

## Ruby master - Feature #8158

### lightweight structure for loaded features index

03/24/2013 04:56 PM - funny\_falcon (Yura Sokolov)

<b>Status:</b>	Closed	
<b>Priority:</b>	Normal	
<b>Assignee:</b>	nobu (Nobuyoshi Nakada)	
<b>Target version:</b>		
<b>Description</b>		
Use lightweight structure for loaded_features index:		
<ul style="list-style-type: none"><li>• use hand made simple hash structure, which uses only one memory chunk,</li><li>• do not store feature name string, only hash of it, since loaded_feature_path will recheck feature name on hash collision</li><li>• use single linked lists instead of arrays for storing features indices.</li><li>• store this lists inside one array, using array's indices as a reference.</li></ul>		
While startup speedup improvement is relatively small compared current implementation, this one does not need any Ruby Objects at all, so that there is no pressure on GC.		
<a href="https://github.com/ruby/ruby/pull/264.patch">https://github.com/ruby/ruby/pull/264.patch</a> <a href="https://github.com/ruby/ruby/pull/264.diff">https://github.com/ruby/ruby/pull/264.diff</a> <a href="https://github.com/ruby/ruby/pull/264">https://github.com/ruby/ruby/pull/264</a>		
<b>Related issues:</b>		
Related to Ruby master - Feature #14460: Speed up `require` and reduce memory...		<b>Closed</b>

### History

#### #1 - 03/24/2013 10:47 PM - funny\_falcon (Yura Sokolov)

Same for ruby\_2\_0\_0 branch

[https://github.com/funny-falcon/ruby/compare/ruby\\_2\\_0\\_0...features\\_index/ruby\\_2\\_0\\_0.diff](https://github.com/funny-falcon/ruby/compare/ruby_2_0_0...features_index/ruby_2_0_0.diff)  
[https://github.com/funny-falcon/ruby/compare/ruby\\_2\\_0\\_0...features\\_index/ruby\\_2\\_0\\_0](https://github.com/funny-falcon/ruby/compare/ruby_2_0_0...features_index/ruby_2_0_0)

#### #2 - 04/05/2013 11:43 AM - zzak (Zachary Scott)

- File 264.patch added  
- Category set to core

#### #3 - 10/01/2013 04:58 PM - nobu (Nobuyoshi Nakada)

- Status changed from Open to Assigned  
- Assignee set to nobu (Nobuyoshi Nakada)

#### #4 - 01/30/2014 06:16 AM - hsb (Hiroshi SHIBATA)

- Target version changed from 2.1.0 to 2.2.0

#### #5 - 10/07/2016 01:20 PM - funny\_falcon (Yura Sokolov)

- File 0001-load.c-reduce-memory-usage-of-loaded\_features\_index.patch added

#### #6 - 10/07/2016 01:24 PM - funny\_falcon (Yura Sokolov)

I've changed implementation a bit:  
Because <https://bugs.ruby-lang.org/issues/12142> likely to be accepted,  
I've used st\_table with numtable instead of separate datastructure.

So patch now is shorter.

[https://bugs.ruby-lang.org/attachments/download/6182/0001-load.c-reduce-memory-usage-of-loaded\\_features\\_index.patch](https://bugs.ruby-lang.org/attachments/download/6182/0001-load.c-reduce-memory-usage-of-loaded_features_index.patch)

#### #7 - 12/12/2016 06:09 AM - normalperson (Eric Wong)

[funny.falcon@gmail.com](mailto:funny.falcon@gmail.com) wrote:

So patch now is shorter.

Thanks; this got broken by trivial whitespace change in r57032

I'm not sure about the portability of initializing structs with non-const values for some compilers. Maybe usa or somebody with more portability experience can comment.

Also, in C Ruby; initial indent is 4 spaces, second indent is hard tab. Not my rule and I hate it; but not much we can do...

**#8 - 12/12/2016 06:41 AM - usa (Usaku NAKAMURA)**

Hi,

In message "[ruby-core:78595] Re: [Ruby trunk Feature#8158] lightweight structure for loaded features index" on Mon, 12 Dec 2016 06:06:13 +0000, [normalperson@yhbt.net](mailto:normalperson@yhbt.net) wrote:

I'm not sure about the portability of initializing structs with non-const values for some compilers. Maybe usa or somebody with more portability experience can comment.

It's OK, at least for VC++.

Regards,

--

U.Nakamura [usa@garbagecollect.jp](mailto:usa@garbagecollect.jp)

**#9 - 12/12/2016 07:20 AM - nobu (Nobuyoshi Nakada)**

Usaku NAKAMURA wrote:

In message "[ruby-core:78595] Re: [Ruby trunk Feature#8158] lightweight structure for loaded features index" on Mon, 12 Dec 2016 06:06:13 +0000, [normalperson@yhbt.net](mailto:normalperson@yhbt.net) wrote:

I'm not sure about the portability of initializing structs with non-const values for some compilers. Maybe usa or somebody with more portability experience can comment.

It's OK, at least for VC++.

Other compilers fail, IIRC, Solaris, AIX, HP-UX, or something.

**#10 - 12/12/2016 08:05 AM - funny\_falcon (Yura Sokolov)**

I'll fix patch today.

**#11 - 12/12/2016 10:39 AM - funny\_falcon (Yura Sokolov)**

- File deleted (264.patch)

**#12 - 12/12/2016 10:39 AM - funny\_falcon (Yura Sokolov)**

- File deleted (0001-load.c-reduce-memory-usage-of-loaded\_features\_index.patch)

**#13 - 12/12/2016 10:39 AM - funny\_falcon (Yura Sokolov)**

- File 0001-load.c-reduce-memory-usage-of-loaded\_features\_index.patch added

**#14 - 12/12/2016 10:44 AM - funny\_falcon (Yura Sokolov)**

I've uploaded fixed patch:

[https://bugs.ruby-lang.org/attachments/download/6293/0001-load.c-reduce-memory-usage-of-loaded\\_features\\_index.patch](https://bugs.ruby-lang.org/attachments/download/6293/0001-load.c-reduce-memory-usage-of-loaded_features_index.patch)

Github branch and diff:

[https://github.com/funny-falcon/ruby/tree/loaded\\_features\\_strings](https://github.com/funny-falcon/ruby/tree/loaded_features_strings)

[https://github.com/ruby/ruby/compare/trunk...funny-falcon:loaded\\_features\\_strings.patch](https://github.com/ruby/ruby/compare/trunk...funny-falcon:loaded_features_strings.patch)

**#15 - 12/12/2016 11:22 AM - funny\_falcon (Yura Sokolov)**

- File load.c-loaded\_features\_numindex.patch added

**#16 - 12/12/2016 11:23 AM - funny\_falcon (Yura Sokolov)**

- File deleted (*load.c-loaded\_features\_numindex.patch*)

**#17 - 12/12/2016 11:24 AM - funny\_falcon (Yura Sokolov)**

- File *load.c-loaded\_features\_numindex.patch* added

**#18 - 12/12/2016 11:30 AM - funny\_falcon (Yura Sokolov)**

I've tried to simplify patch: remove struct `feature_str` and use just pointer with `len`.

[https://bugs.ruby-lang.org/attachments/download/6295/load.c-loaded\\_features\\_numindex.patch](https://bugs.ruby-lang.org/attachments/download/6295/load.c-loaded_features_numindex.patch)

[https://github.com/ruby/ruby/compare/trunk...funny-falcon:loaded\\_features\\_index.patch](https://github.com/ruby/ruby/compare/trunk...funny-falcon:loaded_features_index.patch)

[https://github.com/funny-falcon/ruby/tree/loaded\\_features\\_index](https://github.com/funny-falcon/ruby/tree/loaded_features_index)

**#19 - 12/21/2016 02:25 PM - shyouhei (Shyouhei Urabe)**

We looked at this issue at today's developer meeting and had positive opinions. Maybe introduced in 2.5.

**#20 - 01/19/2017 05:48 AM - matz (Yukihiro Matsumoto)**

The patch seems OK now. [funny\\_falcon \(Yura Sokolov\)](#) do you want to check in by yourself?

Matz.

**#21 - 01/19/2017 11:11 AM - funny\_falcon (Yura Sokolov)**

Matz,

I don't know English enough to clearly understand your suggestion.

Is it suggestion of commit rights?

If it is, it will be a great honor for me.

But I'm not even-tempered person, and I fear I will put a mess into repository if I will have commit rights.

If your suggestion means some-thing else, then may you formulate it in some other words so I can understand?

With regards,

Yura.

**#22 - 02/13/2018 09:39 PM - sam.saffron (Sam Saffron)**

Hi Yura,

I think the commit rights here reduce the amount of work for the rest of the Ruby team.

This change is approved and reviewed, so you can commit it directly once you have the rights, no need to make other people do the committing.

Very keen to have this change included in Ruby.

Sam

**#23 - 02/14/2018 04:56 AM - duerst (Martin Dürst)**

- Related to Feature #14460: *Speed up `require` and reduce memory usage* added

**#24 - 02/14/2018 10:53 AM - funny\_falcon (Yura Sokolov)**

Year after I think I am more stable person, so I'm going to accept commit rights with gratitude, honor and great responsibility.

But until I passed procedure, I'd be thankful if [tenderlovmaking \(Aaron Patterson\)](#) will commit this one.

With regards,

Yura.

**#25 - 02/14/2018 06:39 PM - tenderlovmaking (Aaron Patterson)**

- Status changed from *Assigned* to *Closed*

Committed in r62404. Thanks Yura!!

**#26 - 02/15/2018 08:12 PM - normalperson (Eric Wong)**

<https://bugs.ruby-lang.org/issues/8158>

I am curious, what is the significance of the 0xfea7009e initializer passed to st\_hash?

Would it be appropriate to use for the `st\_hash` call in hash.c::hash\_i as well?

Thanks.

**#27 - 02/20/2018 06:34 AM - matz (Yukihiro Matsumoto)**

[funny\\_falcon \(Yura Sokolov\)](#)? Will you be a committer?

Matz.

**#28 - 02/20/2018 10:40 AM - funny\_falcon (Yura Sokolov)**

[matz \(Yukihiro Matsumoto\)](#) , yes I will.

Excuse me for the delay. I'll try to proceed with steps from CommitterHowto this week.

With regards,  
Yura.

**#29 - 02/20/2018 10:49 AM - funny\_falcon (Yura Sokolov)**

normalperson (Eric Wong) wrote:

<https://bugs.ruby-lang.org/issues/8158>

I am curious, what is the significance of the 0xfea7009e initializer passed to st\_hash?

0xfea7009e - is a "feature" hex-spelled :-). Nothing special. It could be omitted, given features are strings without zero bytes. I put it here just for fun.

Would it be appropriate to use for the `st\_hash` call in hash.c::hash\_i as well?

I believe, there's no real needs.

Thanks.

With regards,  
Yura.

**#30 - 02/23/2018 12:58 PM - hsbt (Hiroshi SHIBATA)**

Hi all.

I did create an account of funny\_falcon on svn.ruby-lang.org.

## Files

0001-load.c-reduce-memory-usage-of-loaded_features_index.patch	5.99 KB	12/12/2016	funny_falcon (Yura Sokolov)
load.c-loaded_features_numindex.patch	4.99 KB	12/12/2016	funny_falcon (Yura Sokolov)