

Servers - Bug #1224

[repo] db-import's repo-add locks up sometimes

2017-03-08 12:34 AM - lukeshu

Status: fixed	% Done: 100%
Priority: bug	
Assignee: lukeshu	
Category:	
Description <p>So this is the second time this has happened.</p> <p>db-import-*, which import packages from Arch/ALARM, calls repo-add, which is a shell program owned by pacman.</p> <p>Sometimes, repo-add itself (that is: the /bin/bash binary) locks up, pegging one of the CPU cores at 100%.</p> <p>The first time, I announced it on IRC, and killed it after about 6 hours.</p> <p>It happened again (at 2017-02-26, 7:00AM EST (UTC-5)); and having 8 cores, the rest of which are pretty idle, I decided to wait, and see what would happen. It's been several days, and nothing has happened.</p> <p>There's a patch out for bash (4.4.011-2 -> 4.4.012-2), but that doesn't seem to be related.</p> <p>I installed gdb, and attached to the process, but without debugging symbols, I didn't have much luck.</p> <p>I'm re-compiling bash with debugging symbols, installing that. I'll kill the process, and wait for it to happen again.</p>	

History

#1 - 2017-03-08 06:17 AM - isacdaavid

debugging symbols!?

Only db-import-archlinuxarm-pkg is calling repo-add: to create pristine package databases (yes, every time db-import-archlinuxarm-pkg is called) and thus avoid the spoiled databases ALARM is serving (missing fields and mismatches between .files and .db databases)

I added that extra step so that certain consistency checks in Parabolaweb wouldn't fail and skip over populating file lists; which in turn was the cause of further errors. ALARM doesn't use something like Archweb, nor their website knows anything about a package's files, so they don't care/don't notice the issue. I knew the process was very slow and wasteful but didn't expect it to hang up completely (or even take more than it takes cron to call the script again). I can't recall exactly how long it took under my test environment, although the figures weren't *that* bad. Only a single-core VM, half a GiB of RAM and, admittedly, an SSD were used. Maybe my expectations were wrong.

It would be nice to try to record the full repo-add command line, as seen in the process tree, the next time feces hit the fan. This could be related to a specific package

#2 - 2017-03-20 06:42 PM - lukeshu

It's going on right now. I've attached to it with gdb. I've identified the problem, but not the root cause.

Bash's bgpids.storage is wrong somehow. It gets stuck the loop in bgp_delete (jobs.c:867). A comment describes bgpids.storage it as being circular, but that seems to be contradicted by the pid != NO_PIDSTAT check; so I'm not entirely sure what "correct" is.

#3 - 2017-03-20 07:05 PM - lukeshu

I've reported the bug to bash upstream, but it hasn't yet shown up in the web view.

#4 - 2017-03-21 03:46 PM - lukeshu

My report: <https://lists.gnu.org/archive/html/bug-bash/2017-03/msg00141.html>

Dup of: <http://lists.gnu.org/archive/html/bug-bash/2017-02/msg00025.html>

There are a couple of possible patches, but I'm not sure how I want to proceed.

#5 - 2018-04-30 02:27 AM - lukeshu

- % Done changed from 0 to 20
- Assignee set to lukeshu
- Status changed from open to in progress

I've just published ~lukeshu/bash (and bash-debug) that have the patch from <https://lists.gnu.org/archive/html/bug-bash/2017-03/msg00144.html> applied.

I've installed it on Winston. We'll see how things go.

#6 - 2018-07-16 09:38 PM - lukeshu

- % Done changed from 20 to 100
- Status changed from in progress to fixed

That patch became <https://ftp.gnu.org/gnu/bash/bash-4.4-patches/bash44-020> which is included in core/bash-4.4.023