

[[prev in list](#)] [[next in list](#)] [[prev in thread](#)] [[next in thread](#)]

List: [ncurses-bug](#)
 Subject: [Re: SPAM: \[security\]\[ncurses\] infocmp -i stack buffer overflow \(CVE-121\), request CVE](#)
 From: [Thomas Dickey <dickey@invisible-island.net>](#)
 Date: [2025-12-11 11:32:24](#)
 Message-ID: [aTqryOWX08FitVx- \(\)_prl-debianold-64 ! jexium-island ! net](#)
 [Download RAW [message](#) or [body](#)]

On Wed, Dec 10, 2025 at 11:07:55PM +0100, Sven Joachim wrote:

> On 2025-12-10 15:00 -0500, Thomas Dickey wrote:

>

> > On Wed, Dec 10, 2025 at 10:40:49PM +0800, Yixuan Cao wrote:

> >> Hi Thomas and ncurses maintainers,

> >>

> >>

> >> I am reporting a stack buffer overflow in ncurses (infocmp -i /
 > >> analyze_string). This is a private, responsible disclosure. Please advise
 > >> your preferred coordination timeline.

> >

> > this is a publicly-viewable mailing list

> >

> >> Summary

> >> - Affected component: progs/infocmp.c, function analyze_string()
 > >> - Issue: copies unbounded CSI parameter substrings into 4096-byte
 > >> stack buffers (buf2/buf3). On glibc builds _nc_STRNCPY falls back to
 > >> strncpy, so a long SGR parameter list overflows buf2.

> > That does not seem to be correct: _nc_STRNCPY falls back to strncpy
 > > rather than strncpy, so no overflow occurs unless I am missing something.

It doesn't matter, actually: the computed length is larger than the buffer.

Just delete the sentence mentioning glibc.

The issue applies only to buf2. Delete the mention of buf3.

The length is computed with strlen ... "unbounded" is unwarranted.

> >> - Impact: Local attacker controlling a terminfo entry can cause
 > >> denial of service and potentially gain code execution under the
 > >> privileges of the user running infocmp. The attack surface is any
 > >> terminfo path accessible to the attacker (e.g., \$HOME/.terminfo or
 > >> TERMINFO).

that's implausible.

Breakage in a rarely used option for analysis won't affect the other
 99.99999% of users with vim and bash.

The report doesn't mention how this function is used. The manpage says:

```
-i Analyze the initialization (is1, is2, is3), and reset (rs1, rs2,
rs3), strings in the entry, as well as those used for start-
ing/stopping cursor-positioning mode (smcup, rmcup) as well as
starting/stopping keymap mode (smkx, rmkx).
```

That is, it's a specialized function.

Someone can make a script and get people to run it.

But see below.

> >> Affected version

> >> - Reproduced on ncurses 6.4 (official tarball from

> >> <https://invisible-mirror.net/archives/ncurses/ncurses-6.4.tar.gz>). Please

> >> confirm other versions/branches.

You have to remember on these reports, that asan isn't being run
 against the packaged version, which uses compiler options such
 as -fstack-protector-strong to address this sort of issue.

--

Thomas E. Dickey <dickey@invisible-island.net>
<https://invisible-island.net>

[["signature.asc" \(application/pgp-signature\)](#)].

5/5/26, 8:21 PM

'Re: SPAM: [security][ncurses] infocmp -i stack buffer overflow (CVE-121), request CVE' - MARC

[\[prev in list\]](#) [\[next in list\]](#) [\[prev in thread\]](#) [\[next in thread\]](#)

[Configure](#) | [About](#) | [News](#) | [Add a list](#) | Sponsored by [KoreLogic](#)