

Fuzzing Project report December 2015

Hanno Böck

December 2015

1 Bignum library testing

The latest OpenSSL release 1.0.2e contains a fix for a carry propagation bug in the `BN_mod_exp()` function that I reported in August ¹. In some cases the function produced wrong results. The exact security implications are unclear, but it might be possible to attack Diffie Hellman key exchanges with this vulnerability.

This was one result of tests I did with several bignum libraries. Bignum calculations are an important building block of cryptographic software and incorrect results can lead to security issues.

I looked into this because earlier this year a similar bug was found in the `BN_sqr()` function of OpenSSL ². Ralph-Philipp Weinmann was able to show that it would've been possible to find this bug using american fuzzy lop ³.

There are two potential strategies to use fuzz testing on bignum libraries. One can check for inconsistencies inside a library itself, for example by doing calculations that should always produce the same result for different input. Alternatively one can compare the output of different implementations of the same operation. I did several fuzzing tests with OpenSSL, nss, libgcrypt, gmp and nettle.

Appart from the OpenSSL issue I reported a bug in the `gcry_mpi_invmod()` function ⁴ in libgcrypt, a function to calculate the multiplicative inverse modulo an integer. According to Werner Koch this was a known issue, a fix is planned for the next version of libgcrypt. It is unlikely that this bug has security implications.

I also reported two bugs in the elliptic curve multiplication operations in nettle, one affecting the NIST P-256 curve operations on all architectures ⁵ and one affecting the NIST P-384 curve operations on 64 bit systems ⁶. Nettle-Developer Niels Möller discovered an additional issue while fixing the first bug.

¹<https://openssl.org/news/secadv/20151203.txt> https://blog.fuzzing-project.org/31-Fuzzing-Math-miscalculations-in-OpenSSLs-BN_mod_exp-CVE-2015-3193.html

²<https://www.openssl.org/news/secadv/20150108.txt>

³<https://www.youtube.com/watch?v=PqW8-MUu09c>

⁴<https://lists.gnupg.org/pipermail/gcrypt-devel/2015-December/003642.html>

⁵<https://lists.lysator.liu.se/pipermail/nettle-bugs/2015/003028.html>

⁶<https://lists.lysator.liu.se/pipermail/nettle-bugs/2015/003024.html>

I think it's an important takeaway that fuzz testing can be a successful strategy to test mathematical library functions. This may seem unexpected, as fuzzing is usually associated with memory corruption bugs. It also highlights that carry propagation bugs are a common problem in bignum implementations. I expect more similar bugs to be found in the future.

2 Real-world testing with ASAN-based system

In the last report I mentioned already that I am working on a Gentoo system with almost everything built with Address Sanitizer.

In the mean time I was able to create a bootable version of it that I both used to test a virtualized desktop environment and a server. On the server side this uncovered a use after free error in the MySQL Perl bindings (DBD::mysql)⁷, a stack overflow in vnstat⁸, another out of bounds read in courier⁹, a stack overflow in Monit¹⁰, a global out of bounds read in screen¹¹, an invalid memory read due to a variable type error in syslog-ng¹² and one more out of bounds read in bash¹³.

Testing this with some desktop applications uncovered a use after free in Hexchat¹⁴ and three out of bounds bugs in Claws-Mail¹⁵.

It can generally be seen that simple testing of software with Address Sanitizer (even without any fuzzing) can uncover bugs, even in major software packages that are in widespread use on Linux systems.

3 Network fuzzing

Recently Dough Birdwell released a modified version of american fuzzy lop that supports network input¹⁶. The lack of network support was a major drawback of american fuzzy lop in the past. I tested the new code on various applications¹⁷, surprisingly I haven't found anything with it yet.

This indicates that networking code is generally much better tested than file parsing code.

⁷<https://github.com/perl5-dbi/DBD-mysql/pull/45>

⁸<https://github.com/vergo/vnstat/issues/34>

⁹<https://www.mail-archive.com/courier-users@lists.sourceforge.net/msg38073.html>

html

¹⁰<https://bitbucket.org/tildeslash/monit/commits/58c562b9de25981060486b38a205dfcfa2235ef4>

¹¹<https://savannah.gnu.org/bugs/index.php?46401>

¹²<https://github.com/balabit/syslog-ng/issues/780>

¹³<https://lists.gnu.org/archive/html/bug-bash/2015-11/msg00040.html>

¹⁴<https://github.com/hexchat/hexchat/issues/1540>

¹⁵http://www.thewildbeast.co.uk/claws-mail/bugzilla/show_bug.cgi?id=3559

http://www.thewildbeast.co.uk/claws-mail/bugzilla/show_bug.cgi?id=3563 http:

http://www.thewildbeast.co.uk/claws-mail/bugzilla/show_bug.cgi?id=3573

¹⁶<https://github.com/jdbirdwell/afl>

¹⁷<https://blog.fuzzing-project.org/27-Network-fuzzing-with-american-fuzzy-lop.html>

html

4 Public talks

I held a talk on the Fuzzing Project at the BSides Vienna conference ¹⁸. I plan to do a session on Fuzzing during the 32C3 in cooperation with the Free Software Foundation Europe ¹⁹. I also have submitted a talk on the Address Sanitizer work (see above) for FOSDEM.

5 Notable bugs (dpkg, libxml2, libpcrc, mutt)

As usual, I discovered and reported a number of bugs via fuzzing. Notable packages that I discovered bugs via fuzzing were DPKG ²⁰, Libxml2 ²¹, PCRE ²² and Mutt ²³.

One bug I discovered that was interesting, but that wasn't related to Fuzzing, was in old versions of GnuTLS that are still in use by Ubuntu LTS and Debian old stable. They failed to check the first byte of the CBC padding, some tests for the POODLE TLS vulnerability would mark such hosts als vulnerable ²⁴.

Further fuzzing- and asan-related bugs were reported in GNU Indent ²⁵, sed ²⁶, jbig2dec ²⁷, troff ²⁸, fltk ²⁹, nasm ³⁰, flex ³¹, mawk ³², gawk ³³, grep ³⁴ (bug was in gnulib), ZStandard ³⁵.

¹⁸<http://bsidesvienna.at/talks/#18>

¹⁹https://events.ccc.de/congress/2015/wiki/Session:The_Fuzzing_Project

²⁰<https://blog.fuzzing-project.org/30-Stack-overflows-and-out-of-bounds-read-in-dpkg-Debian.html>

²¹<https://blog.fuzzing-project.org/28-Libxml2-Several-out-of-bounds-reads.html>

²²<https://blog.fuzzing-project.org/29-Heap-Overflow-in-PCRE.html>

²³<http://dev.mutt.org/trac/ticket/3776> <http://dev.mutt.org/trac/ticket/3787>

²⁴<https://blog.hboeck.de/archives/877-A-little-POODLE-left-in-GnuTLS-old-versions.html>

html

²⁵<https://savannah.gnu.org/bugs/index.php?46626> <https://savannah.gnu.org/bugs/index.php?46627>

index.php?46627

²⁶<http://debbugs.gnu.org/cgi/bugreport.cgi?bug=22127>

²⁷http://bugs.ghostscript.com/show_bug.cgi?id=696452 http://bugs.ghostscript.com/show_bug.cgi?id=696453

²⁸<https://savannah.gnu.org/bugs/index.php?46401>

²⁹<http://www.fltk.org/str.php?L3263+P0+S-2+CO+IO+E0+V1>

³⁰http://bugzilla.nasm.us/show_bug.cgi?id=3392328 http://bugzilla.nasm.us/show_bug.cgi?id=3392329 http://bugzilla.nasm.us/show_bug.cgi?id=3392330

³¹<https://sourceforge.net/p/flex/bugs/193/> <https://sourceforge.net/p/flex/bugs/194/>

³²<https://github.com/ThomasDickey/original-mawk/issues/36>

³³<http://git.savannah.gnu.org/cgit/gawk.git/commit/?id=611dc46d8b216c22f05604e8df6bee7aa59e5977>

³⁴<https://debbugs.gnu.org/cgi/bugreport.cgi?bug=21513>

³⁵<https://blog.fuzzing-project.org/26-Two-out-of-bounds-reads-in-Zstandard-zstd.html>

html

6 Misc

There is a new kernel fuzzing tool called syzcaller ³⁶. I haven't tried it yet, but it looks very promising and much more advanced than previous similar tools. It currently depends on the development code of gcc and requires a patched kernel.

I did some testing with libfuzzer ³⁷, which is a fuzzing tool with similar methods as american fuzzy lop, but it operates on a function level. It is part of LLVM.

³⁶<https://github.com/google/syzkaller>

³⁷<http://llvm.org/docs/LibFuzzer.html>