

# Updating callout\_reset

**Source:** <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2004-01/0102.html>

---

**From:** Mike Silbersack (*silby\_at\_silby.com*)

**Date:** 01/26/04

Date: Sun, 25 Jan 2004 23:29:33 -0600 (CST)

To: arch@freebsd.org

Based on the work I've been doing with the rtc driver and some issues in the tcp stack with high hz values, I'd like to propose a change in the callout\_reset interface. Currently, the timeout value is specified in ticks. This creates two problems:

#1 – At high hz rates, the maximum timeout possible is reduced; we have at least one example of this causing problems in the kernel at present (16-bit sbtimeout values getting reduced from 327 to 32 seconds), there could be others we haven't caught. Many of these are likely to result from integer math errors introduced by authors who hadn't considered larger hz values occurring.

#2 – Using ticks reduces the potential accuracy of wakeups with our default hz setting. For example, if one wishes to sleep for 5ms when hz=100, the only choice is to request one tick, or 10ms. However, if 5ms could be specified, the callout subsystem would be able to schedule more precise wakeups in places where the next timer interval was between 5 and 10ms away.

Case #1 also provides a second justification for a change; any piece of code which requests a timeout must be aware of the system hz; moving to a standard time format would create more straightforward code.

In order to ensure that modules and the like do not break, I would like to propose that we leave callout\_reset defined as is and create a new function, callout\_set (or some other clever name) which takes a struct \*timespec in place of ticks.

Optimizations as suggested in example #2 above would not be implemented right away, but would be possible due to the better interface.

Thoughts?

Mike "Silby" Silbersack

---

freebsd-arch@freebsd.org mailing list

freebsd-arch: Updating callout\_reset

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@freebsd.org"