

Ruby - Feature #6752

Replacing ill-formed subsequences

07/19/2012 11:42 AM - naruse (Yui NARUSE)

Status:	Closed
Priority:	Normal
Assignee:	matz (Yukihiro Matsumoto)
Target version:	2.6

Description

```
=begin
==  []
String[]
```

```
== 000000  
000000000000000000000000000000000000000000
```

- twitterのtitle
- IRCの
- のAPI
- Webの
 - の
 - の
 - の
- <https://twitter.com/takahashim/status/18974040397>
- <https://twitter.com/n0kada/status/215674740705210368>
- <https://twitter.com/n0kada/status/215686490070585346>
- <https://twitter.com/hajimehoshi/status/215671146769682432>
- <http://po-ru.com/diary/fixing-invalid-utf-8-in-ruby-revisited/>
- <http://stackoverflow.com/questions/2982677/ruby-1-9-invalid-byte-sequence-in-utf-8>

[illegible]

http://unicode.org/reports/tr36/#UTF-8_Exploit

```
== API
--- str.encode(str.encoding, invalid: replace, [replace: ""])
```

- CSI
- iconv glibc iconv GNU libiconv //IGNORE
- (

==

- 000000000000
- fix/repair invalid/illegal bytes/sequence 00000000

```
== 00
== 000000
int ret = rb_enc_precise_mbclen(p, e, enc); 000
MBCLEN_INVALID P(ret) 00000000000000000000000000000000
```

ONIGENC_CONSTRUCT_MBCLEN_INVALID() []
[]
[]

=== transcode[]
UCS[]glibc iconv, GNU libiconv, Perl Encode[]
CS[]transcode[]
[]

[]

```
diff --git a/string.c b/string.c
index d038835..4808f15 100644
--- a/string.c
+++ b/string.c
@@ -7426,6 +7426,199 @@ rb_str_ellipsis(VALUE str, long len)
return ret;
}
```

+/

- call-seq:
- str.fix_invalid -> new_str
- If the string is well-formed, it returns self.
- If the string has invalid byte sequence, repair it with given replacement
- character.
- */
+VALUE
+rb_str_fix_invalid(VALUE str)
+{
• int cr = ENC_CODERANGE(str);
• rb_encoding *enc;
• if (cr == ENC_CODERANGE_7BIT || cr == ENC_CODERANGE_VALID)
• return rb_str_dup(str);

• enc = STR_ENC_GET(str);
• if (rb_enc_asciicompat(enc)) {
• const char *p = RSTRING_PTR(str);
• const char *e = RSTRING_END(str);
• const char *p1 = p;
• /* 10 should be enough for the usual use case,
 - fixing a wrongly chopped character at the end of the string
• */
• long room = 10;
• VALUE buf = rb_str_buf_new(RSTRING_LEN(str) + room);
• const char *rep;
• if (enc == rb_utf8_encoding())

•

• else

•

• cr = ENC_CODERANGE_7BIT;

• p = search_nonascii(p, e);
• if (!p) {

- [illegible]

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- }
- if ($p_1 < p$) {

• [REDACTED]

• }
• if ($p < e$) {

• [REDACTED]

• [REDACTED]

• }
• ENCODING_CODERANGE_SET(buf, rb_enc_to_index(enc), cr);
• return buf;
• }
• else if (rb_enc_dummy_p(enc)) {
• return rb_str_dup(str);
• }
• else {
• /* ASCII incompatible */
• const char *p = RSTRING_PTR(str);
• const char *e = RSTRING_END(str);
• const char *p1 = p;
• /* 10 should be enough for the usual use case,
◦ fixing a wrongly chopped character at the end of the string
• */
• long room = 10;
• VALUE buf = rb_str_buf_new(RSTRING_LEN(str) + room);
• const char *rep;

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

● [REDACTED]

● [REDACTED]

- [REDACTED]

- [REDACTED]

- [REDACTED]

• [REDACTED]

• [REDACTED]

● [REDACTED]

• [REDACTED]

● [REDACTED]

• [REDACTED]

• [REDACTED]

-
- A horizontal bar chart consisting of 30 rows. Each row begins with a black bullet point, followed by a black horizontal bar of varying length. The bars represent different values, with some being significantly longer than others. The bars are set against a white background.

- if (p1 < p) {

- [REDACTED]

- }

- if (p < e) {

- [REDACTED]

- }

- ENCODING_CODERANGE_SET(buf, rb_enc_to_index(enc), ENC_CODERANGE_VALID);

- return buf;

- }

- +}

```
/******
```

- Document-class: Symbol

```
@@ -7882,6 +8075,7 @@ Init_String(void)
```

```
rb_define_method(rb_cString, "getbyte", rb_str_getbyte, 1);
```

```
rb_define_method(rb_cString, "setbyte", rb_str_setbyte, 2);
```

```
rb_define_method(rb_cString, "byteslice", rb_str_byteslice, -1);
```

- rb_define_method(rb_cString, "fix_invalid", rb_str_fix_invalid, 0);

```
rb_define_method(rb_cString, "to_i", rb_str_to_i, -1);
```

```
rb_define_method(rb_cString, "to_f", rb_str_to_f, 0);
```

```
diff --git a/test/ruby/test_string.rb b/test/ruby/test_string.rb
```

```
index 47f349c..2b0cfcb 100644
```

```
--- a/test/ruby/test_string.rb
```

```
+++ b/test/ruby/test_string.rb
```

```
@@ -2031,6 +2031,29 @@ class TestString < Test::Unit::TestCase
```

```
assert_equal(u("\x82")+("\u3042"*9), ("\u3042"*10).byteslice(2, 28))
```

```
end
```

- def test_fix_invalid

- assert_equal("\uFFFF\uFFFF\uFFFF", "\x80\x80\x80".fix_invalid)

- assert_equal("\uFFFDA", "\xF4\x80\x80A".fix_invalid)

• exapmles in Unicode 6.1.0 D93b

- assert_equal("\x41\uFFFF\uFFFF\x41\uFFFF\x41",

- [REDACTED]

- assert_equal("\x41\uFFFF\uFFFF\uFFFF\x41",

- [REDACTED]

- assert_equal("\u0061\uFFFF\uFFFF\uFFFF\u0062\uFFFF\u0063\uFFFF\uFFFF\u0064",

- [REDACTED]

- assert_equal("abcdefghijklmnopqrstuvwxyz\u0061\uFFFF\uFFFF\uFFFF\u0062\uFFFF\u0063\uFFFF\uFFFF\u0064",

- [REDACTED]

- string.c (rb_str_scrub): fix for UTF-32. strlen() on strings contain NUL returns wrong result, use sizeof operator instead. [ruby-dev:45975] [Feature #6752]

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/trunk@40417 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

History

#1 - 10/05/2012 09:47 AM - nobu (Nobuyoshi Nakada)

- Description updated

#2 - 10/05/2012 11:10 AM - kosaki (Motohiro KOSAKI)

XXXXXXXXXXXXXXXXXXXXTwitterXXXXXXXXXXXXXXXXXXXXXXXXXXXX

```
XXXXXXXXXXencode()XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXXXXXXX replace_invalid_character XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
XXXXencode X invalid => replace XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX replace_invalid XX
XXXXXXXXXXXXXXXXXXXX( )XXXXXXXXXXXXXXXXXXXX UnicodeXXXXU+FFFD XXXXXXXXXXXXXXXXXXXXXXX
```

#3 - 11/10/2012 05:23 PM - duerst (Martin Dürst)

- Target version changed from 2.0.0 to 2.6

I have thought about this a bit. Yui's patch to string treats this as a problem separat from transcoding. I think it is preferable to use the transcoding logic to implement this. The advantage is that exactly the same options and fallbacks can be used, and if we add a new option or fallback to transcode, it will be usable, too.

Some more notes: The checks for converting from one encoding to the same encoding are in str_transcode0. Anywhere else? We need some data to drive the conversion, but this should be easy to generate, and will be the same for many 8-bit encodings.

It will be easy to catch invalid byte sequences, but I'm not sure it's worth to check unassigned codepoints, at least not in Unicode.

I have changed the target version from 2.0.0 to next minor, because I don't think this will be ready for 2.0.0. But please change back if somebody can do it faster.

#4 - 01/09/2013 05:58 PM - knu (Akinori MUSHASHI)

=begin
(('+1')) for the functionality.

What we currently have (Encoding::Converter) is not enough to solve this problem because 1) it is mandatory to pick a different encoding to convert to for nothing, and even if you have decided to pick one 2) #primitive_errinfo does not give you offset information that is necessary to locate where the found invalid bytes are.

In addition to this proposal, I'd like String#encode to accept a proc instead of a fixed string as a :replace option to get a callback for each invalid byte so you can dynamically compose a replace string for the given invalid byte. Perl's encode() and decode() have this feature and it's very handy to investigate and visualize how a string is garbled. (e.g. ({replace: ->(byte) { "e[7m<%02X>\e[m" % byte }]))

I don't have a particular opinion about the API, but having self-transcoders perform validation as [@duerst \(Martin Dürst\)](#) implies sounds great to me if it could be properly implemented.

My tentative solution that is known to be very slow is here: ([URL:https://gist.github.com/4491446](https://gist.github.com/4491446))
=end

#5 - 04/09/2013 04:38 AM - naruse (Yui NARUSE)

duerst (Martin Dürst) wrote:

I have thought about this a bit. Yui's patch to string treats this as a problem separat from transcoding. I think it is preferable to use the transcoding logic to implement this. The advantage is that exactly the same options and fallbacks can be used, and if we add a new option or fallback to transcode, it will be usable, too.

This method doesn't need same options and fallbacks.
It need only invalid related, doesn't need undef related.
Moreover transcoder is usable only if Ruby has related transcoder of the target encoding.
But Ruby has some encodings which doesn't have transcoder for example emacs-mule.
Therefore this can't be built on transcoder.

Some more notes: The checks for converting from one encoding to the same encoding are in str_transcode0. Anywhere else? We need some data to drive the conversion, but this should be easy to generate, and will be the same for many 8-bit encodings.

Yeah, I came to str_transcode0 and it is correct place.

The date we need is problem.

transcode doesn't have all the data though tool/transcode-tblgen.rb has some base data.

The only one which has all the data we need is enc/*.

It will be easy to catch invalid byte sequences, but I'm not sure it's worth to check unassigned codepoints, at least not in Unicode.

If we need unassigned codepoints, we must define encodings more strictly.

Even if it is Unicode, it needs versions.

I don't think it's worth to check.

#6 - 04/09/2013 06:24 PM - naruse (Yui NARUSE)

I wrote a updated patch which include String#scrub and String#encode with extension.

String#scrub allows replacement as both argument or block.

```
diff --git a/string.c b/string.c
index 8b85739..bc973dc 100644
--- a/string.c
+++ b/string.c
@@ -7741,6 +7741,271 @@ rb_str_ellipsize(VALUE str, long len)
 return ret;
 }

+static VALUE
+str_compat_and_valid(VALUE str, rb_encoding *enc)
+{
+    • int cr;
+    • str = StringValue(str);
+    • cr = rb_enc_str_coderange(str);
+    • if (cr == ENC_CODERANGE_BROKEN) {
+    • rb_raise(rb_eArgError, "replacement must be valid byte sequence '%+PRIsVALUE'", str);
+    • }
+    • else if (cr == ENC_CODERANGE_7BIT) {
+    • rb_encoding *e = STR_ENC_GET(str);
+    • if (!rb_enc_asciicompat(enc)) {
+
+    •
+
+    • }
+    • }
+    • else { /* ENC_CODERANGE_VALID */
+    • rb_encoding *e = STR_ENC_GET(str);
+    • if (enc != e) {
+
+    •
+
+    • }
+    • }
+    • return str;
+    +}
+}
```

+/

◦ call-seq:

◦ str.scrub -> new_str

◦ str.scrub(repl) -> new_str

◦ str.scrub[|bytes|] -> new_str

◦ If the string is invalid byte sequence then replace invalid bytes with given replacement

- character, else returns self.
- */
 - +VALUE
 - +rb_str_scrub(int argc, VALUE *argv, VALUE str)
 - +{
- int cr = ENC_CODERANGE(str);
- rb_encoding *enc;
- VALUE repl;
- if (cr == ENC_CODERANGE_7BIT || cr == ENC_CODERANGE_VALID)
- return rb_str_dup(str);
- enc = STR_ENC_GET(str);
- rb_scan_args(argc, argv, "01", &repl);
- if (argc != 0) {
- repl = str_compat_and_valid(repl, enc);
- }
- if (rb_enc_dummy_p(enc)) {
- return rb_str_dup(str);
- }
- if (rb_enc_asciicompat(enc)) {
- const char *p = RSTRING_PTR(str);
- const char *e = RSTRING_END(str);
- const char *p1 = p;
- const char *rep;
- long replen;
- int rep7bit_p;
- VALUE buf = rb_str_buf_new(RSTRING_LEN(str));
- if (rb_block_given_p()) {
-
- }
- else if (!NIL_P(repl)) {
-
-
-
- }
- else if (enc == rb_utf8_encoding()) {
-
-
-
- }
- else {
-
-
-
- }
- cr = ENC_CODERANGE_7BIT;
- p = search_nonascii(p, e);
- if (!p) {
-

- }
- while (p < e) {

-
- | Response | Percentage |
|---------------------|------------|
| Current government | 85% |
| Previous government | 15% |

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- }
- if (p1 < p) {
- [REDACTED]
- }
- if (p < e) {
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- }
- ENCODING_CODERANGE_SET(buf, rb_enc_to_index(enc), cr);
- return buf;
- }
- else {
- /* ASCII incompatible */
- const char *p = RSTRING_PTR(str);
- const char *e = RSTRING_END(str);
- const char *p1 = p;

- VALUE buf = rb_str_buf_new(RSTRING_LEN(str));
- const char *rep;
- long replen;
- long mbminlen = rb_enc_mbminlen(enc);
- static rb_encoding *utf16be;
- static rb_encoding *utf16le;
- static rb_encoding *utf32be;
- static rb_encoding *utf32le;
- if (!utf16be) {

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

- }
- if (!NIL_P(repl)) {

- [REDACTED]
- [REDACTED]

- }
- else if (enc == utf16be) {

- [REDACTED]
- [REDACTED]

- }
- else if (enc == utf16le) {

- [REDACTED]
- [REDACTED]

- }
- else if (enc == utf32be) {

- [REDACTED]
- [REDACTED]

- }
- else if (enc == utf32le) {

- [REDACTED]
- [REDACTED]

- }
- else {

- [REDACTED]
- [REDACTED]

- }

- while (p < e) {

- [REDACTED]
- [REDACTED]

- [illegible]

- [REDACTED]
- }
- if (p < e) {
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- }
- ENCODING_CODERANGE_SET(buf, rb_enc_to_index(enc), ENC_CODERANGE_VALID);
- return buf;
- }
- +}

/******

- Document-class: Symbol

```
@@ -8222,6 +8487,7 @@ Init_String(void)
rb_define_method(rb_cString, "getbyte", rb_str_getbyte, 1);
rb_define_method(rb_cString, "setbyte", rb_str_setbyte, 2);
rb_define_method(rb_cString, "byteslice", rb_str_byteslice, -1);
```

- rb_define_method(rb_cString, "scrub", rb_str_scrub, -1);

```
rb_define_method(rb_cString, "to_i", rb_str_to_i, -1);
rb_define_method(rb_cString, "to_f", rb_str_to_f, 0);
diff --git a/test/dl/test_callback.rb b/test/dl/test_callback.rb
index 8ae652b..ef24235 100644
--- a/test/dl/test_callback.rb
+++ b/test/dl/test_callback.rb
@@ -61,9 +61,11 @@ module DL
  addr = set_callback(TYPE_VOID, 1) do |foo|
    called = true
  end
```

- [REDACTED]

```
func = CFunc.new(addr, TYPE_VOID, 'test')
f = Function.new(func, [TYPE_VOIDP])
- [REDACTED]


```
f.call(nil)

assert called, 'function should be called'
```


```

```
diff --git a/test/ruby/test_m17n.rb b/test/ruby/test_m17n.rb
index a8d56a4..60834bb 100644
--- a/test/ruby/test_m17n.rb
+++ b/test/ruby/test_m17n.rb
@@ -1489,4 +1489,38 @@ class TestM17N < Test::Unit::TestCase
  s.untrust
  assert_equal(true, s.b.untrusted?)
end
+
```

- def test_scrub
- assert_equal("\uFFFF\uFFFF\uFFFF", u("\x80\x80\x80").scrub)

- `assert_equal("\uFFFDA", u("\xF4\x80\x80A").scrub)`

• examples in Unicode 6.1.0 D93b

- `assert_equal("\x41\uFFFD\uFFFD\x41\uFFFD\x41",`

- `assert_equal("\x41\uFFFD\uFFFD\uFFFD\x41",`

- `assert_equal("\x41\uFFFD\uFFFD\uFFFD\x41",`

- `assert_equal("\u0061\uFFFD\uFFFD\uFFFD\u0062\uFFFD\u0063\uFFFD\uFFFD\u0064",`

- `assert_equal("abcdefghijklmnopqrstuvwxyz\u0061\uFFFD\uFFFD\uFFFD\u0062\uFFFD\u0063\uFFFD\uFFFD\u0064",`

- `assert_equal("abcdefghijklmnopqrstuvwxyz\u0061\uFFFD\uFFFD\uFFFD\u0062\uFFFD\u0063\uFFFD\uFFFD\u0064",`

- `assert_equal("abcdefghijklmnopqrstuvwxyz\u0061\uFFFD\uFFFD\uFFFD\u0062\uFFFD\u0063\uFFFD\uFFFD\u0064",`

- `assert_equal("\u3042\u3013", u("\xE3\x81\x82\xE3\x81").scrub("\u3013"))`

- `assert_raise(Encoding::CompatibilityError){ u("\xE3\x81\x82\xE3\x81").scrub(e("\xA4\xA2")) }`
- `assert_raise(TypeError){ u("\xE3\x81\x82\xE3\x81").scrub(1) }`
- `assert_raise(ArgumentError){ u("\xE3\x81\x82\xE3\x81\x82\xE3\x81").scrub(u("\x81")) }`
- `assert_equal(e("\xA4\xA2\xA2\xAE"), e("\xA4\xA2\xA4").scrub(e("\xA2\xAE")))`

- `assert_equal("\u3042", u("\xE3\x81\x82\xE3\x81").scrub{|x|<'<'+x.unpack('H*')[0]+'>'})`
- `assert_raise(Encoding::CompatibilityError){ u("\xE3\x81\x82\xE3\x81").scrub{e("\xA4\xA2")) }`
- `assert_raise(TypeError){ u("\xE3\x81\x82\xE3\x81").scrub{1} }`
- `assert_raise(ArgumentError){ u("\xE3\x81\x82\xE3\x81\x82\xE3\x81").scrub{u("\x81")) }`
- `assert_equal(e("\xA4\xA2\xA2\xAE"), e("\xA4\xA2\xA4").scrub{e("\xA2\xAE"))}`

- `assert_equal("\uFFFD\u3042".encode("UTF-16BE"),`

- `assert_equal("\uFFFD\u3042".encode("UTF-16BE"),`

- `assert_equal("\uFFFD\u3042".encode("UTF-16LE"),`

- `assert_equal("\uFFFD\u3042".encode("UTF-16LE"),`

- `assert_equal("\uFFFD\u3042".encode("UTF-16LE"),`

- `assert_equal("\uFFFD\u3042".encode("UTF-16LE"),`

- `end`
- `end`
- `diff --git a/transcode.c b/transcode.c`
- `index de12c04..9c940ed 100644`
- `--- a/transcode.c`
- `+++ b/transcode.c`
- `@@ -2652,6 +2652,8 @@ str_transcode_enc_args(VALUE str, volatile VALUE *arg1, volatile VALUE *arg2,`
- `return dencidx;`
- `}`

+VALUE rb_str_scrub(int argc, VALUE *argv, VALUE str);

```
+
static int
str_transcode0(int argc, VALUE *argv, VALUE *self, int ecflags, VALUE ecopts)
{
  @@ -2686,6 +2688,17 @@ str_transcode0(int argc, VALUE *argv, VALUE *self, int ecflags, VALUE ecopts)
  ECONV_XML_ATTR_CONTENT_DECORATOR|
  ECONV_XML_ATTR_QUOTE_DECORATOR)) == 0) {
    if (senc && senc == denc) {
```

- `assert_equal("\uFFFD\u3042".encode("UTF-16LE"),`

- [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
 - [REDACTED]
- ```

return NIL_P(arg2) ? -1 : dencidx;
}
if (senc && denc && rb_enc_asciicompat(senc) && rb_enc_asciicompat(denc)) {

```

## #7 - 04/10/2013 08:04 PM - naruse (Yui NARUSE)

パッチ [\[ruby-dev:47241\]](https://github.com/ruby-dev/ruby-dev/pull/47241) による patch による文字列エンコーディングの修正

= String#encode

== String#encode  
String#encode API による API による文字列エンコーディング

== String

- API による String#encode による文字列エンコーディング
- String#scrub による

== API

API による String#encode による文字列エンコーディング String#scrub による

== String#encode

API による iconv による API による文字列エンコーディング  
String#encode による CSI による Ruby による文字列エンコーディング  
UTF-8 による

str.encode("UTF-8", invalid: :replace)

String#encode

str.encode("UTF-8", invalid: :replace, replace: "\u3013")

String#encode  
String#encode fallback による  
Ruby M17N による CSI による fallback による  
String#encode による UTF-8 による

- invalid による undef による
- from encoding (による)
- to encoding (による)
- による
- による

== String#scrub

String#scrub API による  
String#scrub fallback による API による  
zpool scrub による

<http://docs.oracle.com/cd/E19253-01/819-6260/gbbxi/>

```
=== str.scrub
```

Unicode `U+FFFD` (Replacement Character) `?`

```
=== str.scrub("")
```

```
=== str.scrub{|x| '<'+x.ord.to_s(16)+'>' }
```

#### #8 - 04/19/2013 12:31 PM - matz (Yukihiro Matsumoto)

OK, I agree with enhancement of `String#encoding` and `String#scrub`.

Matz.

#### #9 - 04/20/2013 05:21 AM - naruse (Yui NARUSE)

- *Status changed from Assigned to Closed*

- *% Done changed from 0 to 100*

This issue was solved with changeset [r40391](#).

Yui, thank you for reporting this issue.

Your contribution to Ruby is greatly appreciated.

May Ruby be with you.

---

Add example for `String#scrub`

[Feature [#6321](#)] [Feature [#6752](#)] [Bug [#7967](#)]