Ruby master - Bug #15935

Memory leak triggered by String#encode, possibly elsewhere too

06/17/2019 08:04 PM - luke-gru (Luke Gruber)

Status:	Closed		
Priority:	Normal		
Assignee:			
Target version:			
ruby -v:		**Backport:**	2.4: UNKNOWN, 2.5: REQUIRED, 2.6: REQUIRED

Description

Hi, I've found a leak that can be reproduced in the following way:

```
loop do
    puts "running..."

    50.times do
        File.open("./test/rexml/data/utf16.xml", external_encoding: 'UTF-16LE', binmode: true) do
|f| # must be in ruby's srcdir
            while line = f.readline()
                line.encode("UTF-8", "UTF-16LE")
            end
        end
    rescue EOFError
    end
end
```

It seems to be due to rb_enc_associate_index(), in encoding.c, which can call rb_str_change_terminator_length() with
the given string. Under certain conditions, str_make_independent_expand() is called in this function. This function
can allocate a new heap buffer if the string is large enough to not be embeddable, but does not free the previous one, if
necessary.

The following patch seems to fix the leak:

```
    oldptr = RSTRING_PTR(str);
    if (oldptr) {
    memcpy(ptr, oldptr, len);
    }
    if (!str_dependent_p(str) && !FL_TEST(str, STR_NOFREE) && FL_TEST(str, STR_NOEMBED)) {
        xfree(oldptr);
    }
```

I can add PR if you want, or you can fix it as you see fit with whichever code you prefer.

NOTE: I found this leak by adding more debug assertions to string.c. Basically I added the same code as above, except instead of
xfree I asserted that the string should never
have a freeable buffer, as this is the assumption the function seemed to make.

Then, when running make test-all, I run into this failed assertion a bunch, especially in rexml tests. This also causes the leak:

```
loop do
  File.open("./test/rexml/data/utf16.xml") do |f|
    REXML::Document.new(f)
  end
end
```

Thanks for your time :)

Associated revisions

Revision 8b3774be - 06/18/2019 04:40 AM - nobu (Nobuyoshi Nakada)

Fix memory leak

- string.c (str_make_independent_expand): free independent buffer. [Bug# 15935]

Co-Authored-By: luke-gru (Luke Gruber) luke.gru@gmail.com

History

#1 - 06/19/2019 06:00 PM - jeremyevans0 (Jeremy Evans)

- Status changed from Open to Closed

#2 - 09/09/2019 03:14 AM - nagachika (Tomoyuki Chikanaga)

- Backport changed from 2.4: UNKNOWN, 2.5: UNKNOWN, 2.6: UNKNOWN to 2.4: UNKNOWN, 2.5: REQUIRED, 2.6: REQUIRED