

Ruby - Bug #7029

yaml#load_file too slow under Psych

09/16/2012 03:16 AM - mattneub (Matt Neuburg)

Status:	Closed	Backport:
Priority:	Normal	
Assignee:	tenderlovmaking (Aaron Patterson)	
Target version:		
ruby -v:	ruby 1.9.3p194 (2012-04-20 revision 35410) [x86_64-darwin10.8.0]	
Description		
<p>Unzip attached folder and run test.rb. Remember time output. Now uncomment the ENGINE line so we switch from Psych to Syck as yaml engine, and do it again. On my machine, time for Psych is about twice the time for Syck. (Time under old yaml under Ruby 1.8.7 is comparable to Syck here.)</p> <p>The example is artificial but in the actual use-case in my application this doubling of the time is killing performance for me. I regard this as a severe bug. I expect at least comparable performance. In my view the adoption of Psych yaml engine in 1.9.3 has been prematurely forced upon users, and should be rolled back until performance is comparable to Syck. At least the default should be reversed; for now, Syck should be the default and users can then choose Psych if they want it.</p> <p>ruby 1.9.3p194 (2012-04-20 revision 35410) [x86_64-darwin10.8.0]</p>		

History

#1 - 09/16/2012 06:31 AM - mattneub (Matt Neuburg)

- File yamlLoadFileTest2.zip added

Please ignore the previous test file. I've come up with a much better test. This is indicative of the real-world use-case. On my machine, the output is:

```
Psych
17.156296968460083
Syck
5.016614198684692
Resulting loaded hash the same? true
```

So Psych is taking over three times as long as Syck to load and parse the file. Times under Ruby 1.8.7 are comparable to the Syck time shown here. So this is a massive slowdown under Ruby 1.9.3.

#2 - 09/16/2012 02:38 PM - drbrain (Eric Hodel)

- Category set to ext
- Status changed from Open to Assigned
- Assignee set to tenderlovmaking (Aaron Patterson)

Please note that syck no longer exists in ruby trunk, so this benchmark cannot run

#3 - 07/11/2013 03:38 AM - tenderlovmaking (Aaron Patterson)

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

```
=begin
Hi,

The latest release of the psych gem cuts the time for this benchmark. It may still be slightly slower than Syck, but the parsers are different (YAML 1.0 vs YAML 1.1), so we can't really compare apples to apples. I'll continue to reduce bottlenecks as I can.

[aaron@higgins yamlLoadFileTest2]$ ruby test.rb
Psych version: 1.3.4
Psych
9.063196897506714
[aaron@higgins yamlLoadFileTest2]$ ruby -I ../lib/psych test.rb
Psych version: 2.0.0
```

Psych
5.260622024536133

=end

Files				
yamlLoadFileTest.zip	1.86 KB	09/16/2012	mattneub (Matt Neuburg)	
yamlLoadFileTest2.zip	12.8 KB	09/16/2012	mattneub (Matt Neuburg)	