

Ruby trunk - Bug #6653

1.9.2/1.9.3 exhibit SEGV with many threads+tcp connections

06/27/2012 08:19 AM - erikh (Erik Hollensbe)

Status: Closed	
Priority: Normal	
Assignee: akr (Akira Tanaka)	
Target version: 2.6	
ruby -v: ruby 1.9.2p290 (2011-07-09 revision 32553) [x86_64-linux]	Backport:
Description the script: https://gist.github.com/4f36f8543ad702861096 the trace + output of the run: https://gist.github.com/cf7dd137ad65802c46ae ruby -v is 1.9.2-p290, but we're seeing this in 1.9.3-p194 as well. This does <i>not</i> exhibit on OS X, only linux, we tested on Ubuntu 12.04. I can get more information if desired. Just guessing, this appears to be a bug in how FD_SETSIZE is handled. Thank you!	
Related issues: Is duplicate of Backport193 - Backport #8080: Segfault in rb_fd_set Closed 03/13/2013	

History

#1 - 06/27/2012 11:23 AM - normalperson (Eric Wong)

"erikh (Erik Hollensbe)" erikh@hollensbe.org wrote:

Issue [#6653](#) has been reported by erikh (Erik Hollensbe).

Bug [#6653](#): 1.9.2/1.9.3 exhibit SEGV with many threads+tcp connections
<https://bugs.ruby-lang.org/issues/6653>

Author: erikh (Erik Hollensbe)
Status: Open
Priority: Normal
Assignee:
Category:
Target version:
ruby -v: ruby 1.9.2p290 (2011-07-09 revision 32553) [x86_64-linux]

the script: <https://gist.github.com/4f36f8543ad702861096>
the trace + output of the run: <https://gist.github.com/cf7dd137ad65802c46ae>

Private gist for public bug reports makes no sense. Private gists requires account + ssh key on github to "git clone" from.

ruby -v is 1.9.2-p290, but we're seeing this in 1.9.3-p194 as well.

This does *not* exhibit on OS X, only linux, we tested on Ubuntu 12.04.

I can't reproduce this on a similar system (Debian testing (wheezy)) with 1.9.3-p194 nor Ruby 1.9.2-p290.

rb_fd_set() should not get called under 1.9.3 on Linux from rb_thread_fd_writable(), can you show a backtrace from 1.9.3?

Are you certain /opt/ruby/lib/libruby.so.1.9 got changed/upgraded to the 1.9.3 version?

The ruby/config.h header for 1.9.3 should have detected ppoll() and set: #define HAVE_PPOLL 1

ppoll() usage would prevent rb_fd_set() usage in your particular code path.

Also, what is the value of HAVE_RB_FD_INIT in ruby/config.h? (it should be 1 on Linux for all Ruby 1.9.x)

If you have build logs handy, can you see if ppoll() got detected on 1.9.3?

#2 - 07/14/2012 06:10 PM - kosaki (Motohiro KOSAKI)

- Status changed from Open to Feedback

#3 - 08/08/2012 10:17 PM - tommy.odom (Tommy Odom)

I've hit a similar issue while using Chef with Ruby 1.9.3 on Ubuntu 12.04 x86_64. I've tried with both the Ubuntu 1.9.3 packages as well as the packages provided by Brightbox (ruby 1.9.3p194 (2012-04-20 revision 35410) [x86_64-linux]) and with both I've hit a very similar stack trace. One thing I have noticed though is that this does not occur if the max open files is set to <= 1700.

You can see the stack trace at: <https://gist.github.com/3294941>

The code in Chef that is failing is: <https://github.com/opscode/mixlib-shellout/blob/master/lib/mixlib/shellout/unix.rb>

** Update **

I figured out that I had a piece of code that was opening a bunch of file handles (around 1700) using File.new and wasn't closing them. So it appears that in my case having 1700 open files was contributing to the issue.

#4 - 10/12/2012 10:14 PM - mame (Yusuke Endoh)

- Priority changed from Normal to 3

Please write a complete reproducing procedure. It requires memcached, right? I cannot repro on Ubuntu 12.04.

--

Yusuke Endoh mame@tsg.ne.jp

#5 - 11/05/2012 09:33 PM - mame (Yusuke Endoh)

Erik Hollensbe, ping?

--

Yusuke Endoh mame@tsg.ne.jp

#6 - 11/16/2012 04:43 AM - erikh (Erik Hollensbe)

Sorry for the abysmally late response -- I can't seem to get the redmine here to send me email for some reason.

Hi Folks, so I actually sorted this out with some help from others. It's not an issue of memcached, or rather, didn't appear to be when I looked into it.

If you adjust the limit (either with ulimit or the Process::tooling) it goes away. Conversely you *should* see this problem if you adjust the ulimit threshold below the amount of descriptors you're trying to work with.

I will also say that it has been a significant amount of time since I had this problem and have changed jobs since then, so I don't have access to specifics on build env, etc anymore.

The problem seems to be the handling of the case where the system says "I can't give you any more descriptors", not any specific value. I was using a lot of threads too, if that matters.

#7 - 11/25/2012 11:41 AM - mame (Yusuke Endoh)

- Status changed from Feedback to Assigned

- Assignee set to akr (Akira Tanaka)

- Target version set to 2.0.0

Erik, thank you for the reply!

Well, it seems that there is something wrong in the handling of file descriptors bigger than FD_SETSIZE.

Akr-san, kosaki-san, ko1, do you have any idea?

--

Yusuke Endoh mame@tsg.ne.jp

#8 - 11/25/2012 01:45 PM - kosaki (Motohiro KOSAKI)

Unfortunately, I've seen nothing wrong even if file descriptor limits are greater than FD_SETSIZE.

#9 - 02/18/2013 09:54 PM - mame (Yusuke Endoh)

- *Target version changed from 2.0.0 to 2.6*

#10 - 03/13/2013 04:43 PM - kosaki (Motohiro KOSAKI)

- *Status changed from Assigned to Closed*

closed. because it is duplicated.