

Ruby master - Feature #11066

symbol.to_s performance - caching and returning a frozen string

04/13/2015 08:33 PM - rdsbhas (Subhas Dandapani)

Status:	Open
Priority:	Normal
Assignee:	
Target version:	
Description	
Hi, this is my first report so please bear with me if its bad.	
Symbol.to_s creates a new string upon every call. This seems to be a huge waste of time. In typical Rails applications, "to_json" is called hundreds of times for ActiveRecord models on each request. Many symbols are being converted to string and GCd again and again. Is it possible to cache and return a Frozen string (string.freeze) when symbol.to_s is called? Now that symbols are also garbage collected, the string can also be garbage collected.	
I tried getting some performance benchmarks, but with Ruby 2.1, I'm not able to modify Symbol class. I believe Symbols are now frozen internally so any modifications are not possible. So I was not able to do any monkey patching or benchmarking. But I can confirm that new strings are getting created every time, please see attached file. I believe caching this would improve performance. It would also increase memory usage, but since symbols are also garbage collected, the memory can be freed up over time, and anyways we can cache only when "symbol.to_s" is invoked first time.	
Related issues:	
Related to Ruby master - Feature #11346: New #intern_str or similarly-named m...	Open

History

#1 - 04/15/2015 06:34 AM - nobu (Nobuyoshi Nakada)

I tried it and found some libraries (ostruct.rb at least) depend on the current behavior.

#2 - 07/13/2015 03:13 AM - nobu (Nobuyoshi Nakada)

- Related to Feature #11346: New #intern_str or similarly-named method of Symbol to return a cached, frozen string (like "string".freeze) added

#3 - 05/17/2016 06:37 AM - naruse (Yui NARUSE)

- Assignee deleted (ruby-core)

Files

symbol_to_s_mem_usage.rb	283 Bytes	04/13/2015	rdsbhas (Subhas Dandapani)
--------------------------	-----------	------------	----------------------------