

FreeNAS - Bug #40768

Revert recent zilstat commit as there is no zil_lwb_write_start() function in FreeNAS 11.1

08/04/2018 09:54 AM - Geoff Seeley

| | |
|----------------------------------|----------------------------------|
| Status: Done | |
| Priority: No priority | |
| Assignee: Alexander Motin | |
| Category: OS | |
| Target version: 11.1-U6 | |
| Seen in: 11.1-U5 | Needs Merging: No |
| Severity: New | 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

Running zilstat on 11.1-U5 results in the error:

```
probe description fbt::zil_lwb_write_issue:entry does not match any probes
```

zilstat appears to have an incorrect probe reference on line 173. It does not match the probe reference in the comment on line 171. I changed the reference to match the one in the comment and zilstat works.

Error output:

```
root@nas:/mnt/nas1 # /usr/local/bin/zilstat
dtrace: invalid probe specifier
#pragma D option quiet
inline int OPT_time = 0;
inline int OPT_txcg = 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_txdg ? 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_txdg ? 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_txdg == 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_txdg == 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",
        "txg", "N-MB", "N-MB/s", "N-Max-Rate",
        "B-MB", "B-MB/s", "B-Max-Rate") :
        printf("%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_txxg == 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_txxg == 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", 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", 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_txg == 0 && counts == 0/
{
    exit(0);
}
fbt::txg_quiesce:entry
/OPT_txg == 1 && counts == 0/
{
    exit(0);
}
: probe description fbt::zil_lwb_write_issue:entry does not match any probes
```

Patch:

```
root@nas:/mnt/nas1 # diff /usr/local/bin/zilstat zilstat-fixed
173c173
< fbt::zil_lwb_write_issue:entry
---
> fbt::zil_lwb_write_start:entry
```

Patched Output:

```
root@nas:/mnt/nas1 # ./zilstat-fixed
  N-Bytes  N-Bytes/s  N-Max-Rate    B-Bytes  B-Bytes/s  B-Max-Rate    ops  <=4kB  4-32kB  >=32kB
      0         0         0         0         0         0         0       0       0       0
      0         0         0         0         0         0         0       0       0       0
      0         0         0         0         0         0         0       0       0       0
      0         0         0         0         0         0         0       0       0       0
      0         0         0         0         0         0         0       0       0       0
^C
```

Related issues:

Related to FreeNAS - Bug #28539: Fix renamed function in zilstat MIB

Done

Associated revisions

Revision d40b72a2 - 08/09/2018 02:20 PM - Alexander Motin

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

This reverts commit 3bda775e16b8d276245d40741f926804f09afa6b.

There is no zil_lwb_write_start() function in FreeNAS 11.1 ZFS.
This commit should be limited to 11.2 only.

Ticket: #40768

History

#1 - 08/04/2018 09:56 AM - Geoff Seeley

- Description updated

#2 - 08/05/2018 01:59 PM - Dru Lavigne

- Category changed from Console to OS

- Assignee changed from Release Council to Alexander Motin

#4 - 08/09/2018 02:12 PM - Alexander Motin

- Related to Bug #28539: Fix renamed function in zilstat MIB added

#5 - 08/09/2018 02:18 PM - Alexander Motin

- Status changed from Unscreened to In Progress

The problem I see in fact that [#28539](#) commit was merged to 11.1-stable, where it does not belong, since mentioned function does not exist there, added only in 11.2.

#6 - 08/09/2018 02:22 PM - Dru Lavigne

- Target version changed from Backlog to 11.1-U6

#7 - 08/09/2018 04:26 PM - Dru Lavigne

- Subject changed from zilstat error: probe description fbt::zil_lwb_write_issue:entry does not match any probes to Revert recent zilstat commit as there is no zil_lwb_write_start() function in FreeNAS 11.1
- Needs Doc changed from Yes to No

#8 - 08/10/2018 04:56 AM - Dru Lavigne

11.1-stable PR: <https://github.com/freenas/freenas/pull/1668>

#9 - 08/10/2018 10:34 AM - Bug Clerk

- Status changed from In Progress to Ready for Testing

PR: <https://github.com/freenas/freenas/pull/1668>

#10 - 08/13/2018 08:34 AM - Dru Lavigne

- Needs Merging changed from Yes to No

#11 - 08/14/2018 07:04 AM - Bonnie Follweiler

- Status changed from Ready for Testing to Passed Testing

Test Passed in FreeNAS 11.1-U6 Internal3

#12 - 08/14/2018 07:11 AM - Bonnie Follweiler

- File zilstat.jpg added

#13 - 08/14/2018 07:15 AM - Bonnie Follweiler

- File zilstat_test.jpg added

#14 - 08/14/2018 07:36 AM - Dru Lavigne

- File deleted (zilstat.jpg)

#15 - 08/14/2018 08:01 AM - Dru Lavigne

- Status changed from Passed Testing to Done
- Needs QA changed from Yes to No

Files

| | | | |
|------------------|---------|------------|-------------------|
| zilstat_test.jpg | 1.05 MB | 08/14/2018 | Bonnie Follweiler |
|------------------|---------|------------|-------------------|