && POOL == args[0]->dp_spa->spa_name/
{
    secs <= 0 ? secs=1 : 1;
    OPT_time ? printf("%-20Y ", walltimestamp) : 1;
    OPT_mega ?
        printf("%10d %10d %10d %10d %10d %10d %10d", args[1],
            nused/MEGA, nused/(secs*MEGA), nused_max_per_sec/MEGA,
            size/MEGA, size/(secs*MEGA), size_max_per_sec/MEGA) :
        printf("%10d %10d %10d %10d %10d %10d %10d", args[1],
            nused, nused/secs, nused_max_per_sec,
            size, size/secs, size_max_per_sec);
    printf(" %6d %6d %6d %6d\n",
        syncops, size_4k, size_4k_32k, size_32k);
    nused = 0;
    nused_per_sec = 0;
    nused_max_per_sec = 0;
    size=0;
    size_max_per_sec=0;
    size_per_sec=0;
}

```

```
    syncops=0;
    size_4k=0;
    size_4k_32k=0;
    size_32k=0;
    secs = 0;
    counts--;
    line--;
}

/*
 * End of program
 */
profile:::tick-1sec
/OPT_txd == 0 && counts == 0/
{
    exit(0);
}
fbt:::txg_quiesce:entry
/OPT_txd == 1 && counts == 0/
{
    exit(0);
}
: "/usr/lib/dtrace/ip.d", line 112: failed to copy typedef in6_addr_t source type: Type information is in parent and unavailable
root@fncertified:/var/log #
```

#18 - 06/25/2018 12:55 PM - Dru Lavigne

- Assignee changed from Bill O'Hanlon to Alexander Motin

#20 - 06/25/2018 02:28 PM - Alexander Motin

The errors at the end are probably caused by different issue ([#34609](#)). It should work (can be tested) on TrueNAS though.

#21 - 06/29/2018 08:57 AM - Dru Lavigne

- Status changed from Failed Testing to Ready for Testing

- Priority changed from Regression to No priority

#22 - 06/29/2018 11:06 AM - Nick Wolff

- Status changed from Ready for Testing to Done

No longer blocked by broken dtrace.

Looks good
test passed

#23 - 06/29/2018 12:35 PM - Dru Lavigne

- Needs QA changed from Yes to No

#24 - 08/09/2018 02:12 PM - Alexander Motin

- Related to Bug #40768: Revert recent zilstat commit as there is no zil_lwb_write_start() function in FreeNAS 11.1 added