Ruby 1.8 - Bug #4493

Patch: MRI 1.8.7: syck: fix buffer overflow when parsing YAML from a String.

03/11/2011 02:39 AM - kstephens (Kurt Stephens)

Status:	Open
Priority:	Normal
Assignee:	
Target version:	Ruby 1.8.7
ruby -v:	ruby 1.8.7 (2011-02-18 patchlevel 334) [i686-linux], MBARI 0x8770, Ruby Enterprise Edition 2011.03

Description

=begin
Certain sequences of tokens will cause syck.c store a NULL string terminator outside the allocated p->buffer when parsing from a large YAML string, causing memory corruption leading to SEGV faults.

The problem was discovered by completely disabling MRI's GC, by changing gc.c:rb_newobj() to call xalloc() directly and returning immediately in gc.c:garbage_collect() and then running REE under valgrind. REE stack clearing code was also disabled. Problem was also visible by directly instrumenting syck.c with mprotect().

The patch is applicable to REE 1.8 and MRI 1.8.7:

1) Replaces the confusing logic in syck.c:syck_io_str_read() with behavior similar to syck_io_file_read().

2) Enables ASSERT() by default.
3) syck_assert() now takes a string msg.
4) syck_assert() calls rb_raise() instead of calling abort().
5) Removes a bogus ASSERT() that always fails under MRI unit tests.

This patch *does not* fix unterminated quoted strings that would normally raise a parsing error under psych.

See http://code.google.com/p/rubyenterpriseedition/issues/detail?id=66 for patch.

Contact me directly for specific test cases, instrumentation patches, etc.

=end