

Ruby master - Feature #6047

read_all: Grow buffer exponentially in generic case

02/20/2012 05:47 AM - MartinBosslet (Martin Bosslet)

Status:	Assigned
Priority:	Normal
Assignee:	openssl
Target version:	
Description	
In the general case, read_all grows its buffer linearly by just the amount that is currently read from the underlying source. This results in a linear number of reallocs, It might turn out beneficial if the buffer were grown exponentially by multiplying with a constant factor (e.g. 1.5 or 2), thus resulting in only a logarithmic number of reallocs.	
I will provide a patch and benchmarks, but I'm already opening this issue so I won't forget.	
See also https://bugs.ruby-lang.org/issues/5353 for more details.	

History

#1 - 10/27/2012 06:38 AM - ko1 (Koichi Sasada)

ping. status?
Do you need helps or comments?

#2 - 11/21/2012 11:49 AM - MartinBosslet (Martin Bosslet)

ko1 (Koichi Sasada) wrote:

ping. status?
Do you need helps or comments?

Thanks for your help, to be honest, I haven't tried so far. Can we leave it at 2.0.0 target for now? If I run into problems, I'll ask here!

#3 - 11/21/2012 12:53 PM - normalperson (Eric Wong)

Martin Bosslet Martin.Bosslet@gmail.com wrote:

In the general case, read_all grows its buffer linearly by just the amount that is currently read from the underlying source. This results in a linear number of reallocs, It might turn out beneficial if the buffer were grown exponentially by multiplying with a constant factor (e.g. 1.5 or 2), thus resulting in only a logarithmic number of reallocs.

I think growing the buffer exponentially makes sense.

I would enforce a hard limit (probably ≤ 8 MB) for each growth, to:

- 1) discourage read_all() for large files, it's very wasteful and usually hurts performance
- 2) prevent memory exhaustion for edge cases (especially on 32-bit)

#4 - 11/24/2012 11:47 AM - mame (Yusuke Endoh)

- Target version changed from 2.0.0 to 2.6

My experience also shows that it is useless to open a ticket for a reminder to myself :-)

I'm setting to next minor tentatively, but if it is really just a performance improvement (i.e., it affects no external modules), you can commit it to 2.0.0 before code freeze.

--
Yusuke Endoh mame@tsg.ne.jp

#5 - 09/13/2015 03:19 AM - zzak (Zachary Scott)

- Assignee changed from *MartinBosslet (Martin Bosslet)* to *openssl*