

Ruby trunk - Feature #9330

[PATCH 0/3] avoid redundant fcntl/fstat syscalls for cloexec sockets

12/31/2013 11:22 AM - normalperson (Eric Wong)

Status:	Closed
Priority:	Normal
Assignee:	
Target version:	2.6
Description	
commit 52525b673669019dc3c03474e613937fd3587187 Author: Eric Wong e@80x24.org Date: Tue Dec 31 02:09:27 2013 +0000	
socket: avoid redundant fcntl calls for FD passing	
If MSG_CMSG_CLOEXEC works on the first recvmsg syscall, we should assume it always works and avoid checking it with every single FD received.	
While we're at it, it is pointless to set FD_CLOEXEC immediately before we close it in discard_cmsg; so only update the max FD (the FD is likely to be recycled, so we should update even if we close it).	
Note: I think it's safe to assume any system with MSG_CMSG_CLOEXEC will have implemented SOCK_CLOEXEC; too, so we reuse use rsock_detect_cloexec instead of defining a new function wrapped by #ifdef MSG_CMSG_CLOEXEC.	
commit aa9ed17a5f10cf53f72dc5ab1f4ff8d8b5e4dd84 Author: Eric Wong e@80x24.org Date: Tue Dec 31 01:36:27 2013 +0000	
socket: avoid redundant calls to rb_update_max_fd	
rsock_init_sock already calls rb_update_max_fd, so do not waste time making redundant fstat() and atomic CAS instructions.	
commit b1694d7bac55bd9f537426624fe599c46b10dee6 Author: Eric Wong e@80x24.org Date: Mon Dec 30 23:58:32 2013 +0000	
socket: avoid redundant fcntl(fd, F_GETFD) with SOCK_CLOEXEC	
If socket or socketpair succeeds with SOCK_CLOEXEC once, we should expect it to succeed for all subsequent invocations of each of those syscalls. This reduces the number of syscalls made and makes strace output cleaner.	
While we're at it, swap some remaining rb_fd_fix_cloexec calls for rb_maygvl_fd_fix_cloexec, as the former calls rb_update_max_fd redundantly. A future commit will reduce rb_update_max_fd calls further.	
The following changes since commit d6c7ab9ffeff31fb5a49776ab74a61b478344107:	
<ul style="list-style-type: none">• 2013-12-31 (2013-12-30 19:48:05 +0000)	
are available in the git repository at:	
git://bogomips.org/ruby.git socket_cloexec_once	
for you to fetch changes up to 52525b673669019dc3c03474e613937fd3587187:	

socket: avoid redundant fcntl calls for FD passing (2013-12-31 02:17:53 +0000)

Eric Wong (3):

socket: avoid redundant fcntl(fd, F_GETFD) with SOCK_CLOEXEC

socket: avoid redundant calls to rb_update_max_fd

socket: avoid redundant fcntl calls for FD passing

```
ext/socket/ancdata.c | 9 ++++++--
ext/socket/init.c    | 49 ++++++-----
ext/socket/rubysocket.h | 2 ++
ext/socket/socket.c  | 31 ++++++-----
ext/socket/unixsocket.c | 10 ++++++--
5 files changed, 77 insertions(+), 24 deletions(-)
```

Associated revisions

Revision 965b947f - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/trunk@44728 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

Revision 44728 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature #9330]

History

#1 - 01/18/2014 01:14 PM - akr (Akira Tanaka)

I feel rb_update_max_fd call should be exist near fd allocation site.
If they are far, it is difficult to confirm that rb_update_max_fd is called for all fd allocation.

#2 - 01/18/2014 09:14 PM - normalperson (Eric Wong)

akr@fsij.org wrote:

I feel rb_update_max_fd call should be exist near fd allocation site.
If they are far, it is difficult to confirm that rb_update_max_fd is called for all fd allocation.

OK. Also, thanks for [r44641](#) (early return from rb_update_max_fd),
I think we can drop patch 2/3.

How about just 1/3 and 3/3?

The following changes since commit 971ef822679dfa6ee63ff83a47b4e4d1aa60d146:

- ext/socket: Avoid unnecessary ppoll/select on Linux. Patch by Eric Wong. [ruby-core:57950] [Bug #9039] (2014-01-18 14:13:22 +0000)

are available in the git repository at:

git://80x24.org/ruby.git socket_cloexec_once-v2

for you to fetch changes up to 205abc14b8ebf84cf9016de260201e3f95044a72:

socket: avoid redundant fcntl calls for FD passing (2014-01-18 20:57:14 +0000)

Eric Wong (2):

socket: avoid redundant fcntl(fd, F_GETFD) with SOCK_CLOEXEC

socket: avoid redundant fcntl calls for FD passing

```
ext/socket/ancdata.c | 9 ++++++--
ext/socket/init.c    | 40 ++++++-----
ext/socket/rubysocket.h | 2 ++
ext/socket/socket.c  | 31 ++++++-----
ext/socket/unixsocket.c | 10 ++++++--
5 files changed, 73 insertions(+), 19 deletions(-)
```

#3 - 01/18/2014 10:53 PM - akr (Akira Tanaka)

The patch 1/3 and 3/3 are also have difficulty with rb_update_max_fd.

For example, rsock_socket0 returns fd without calling rb_update_max_fd.

This style makes us difficult to confirm that rb_update_max_fd is called for all fd allocation.

#4 - 01/18/2014 11:03 PM - normalperson (Eric Wong)

akr@fsij.org wrote:

The patch 1/3 and 3/3 are also have difficulty with rb_update_max_fd.

For example, rsock_socket0 returns fd without calling rb_update_max_fd.

This style makes us difficult to confirm that rb_update_max_fd is called for all fd allocation.

Sorry, I'll reroll the series.

#5 - 01/18/2014 11:52 PM - normalperson (Eric Wong)

Eric Wong normalperson@yhbt.net wrote:

Sorry, I'll reroll the series.

Updated patches also downloadable here:

1/2 <http://bogomips.org/ruby.git/patch/?id=b027a8b530ab6666>

2/2 <http://bogomips.org/ruby.git/patch/?id=1668aeb58a7bfe6f>

Pull request:

The following changes since commit d1d7f12c89cf683947740727fa37b6ce7fe512d4:

- compar.c (cmp_equal): warn for this release and still rescue standard exceptions for a nicer transition. See #7688. Partly reverts r44502. * test/ruby/test_comparable.rb: adapt assertion to match new behavior. (2014-01-18 21:40:58 +0000)

are available in the git repository at:

git://80x24.org/ruby.git socket_cloexec_once-v3

for you to fetch changes up to 1668aeb58a7bfe6f318f4fad3f8a4acdff3fbb99:

socket: avoid redundant fcntl calls for FD passing (2014-01-18 23:33:34 +0000)

Eric Wong (2):

socket: avoid redundant fcntl(fd, F_GETFD) with SOCK_CLOEXEC

socket: avoid redundant fcntl calls for FD passing

```
ext/socket/ancdata.c | 10 ++++++--
ext/socket/init.c    | 48 ++++++-----
ext/socket/rubysocket.h | 2 ++
ext/socket/socket.c  | 35 ++++++-----
ext/socket/unixsocket.c | 8 ++++++
5 files changed, 83 insertions(+), 20 deletions(-)
```

--
Eric Wong

#6 - 01/28/2014 02:37 PM - akr (Akira Tanaka)

- Status changed from Open to Closed
- % Done changed from 0 to 100

Applied in changeset [r44728](#).

-
- ext/socket: Avoid redundant fcntl/fstat syscalls for cloexec sockets. Patch by Eric Wong. [ruby-core:59429] [Feature [#9330](#)]

Files

0001-socket-avoid-redundant-fcntl-fd-F_GETFD-with-SOCK_CL.patch	5.87 KB	12/31/2013	normalperson (Eric Wong)
0002-socket-avoid-redundant-calls-to-rb_update_max_fd.patch	1.35 KB	12/31/2013	normalperson (Eric Wong)
0003-socket-avoid-redundant-fcntl-calls-for-FD-passing.patch	3.27 KB	12/31/2013	normalperson (Eric Wong)