

Re: [Lit.] Buffer overruns

Source: <http://www.derkeiler.com/Newsgroups/sci.crypt/2004-12/0896.html>

From: D. J. Bernstein (*djb_at_cr.yp.to*)

Date: 12/14/04

Date: Tue, 14 Dec 2004 09:48:23 +0000 (UTC)

Douglas A. Gwyn wrote:

- > *if you want to encapsulate all buffer operations so that sizes are*
- > *tracked and automatically checked, etc. you can do that readily enough*
- > *using C. But that doesn't address the actual problem! Why is*
- > *the algorithm attempting to copy too much data in the first place?*

Please explain exactly what you think the program should do instead.

Perhaps you're trying to say that buffers should be automatically sized to handle the amount of data available. For example, instead of setting aside 4096 bytes for an input line, and then checking whether the line has overflowed 4096 bytes, the program should dynamically allocate the right amount of memory for the line.

But that's missing the point. Dynamic allocations can fail too! There's always some limit: a megabyte, a gigabyte, whatever. You have to check whether memory is available. You have to figure out what to do when memory isn't available.

I don't mean to say that dynamic allocations are bad. I wince when I see a 60-byte line being stored in a 4096-byte buffer; dynamic allocation is much less wasteful. I also wince when I see a legitimate 10000-byte line bumping into a static 4096-byte limit; dynamic allocation generally has much larger limits—and can afford to have them because larger limits don't mean larger waste.

But those are efficiency issues, not security issues. From a security perspective, checked static allocation and checked dynamic allocation are both fine.

(The problem is that there are other approaches that aren't fine—and those other approaches are *_easier_* for a typical programmer than any safe approach. The programming environment encourages errors! Yes, the right things are *_doable_*, but they aren't as *_easy_* as all the usual wrong things.)

—D. J. Bernstein, Associate Professor, Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago