

The Fuzzing Project

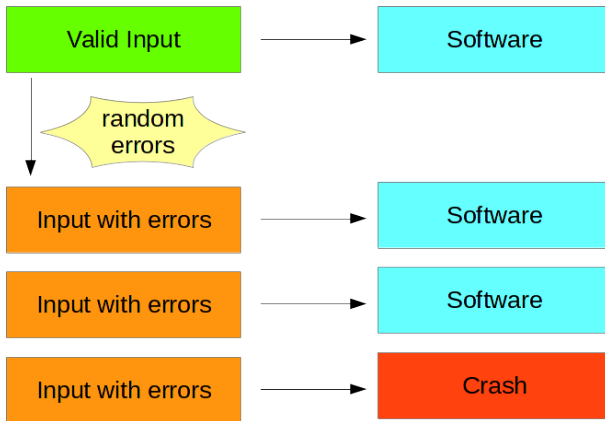
<https://fuzzing-project.org/>

Hanno Böck

Motivation

- Do you use tools like strings, less, file, convert, ldd, unzip, ...?
- Would you use these tools on untrusted input?

Fuzzing



C Memory Bugs

- Buffer Overflow, Stack Overflow, Heap Overflow, Use-after-Free, Out-of-bounds, Memory Corruption, Off-by-1, ...
- Summarize: Software reads or writes the wrong memory
- Many security vulnerabilities are bugs in C memory handling

Invalid memory access example

```
int main() {  
int a[2] = {3, 1};  
int b = a[2];  
}
```

Example: binutils

- October 2014: Michal Zalewski reports a crash in strings
- strings is part of binutils and parses executables (ELF, PE and others) - did you know that?
- Followup: Various people started fuzzing binutils (nm, ld, objdump, readelf, ...) and found hundreds of memory corruption issues - and we're still not done
- binutils 2.25: strings doesn't parse executables by default any more

Example: less

- less pipes input through lesspipe, a script that calls other applications depending on the filetype
- unzip, cpio, lha, antiword, catdoc, unrtf, rpm, msgunfmt, dpkg, identify (ImageMagick), cabextract, readelf (binutils!), isoinfo, ...
- Many of these tools have or had memory corruption bugs that are trivial to find via fuzzing
- less itself has unfixed memory access issues (CVE-2014-9488)

Let's start fuzzing

- Fuzzing finds real security vulnerabilities
- It's easy!
- If you take a random piece of software that parses complex data chances are very high that you will find crashes within minutes
- We should just fuzz everything and fix this

American Fuzzy Lop (afl)

```

american fuzzy lop 0.94b (unrtf)

-- process timing --
  run time : 0 days, 0 hrs, 0 min, 37 sec
  last new path : 0 days, 0 hrs, 0 min, 0 sec
  last uniq crash : 0 days, 0 hrs, 0 min, 21 sec
  last uniq hang : none seen yet
-- cycle progress --
  now processing : 0 (0.00%)
  paths timed out : 0 (0.00%)
-- stage progress --
  now trying : bitflip 2/1
  stage execs : 7406/13.3k (55.57%)
  total execs : 24.2k
  exec speed : 646.5/sec
-- fuzzing strategy yields --
  bit flips : 220/13.3k, 0/0, 0/0
  byte flips : 0/0, 0/0, 0/0
  arithmetics : 0/0, 0/0, 0/0
  known ints : 0/0, 0/0, 0/0
  havoc : 0/0, 0/0
  trim : 4 B/820 (0.24% gain)

-- overall results --
  cycles done : 0
  total paths : 268
  uniq crashes : 1
  uniq hangs : 0
-- map coverage --
  map density : 1360 (2.08%)
  count coverage : 2.62 bits/tuple
-- findings in depth --
  favored paths : 1 (0.37%)
  new edges on : 118 (44.03%)
  total crashes : 5 (1 unique)
  total hangs : 0 (0 unique)
-- path geometry --
  levels : 2
  pending : 268
  pend fav : 1
  own finds : 267
  imported : 0
  variable : 0

[cpu: 29%]
```



American Fuzzy Lop (afl)

- Currently most powerful free tool for fuzzing
- Adds compile time instrumentation and identifies promising code paths
- Developed by Michal Zalewski (lcamtuf), found some of the post Shellshock Bash bugs and issues in gnupg, openssh, libjpeg, libpng, ...



Address Sanitizer (asan)

- Not every invalid memory access causes a crash
- Addressf Sanitizer: Compile time feature to add additional bounds checks (clang, gcc - CFLAGS="-fsanitize=address")
- afl/asan combination is currently the gold standard of fuzzing



Make fuzzing part of development

- Ideally free software projects should integrate fuzzing into their development process
- Make software fuzzing friendly!
- Should not break with Address Sanitizer
- Provide simple command line tools with libraries to expose parsers

Deprecate C

- Shouldn't we deprecate C and rewrite everything in [some other programming language]?
- Answer: Moving away from C is good for new projects
- Projects like miTLS, Servo (browser engine), MirageOS are valuable
- But: We won't deprecate C any time soon



Fix C

- Shouldn't we use mitigations like ASLR because we can't fix all buffer overflows?
- Answer: Yes! Unfortunately state right now is sad. Most Linux distributions don't enable position independent executables by default and have weak ASLR. Better exploit mitigations (Levee) are coming.
- Exploit mitigations are either incomplete or too expensive for real applications - fixing bugs still reduces attack surface
- <http://oss-security.openwall.org/wiki/exploit-mitigation>



Not everything is bad!

- In most cases upstream developers were happy about reports and fixed them quickly, many start fuzzing themselves
- Many people right now flood upstream devs with fuzzing-related bug reports
- Some projects that didn't have releases for a long time were revived (unrtf, cabextract)
- bc/dc had last stable release in 200x, will soon have a new release with fixes for fuzzing-related bugs

Hall of shame

- less: developers didn't answer, new releases didn't fix reported issues
- poppler: several unfixed open bugs, no visible activity on them
- unzip: Public forum has information about memory corruption issues posted several years ago, unfixed in current release
- Dead projects are a problem (no development but active use - e. g. procmail)

The Fuzzing Project

- Tutorial for beginners (Fuzzing is easy!)
- Software list, the good and the bad
- File samples (if you want to fuzz a Microsoft Works importer and don't have an input sample at hand)

Takeaway messages

- Fuzzing is easy - everyone involved in software development should use it
- We have powerful free software tools (american fuzzy lop, address sanitizer)
- If your software is listed on the Fuzzing Project webpage and has no green "OK" - do something about it!

`https://fuzzing-project.org/`

`http://lcamtuf.coredump.cx/afl/`

`https://code.google.com/p/address-sanitizer/`

