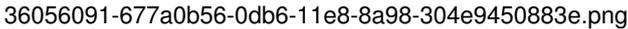


Ruby master - Feature #14460

Speed up `require` and reduce memory usage

02/10/2018 12:31 AM - tenderlovmaking (Aaron Patterson)

Status:	Closed
Priority:	Normal
Assignee:	
Target version:	
Description	
Hi,	
I've attached a patch that speeds up require and reduces memory usage. Before this patch, we could copy copy keys using <code>rb_strdup</code> for the <code>loaded_features_index</code> table. This patch changes the hash population functions to just index in to the Ruby strings without copying them. Some of the keys for the index require that the file names do not have a file extension. For that case, I create a new string (so that it's NULL padded), then index in to that string. Rather than coping the string for each / in the path, this patch just indexes in to the string. It relies on the loaded features array and the new array I introduced for liveness. <code>get_loaded_features_index</code> handles cache invalidation, so I believe it is safe to rely on these two arrays for liveness of the <code>loaded_features_index</code> keys.	
I tested memory usage by booting a Rails application in production mode (the same test I'm doing with the ISeq patches), and found that it saves about 4.2% memory (59580203 bytes in trunk, 57057395 bytes with this patch). I also found that this patch eliminated 76212 system calls (presumably calls to <code>malloc</code> / <code>free</code> for the hash table keys). Finally, I found that this patch sped up require by about 35%.	
Here is a graph of the memory usage:	
	
Each sample is a call to <code>malloc</code> , so you can see from the graph that this patch reduces memory usage and system calls.	
I ran the <code>bm_require.rb</code> benchmark to compare times:	
<pre>[aaron@TC-3 ruby (no-dup-string)]\$ make benchmark OPTS='-p bm_require.rb' /Users/aaron/.rbenv/shims/ruby --disable=gems ./benchmark/driver.rb -v \ --executables="/Users/aaron/.rbenv/shims/ruby --disable=gems -I./lib -I. -I.ext/co mmon --disable-gem; built-ruby:./miniruby -I./lib -I. -I.ext/common --disable-gem" \ --pattern='bm_' --directory=./benchmark -p bm_require.rb 2018-02-09 16:25:37 -0800 target 0: ruby 2.6.0dev (2018-02-07 trunk 62271) [x86_64-darwin17] at "/Users/aaron/.rbenv/shims/r uby --disable=gems -I./lib -I. -I.ext/common --disable-gem" target 1: built-ruby (ruby 2.6.0dev (2018-02-07 trunk 62271) [x86_64-darwin17]) at "./miniruby -I. /lib -I. -I.ext/common --disable-gem" measure target: real ----- require \$:push File.join(File.dirname(__FILE__), "bm_require.data") 1.upto(10000) do i require "c#{i}" end \$:pop ruby 2.6.0dev (2018-02-07 trunk 62271) [x86_64-darwin17] 3.043203999997786 built-ruby 1.9892130000007455 ----- raw data: [["require", [[3.043203999997786], [1.9892130000007455]]]] Elapsed time: 5.036855 (sec)</pre>	

benchmark results:

Execution time (sec)

name	ruby 2.6.0dev (2018-02-07 trunk 62271) [x86_64-darwin17]	built-ruby
require		3.043 1.989

Speedup ratio: compare with the result of `ruby 2.6.0dev (2018-02-07 trunk 62271) [x86_64-darwin17]` (greater is better)

name	built-ruby
require	1.530

Log file: bmlog-20180209-162537.60174.tsv

Related issues:

Related to Ruby master - Feature #8158: lightweight structure for loaded feat...

Closed

History

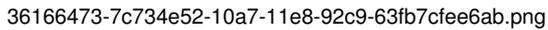
#1 - 02/13/2018 10:41 AM - funny_falcon (Yura Sokolov)

Did you compare it with <https://bugs.ruby-lang.org/issues/8158> ?

#2 - 02/13/2018 06:23 PM - tenderlovmaking (Aaron Patterson)

Hi,

I didn't see that ticket, but I made a comparison now:

36166473-7c734e52-10a7-11e8-92c9-63fb7cfee6ab.png

They are very close. My patch ends slightly larger 56867086 bytes vs 56502412 bytes. I like you patch much more. I think Matz offered you commit access in that ticket. Are you going to accept it? If so, I think you should apply that patch. If not, I will apply your patch.

Thanks!

#3 - 02/14/2018 04:56 AM - duerst (Martin Dürst)

- Related to Feature #8158: lightweight structure for loaded features index added

#4 - 02/14/2018 10:56 AM - funny_falcon (Yura Sokolov)

Aaron, while I'm going to accept commit right, I'm asking you to apply my patch.

With regards,
Yura.

#5 - 02/14/2018 06:39 PM - tenderlovmaking (Aaron Patterson)

- Status changed from Open to Closed

Fixed in r62404. See Feature [#8158](#)

Files

0001-Use-shared-substrings-in-feature-index-cache-hash.patch	6.05 KB	02/10/2018	tenderlovmaking (Aaron Patterson)
--	---------	------------	-----------------------------------