

Ruby trunk - Feature #14426

[PATCH] openssl: reduce memory allocation in OpenSSL::Buffering#do_write

01/31/2018 12:36 PM - janko (Janko Marohnić)

Status:	Closed
Priority:	Normal
Assignee:	rhenium (Kazuki Yamaguchi)
Target version:	
Description	
<p>When writing data to an SSLSocket, there are a lot of, in my opinion, unnecessary strings being allocated, concretely in OpenSSL::Buffering#do_write.</p> <p>When the buffer would be written, it would always be copied into a new string first, regardless of whether the write was partial or not. And in case of partial writes, it's not necessary to create copies of remaining data, we could just use the String[from, length] = "" trick immediately which modifies the string in-place.</p> <p>I also thought that splitting writes on newlines was adding unnecessary memory allocations, so I removed that.</p> <p>I tested uploading a 5MB file using HTTP.rb, and memory allocation went from 7.7 MB to 0.2 MB with this change.</p> <pre>require "http" require "memory_profiler" require "stringio" body = StringIO.new("a" * 5*1024*1024) MemoryProfiler.report do HTTP.post("https://example.com", body: body) end.pretty_print</pre>	

History

#1 - 01/31/2018 07:11 PM - normalperson (Eric Wong)

janko.marohnic@gmail.com wrote:

<https://bugs.ruby-lang.org/issues/14426>

I also thought that splitting writes on newlines was adding unnecessary memory allocations, so I removed that.

Agreed.

I tested uploading a 5MB file using HTTP.rb, and memory allocation went from 7.7 MB to 0.2 MB with this change.

Nice! This patch looks good to me, thanks.

#2 - 02/01/2018 12:04 AM - hsbt (Hiroshi SHIBATA)

- Assignee set to rhenium (Kazuki Yamaguchi)

- Status changed from Open to Assigned

#3 - 04/18/2018 01:16 PM - janko (Janko Marohnić)

Any updates on this one? I think it would be nice to have this included in the next patch version. This will make memory usage during large file uploads consistent regardless of whether it's going to HTTP or HTTPS endpoint; currently uploads over HTTPS consume a lot more memory, while HTTP uploads don't really allocate extra memory.

#4 - 05/31/2018 07:37 PM - janko (Janko Marohnić)

- File deleted (openssl-memory-allocation.patch)

- File openssl-memory-allocation.patch added

#5 - 07/07/2018 05:04 PM - janko (Janko Marohnić)

- File deleted (*openssl-memory-allocation.patch*)

#6 - 07/07/2018 05:07 PM - janko (Janko Marohnić)

- File *openssl-memory-allocation.patch* added

The patch got outdated in the latest trunk, so I updated it again. It's a really small change with big impact on memory usage when uploading files.

#7 - 08/08/2018 02:28 PM - rhenium (Kazuki Yamaguchi)

- Status changed from *Assigned* to *Closed*

The patch is now in ruby/openssl (upstream)'s master branch and Ruby trunk.

For the record, the GitHub Pull Requests are:

<https://github.com/ruby/ruby/pull/1924>

<https://github.com/ruby/openssl/pull/212>

Files

openssl-memory-allocation.patch	913 Bytes	07/07/2018	janko (Janko Marohnić)
---------------------------------	-----------	------------	------------------------