Ruby - Bug #4103

String#hash not returning consistent values in different sessions

12/01/2010 12:30 AM - ryanong (Ryan Ong)

Status:	Closed		
Priority:	Normal		
Assignee:			
Target version:	1.9.2		
ruby -v:	ruby 1.9.2p0 (2010-08-18 revision 29036) [i386-darwin10.4.0]	**Backport:**	

Description

=begin
I open one irb session

ruby-1.9.2-p0 > 'test'.hash
=> -658842761
ruby-1.9.2-p0 > 'test'.hash
=> -658842761

The second time I open it

ruby-1.9.2-p0 > 'test'.hash
=> 11032433
ruby-1.9.2-p0 > 'test'.hash
=> 11032433

I have no clue if this is on purpose or not but in 1.8.7 it was consistent across different sessions.
=end

History

#1 - 12/01/2010 02:07 AM - naruse (Yui NARUSE)

=begin
Hi,

(2010/12/01 0:30), Ryan Ong wrote:

> I open one irb session
>
> ruby-1.9.2-p0> 'test'.hash => -658842761 ruby-1.9.2-p0>
> 'test'.hash => -658842761
>
> The second time I open it
>
> ruby-1.9.2-p0> 'test'.hash => 11032433 ruby-1.9.2-p0> 'test'.hash
> => 11032433
>
> I have no clue if this is on purpose or not but in 1.8.7 it was
> consistent across different sessions.

It is intended. Ruby 1.9 explicitly use session local random seed
to calculate a hash for strings (and some other objects).

This is because the implementation of Object#hash is different
between versions (like 1.9.1 and 1.9.2) and implementations
(like JRuby, Rubinius, IronRuby, and so on). We want people to
write portable code around Object#hash, so we did so.

You should use Digest::SHA256 or some other digest routines
when you want some hash value (message digest).

--
NARUSE, Yui naruse@airemix.jp

=end

#2 - 12/01/2010 03:40 AM - shyouhei (Shyouhei Urabe)

- Status changed from Open to Closed

=begin
See also: http://perldoc.perl.org/perlsec.html#Algorithmic-Complexity-Attacks
=end

#3 - 12/08/2010 11:19 AM - duerst (Martin Dürst)

=begin
On 2010/12/01 2:07, NARUSE, Yui wrote:

> It is intended. Ruby 1.9 explicitly use session local random seed
> to calculate a hash for strings (and some other objects).
>
> This is because the implementation of Object#hash is different
> between versions (like 1.9.1 and 1.9.2) and implementations
> (like JRuby, Rubinius, IronRuby, and so on). We want people to
> write portable code around Object#hash, so we did so.

Also, it helps to avoid some denial of service attacks, such as registering hundreds and thousands of users with usernames that have the same hash code.

Regards, Martin.

--
#-# Martin J. Dürst, Professor, Aoyama Gakuin University
#-# http://www.sw.it.aoyama.ac.jp mailto:duerst@it.aoyama.ac.jp

=end