

Ruby - Bug #18486

Enumerable#group_by ordering appears to have changed and doesn't match docs.

01/13/2022 04:05 PM - masukomi@masukomi.org (kay rhodes)

Status:	Closed	
Priority:	Normal	
Assignee:		
Target version:		
ruby -v:	3.0.2	Backport: 2.6: UNKNOWN, 2.7: UNKNOWN, 3.0: UNKNOWN, 3.1: UNKNOWN
Description <p>The documentation for Enumerable#group_by suggests the following:</p> <pre>(1..6).group_by { i i%3 } #=> {0=>[3, 6], 1=>[1, 4], 2=>[2, 5]}</pre> <p>it actually produces this (same values different order). Tested in 2.6.3 and 3.0.2</p> <pre>(1..6).group_by { i i%3 } #=> {1=>[1, 4], 2=>[2, 5], 0=>[3, 6]}</pre> <p>as hash table ordering is something people depend on these days I wonder if the reordering of the keys may be more than just a minor documentation bug.</p> <p>I noticed this in 2.6.3. Assuming the docs matched reality when written this implies we've been living with this for a while. I don't know how long before that this actually changed.</p>		

History

#1 - 01/13/2022 04:31 PM - jeremyevans0 (Jeremy Evans)

This is only a minor documentation issue. I assume the documentation is a left over from pre-1.9, when hashes were unordered. I'll fix it.

#2 - 01/13/2022 04:33 PM - jeremyevans0 (Jeremy Evans)

- Status changed from Open to Closed

This was already fixed in [1e10099e091](#)

#3 - 01/13/2022 04:34 PM - zverok (Victor Shepelev)

I think it is docs bug.

group_by emerged in Ruby 1.8.7, but my RVM can't install it :(
The earliest I could check was 1.9.3, and it already has order of keys 1, 2, 0.

I think it is a natural reason of keys being in the order they are inserted: the first value for grouping is 1 (1 % 3), so that's why it appears first in the result.

With the preservation of Hash keys order introduced at 1.9, I think that's when the behavior changed :)))
In Ruby 1.8.6, imitation of group_by with inject does the initially documented order:

```
(1..6).to_a.inject({}) { |r, i| (r[i % 3] ||= []) << i; r }  
# => {0=>[3, 6], 1=>[1, 4], 2=>[2, 5]}
```

#4 - 01/13/2022 07:12 PM - matsuda (Akira Matsuda)

group_by emerged in Ruby 1.8.7, but my RVM can't install it :(
The earliest I could check was 1.9.3, and it already has order of keys 1, 2, 0.

off-topic: all-ruby is your friend :) <https://github.com/akr/all-ruby>

```
$ docker run -it --rm rubylang/all-ruby ./all-ruby -e 'p (1..6).to_a.inject({}) { |r, i| (r[i % 3] ||= []) << i; r }'  
ruby-0.49 -e:1: syntax error  
(snip)
```

```
ruby-1.6.8      1..6
                -e:1: undefined method `inject' for []:Array (NameError)
                exit 1
ruby-1.8.0      {0=>[3, 6], 1=>[1, 4], 2=>[2, 5]}
...
ruby-1.8.7-p374 {0=>[3, 6], 1=>[1, 4], 2=>[2, 5]}
ruby-1.9.0-0    {1=>[1, 4], 2=>[2, 5], 0=>[3, 6]}
...
ruby-3.1.0      {1=>[1, 4], 2=>[2, 5], 0=>[3, 6]}
```