

Ruby 1.8 - Bug #2722

gets on a large file takes a very very long time

02/08/2010 05:32 PM - ghazel (Greg Hazel)

Status:	Open
Priority:	Normal
Assignee:	
Target version:	
ruby -v:	ruby 1.8.6 (2009-08-04 patchlevel 383) [i386-mingw32]
Description	
=begin This problem occurs on 1.8.6, 1.8.7 and 1.9.1 (all the versions I tested) This simple script demonstrates the problem: # 100 MB n = 100 * 1000 * 1000 puts "writing" File.open("foo", 'wb'){ f f.write(" " * n) } puts "reading" File.open("foo", 'rb') do io io.gets end This takes about 1911 seconds. Using a 10MB file completes in 19 seconds, instead of in 1/10th the time as you would imagine. Similarly a 1MB file completes in 0.18 seconds. =end	
Related issues:	
Related to Ruby 1.8 - Bug #2741: gets with large file is slow in windoze	Open 02/13/2010

History

#1 - 02/09/2010 01:35 AM - matz (Yukihiko Matsumoto)

=begin
Hi,

In message "Re: [ruby-core:28103] [Bug #2722] gets on a large file takes a very very long time"
on Mon, 8 Feb 2010 17:33:09 +0900, Greg Hazel redmine@ruby-lang.org writes:

[This problem occurs on 1.8.6, 1.8.7 and 1.9.1 (all the versions I tested)

I could reproduce the problem on 1.9, but not on 1.8. 1.9 has been fixed by r26622. Thank you for the report.

matz.

=end

#2 - 02/10/2010 07:45 AM - matz (Yukihiko Matsumoto)

=begin
Hi,

In message "Re: [ruby-core:28138] Re: [Bug #2722] gets on a large file takes a very very long time"
on Wed, 10 Feb 2010 06:09:54 +0900, Roger Pack rogerdpack2@gmail.com writes:

[For me it performs fast when I replace .gets with .read
[Perhaps there is a reason for this?

gets need to scan for the newline. Here's my numbers

```
ruby 1.9.2dev (2010-02-09 trunk 26623) [i686-linux]  
gets: 0.13s user 0.42s system 88% cpu 0.626 total  
read: 0.08s user 0.46s system 93% cpu 0.578 total
```

ruby 1.8.8dev (2010-02-07 revision 26612) [i486-linux]
gets: 2.54s user 0.44s system 97% cpu 3.073 total
read: 2.47s user 0.45s system 97% cpu 3.007 total

1.8 has bottleneck on String#times, maybe we can work on it.

```
matz.
```

=end

#3 - 02/12/2010 11:21 AM - naruse (Yui NARUSE)

- Priority changed from Normal to 3

=begin

=end