

Ruby - Feature #3478

Excruciatingly slow pathname implementation

06/25/2010 05:33 AM - stouset (Stephen Touset)

Status:	Closed	
Priority:	Normal	
Assignee:	akr (Akira Tanaka)	
Target version:	2.0.0	
Description		
<pre>=begin</pre> <p>In writing a pure-Ruby git implementation, I discovered with rprof that using Pathname was the source of a huge proportion of my library's running time. Recently, Rails contributor José Valim discovered a similar situation in Rails. Anecdotally, removing Pathname resulted in a speedup from 6s to 3s on a single request. Here's his commit removing some Pathname usage from Rails:</p> <pre>http://github.com/rails/rails/commit/69abbe893413c99808c8cebd2f1d468c7921c573</pre> <p>I have rewritten Pathname, and the project is on GitHub.</p> <pre>http://github.com/stouset/pathname3/</pre> <p>I've included the rewrite as a patch. It's backwards-compatible with the existing Pathname class, and passes RubySpec and MRI tests.</p> <pre>=end</pre>		

History

#1 - 06/25/2010 08:26 AM - mame (Yusuke Endoh)

- Assignee set to akr (Akira Tanaka)

- Target version set to 2.0.0

```
=begin
```

Hi,

This is not a bug but feature request. This ticket is moved to Feature tracker.

Your patch might be good, but apparently too big. I guess it will take long time to review it.

I recommend you to explain essential changes for performance improvement, and to split the patch for each change. And, do not remove documents without reason.

--
Yusuke Endoh mame@tsg.ne.jp
=end

#2 - 06/25/2010 01:29 PM - stouset (Stephen Touset)

```
=begin
```

The patch would be very difficult to do change by change, partly because the entire class is rewritten to descend from String. Suffice it to say, virtually all of the methods in the original Pathname library tried to implement logic themselves, sometimes in a very inefficient manner, rather than falling back upon the well-optimized implementations in File. This rewrite is akin to the switch from CSV to FasterCSV, which I don't believe was submitted as a series of patches. As the code passes all of the related Ruby specs, I wouldn't think this would take too much effort to review and merge in.

The documentation was different because I rewrote the class without the anticipation of having it become part of standard Ruby. I can change the code to mimic the original comments in Pathname without too much effort, if it would help get the patch accepted.

```
=end
```

#3 - 06/25/2010 02:38 PM - shyouhei (Shyouhei Urabe)

```
=begin
```

-1. Seems your lib lacks Windows support.

```
=end
```

#4 - 06/25/2010 09:55 PM - stouset (Stephen Touse)

=begin
As I understand it, so does the original. Again, it passes all existing tests.
=end

#5 - 06/25/2010 11:24 PM - shyuhei (Shyuehei Urabe)

=begin
No, your understanding is wrong.

C:\Users\shyuehei\Documents>C:\ruby\bin\ruby.exe -rpathname -e"p Pathname('C:\').relative_path_from Pathname('D:\')"
C:/ruby/lib/ruby/1.8/pathname.rb:723:in `relative_path_from': different prefix:
"C:\" and "D:\" (ArgumentError)
from -e:1

C:\Users\shyuehei\Documents>C:\ruby\bin\ruby.exe -rpathname3 -e"p Pathname('C:\').relative_path_from Pathname('D:\')"
"./C:"

=end

#6 - 06/25/2010 11:31 PM - stouset (Stephen Touse)

=begin
I'll be glad to address those issues, then. Give me a day or two to add some Windows support.
=end

#7 - 07/10/2010 06:48 AM - stouset (Stephen Touse)

=begin
I haven't forgot about fixing the implementation to support Windows. I had to move apartments earlier than expected, and haven't had the free time to add the necessary fixes.
=end

#8 - 09/09/2010 09:30 AM - naruse (Yui NARUSE)

- Status changed from Open to Closed

=begin
This problem will be resolved by C implementation of Pathname; see and test trunk.
=end

Files

pathname.diff	42.2 KB	06/25/2010	stouset (Stephen Touse)
---------------	---------	------------	-------------------------