

Fuzzing Low-Level Code



Mathias Payer <mathias.payer@epfl.ch>

<https://hexhive.github.io>

HexHive is hiring!



European Research Council

Established by the European Commission

EPFL

Challenge: software complexity

Google Chrome: 76 MLoC


Gnome: 9 MLoC

Xorg: 1 MLoC


glibc: 2 MLoC

Linux kernel: 17 MLoC

Chrome and OS
~100 mLoC,
27 lines/page,
0.1mm/page \approx 370m



Margaret Hamilton
with code for Apollo
Guidance Computer
(NASA, '69)



Brian Kernighan holding
Lion's commentary on
BSD 6 (Bell Labs, '77)

Defense: Testing *OR* Mitigating?



Software Testing

```
vuIn("AAA");
```

```
vuIn("ABC");
```

```
vuIn("AAAABBBB");
```

```
strcpy_chk(buf, 4, str);
```




Mitigations

```
C/C++  
void log(int a) {  
    printf("A: %d", a);  
}  
  
void vuIn(char *str) {  
    char *buf[4];  
    void (*fun)(int) = &log;  
    strcpy(buf, str);  
  
    fun(15);  
}  
  
CHECK(fun, tgtSet);
```

Status of deployed defenses

- Data Execution Prevention (DEP)
- Address Space Layout Randomization (ASLR)
- Stack canaries
- Safe exception handlers
- Control-Flow Integrity (CFI):
Guard indirect control-flow



Assessing exploitability



Which crash to focus on first?

```
american fuzzy lop 2.32b (test_decode_bmp)

process timing |-----| overall results
  run time : 0 days, 1 hrs, 53 min, 36 sec | cycles done : 2
  last new path : 0 days, 0 hrs, 0 min, 35 sec | total paths : 939
  last uniq crash : 0 days, 0 hrs, 6 min, 18 sec | uniq crashes : 124
  last uniq hang : 0 days, 0 hrs, 16 min, 41 sec | uniq hangs : 128
-----|-----|-----
cycle progress |-----| map coverage
now processing : 120* (12.78%) | map density : 0.23% / 1.45%
paths timed out : 0 (0.00%) | count coverage : 4.75 bits/tuple
-----|-----|-----
stage progress |-----| findings in depth
now trying : bitflip 1/1 | favored paths : 126 (13.42%)
stage execs : 923/5152 (17.92%) | new edges on : 185 (19.70%)
total execs : 11.2M | total crashes : 7089 (124 unique)
exec speed : 3487/sec | total hangs : 68.3k (128 unique)
-----|-----|-----
fuzzing strategy yields |-----| path geometry
bit flips : 291/1.27M, 56/1.26M, 22/1.26M | levels : 11
byte flips : 7/158k, 16/29.9k, 23/30.3k | pending : 644
arithmetics : 100/1.66M, 8/1.59M, 99/1.18M | pend fav : 0
known ints : 4/93.8k, 22/395k, 61/768k | own finds : 938
dictionary : 0/0, 0/0, 0/0 | imported : n/a
havoc : 353/1.43M, 0/0 | stability : 100.00%
trim : 19.01%/76.7k, 80.64%
-----|-----|-----
^C |-----| [cpu000: 50%]
```

Residual Attack Surface Probing


- State-of-the-art mitigations complicate attacks
 - Mitigations have limitations but these are hard to assess and explore systematically (and globally)
- Let's infer the *Residual Attack Surface*
 - Given a crash/bug what can an adversary still do?
 - Residual attack surface depends on program, environment, and input

Approach in a nutshell

- Given: crash that results in arbitrary write
- Assume: mitigations make exploitation hard
- Perform ***Code Reuse*** using ***Data-Only Attack***
 - Leverage memory corruption to corrupt state
 - Build Turing-complete payloads as execution traces
 - Express execution traces as memory writes

BOP Gadget: basic block sequence

- **Functional:**
compute (rax = 7)
- **Dispatcher:**
connect
functional blocks
- **Clobbering:**
destroy context



***SPL
payload***

```
graph TD; A[SPL payload] --> B[Selecting functional blocks]; B --> C[Searching for dispatcher blocks]; C --> D[Stitching BOP gadgets];
```

The diagram is a flowchart with four rounded rectangular boxes. The first box, 'SPL payload', is highlighted with a yellow border. An arrow points from it to the second box, 'Selecting functional blocks'. A second arrow points from the second box down to the third box, 'Searching for dispatcher blocks'. A final arrow points from the third box left to the fourth box, 'Stitching BOP gadgets'.

Selecting
functional blocks

Stitching
BOP gadgets

Searching for
dispatcher blocks

SPL payload

- Payload language
- Subset of C
- Library Calls
- Abstract registers as volatile vars

```
void payload() {  
    string prog = "/bin/sh\0";  
    int64* argv = {&prog, 0x0};  
  
    __r0 = &prog;  
    __r1 = &argv;  
    __r2 = 0;  
  
    execve(__r0, __r1, __r2);  
}
```

***SPL
payload***

```
graph TD; A[SPL payload] --> B[Selecting functional blocks]; B --> C[Searching for dispatcher blocks]; C --> D[Stitching BOP gadgets];
```

***Selecting
functional blocks***

**Stitching
BOP gadgets**

**Searching for
dispatcher blocks**

Functional block selection

- Find set of candidate blocks for SPL statement
- Candidate blocks “*could be*” functional blocks as they execute the correct computation
- What about other side effects? What about chaining functional blocks?

Functional block selection (example)

```
__r0 = 10;  
__r1 = 20;
```


```
rax = 10
```

```
rdi = 10
```

```
rax = 20
```


```
rcx = 10
```

```
rcx = 30
```



Functional block selection (example)

```
__r0 = 10;  
__r1 = 20;
```



```
rax = 10
```

Clobbering Clobbering

```
rdi = 10
```

Dispatcher Functional

```
rax = 20
```

Functional Functional


```
rcx = 10
```

Functional Dispatcher

```
rcx = 30
```

Clobbering Dispatcher


***SPL
payload***



***Selecting
functional blocks***



**Stitching
BOP gadgets**




***Searching for
dispatcher blocks***

Dispatcher block search

- BOP gadgets are *brittle*
- Side-effects make gadgets hard to chain
 - Stitching gadgets is NP-hard
 - There is no approximative solution
- Our approach: back tracking and heuristics

BOP gadgets are brittle




Statement #1

Statement #2

Statement #3

Delta Graph: keeping track of blocks

- Squares: Functional blocks for SPL statements
- Nodes: Functional blocks
- Edges: Length of dispatcher chain
- Goal: Select one “node” from each layer (yellow)



***SPL
payload***

```
graph TD; A[SPL payload] --> B[Selecting functional blocks]; B --> C[Searching for dispatcher blocks]; C --> D[Stitching BOP gadgets];
```

***Selecting
functional blocks***

***Stitching
BOP gadgets***

***Searching for
dispatcher blocks***

Stitching BOP gadgets

- Each path is a candidate exploit
- Check and validate constraints along paths
 - Goal: find a valid configuration
 - Constraints come from environment, SPL program, or execution context
 - Verify using concolic execution & constraint solving

Payload synthesis

Program	SPL payload												
	<i>regset4</i>	<i>regref4</i>	<i>regset5</i>	<i>regref5</i>	<i>regmod</i>	<i>memrd</i>	<i>memwr</i>	<i>print</i>	<i>execve</i>	<i>abloop</i>	<i>inloop</i>	<i>ifelse</i>	<i>loop</i>
ProFTPd	✓	✓	✓	✓	✓	✓	✓	✓ 32	X ₁	✓ 128+	✓ ∞	✓	✓ 3
nginx	✓	✓	✓	✓	✓	✓	✓	X ₄	✓	✓ 128+	✓ ∞	✓	✓ 128
sudo	✓	✓	✓	✓	✓	✓	✓	✓	✓	X ₄	✓ 128+	X ₄	X ₄
orzhttpd	✓	✓	✓	✓	✓	✓	✓	X ₄	X ₁	X ₄	✓ 128+	X ₄	X ₃
wuftdp	✓	✓	✓	✓	✓	✓	✓	✓	X ₁	✓ 128+	✓ 128+	X ₄	X ₃
nullhttpd	✓	✓	✓	✓	✓	✓	X ₃	X ₃	✓	✓ 30	✓ ∞	X ₄	X ₃
opensshd	✓	✓	✓	✓	✓	✓	X ₄	X ₄	X ₄	✓ 512	✓ 128+	✓	✓ 99
wireshark	✓	✓	✓	✓	✓	✓	✓	✓ 4	X ₁	✓ 128+	✓ 7	✓	✓ 8
apache	✓	✓	✓	✓	✓	✓	✓	X ₄	X ₄	✓ ∞	✓ 128+	✓	X ₄
smbclient	✓	✓	✓	✓	✓	✓	✓	✓ 1	X ₁	✓ 1057	✓ 128+	✓	✓ 256

✓ The SPL payload was successfully executed on the target binary

X₁ Not enough candidate blocks

X₂ No valid register/variable mappings


X₃ No valid paths between functional blocks

X₄ Un-satisfiable constraints or solver timeout



Success Rate: 81%

Case study: inf loop on nginx

```
ngx_signal_handler()  
41C765: signals.signo == 0  
40E10F: ngx_time_lock != 0  
41C7B1: ngx_process 3 > 1  
41C9AC: ngx_cycle = $alloc_1  
         $alloc_1 > log = $alloc_2  
         $alloc_2 > log_level <= 5  
  
41CA18: signo == 17  
41CA4B: waitpid() return value != {0, 1}  
41cA50: ngx_last_process == 0  
41CB50: *($stack 0x03C) & 0x7F != 0  
41CB5B: $alloc_2 > log_level <= 1  
41CBE6: *($stack 0x03C + 1) != 2  
41CC48: ngx_accept_mutex_ptr == 0  
41CC5F: ngx_cycle > shared_memory.part.elts = 0  
         __r0 = r14 = 0  
41CC79: ngx_cycle > shared_memory.part.nelts <= 0  
41CC7F: ngx_cycle > shared_memory.part.next == 0
```



Case study: if-else in nginx




BOP summary

- Block Oriented Programming
 - Automates Data-Only attacks
 - SPL: A language to express exploit payloads
 - Concolic execution algorithm stitches BOP gadgets
- We build exploits for 81% of the case studies
- Open source implementation (~14,000 LoC)


Software testing: discover bugs

security



Fuzz testing

- A random testing technique that mutates input to improve test coverage




- State-of-art fuzzers use coverage as feedback to evolutionarily mutate the input

Academic fuzzing research



USBfuzz: explore peripheral space

Virtual Environment



USB fuzz Evaluation

- ~60 new bugs discovered in recent kernels
- 36 memory bugs (UaF / BoF)
- ~12 bugs fixed (with 9 CVEs)
- Bug reporting in progress

Type	Bug Info	#
Memory Bugs (36)	double-free	2
	NULL pointer dereference	8
	general protection	6
	slab-out-of-bounds access	6
	user-after-free access	16
Unexpected state reached (17)	INFO	6
	WARNING	9
	BUG	2

Vulnerable apps

Program	Vulnerability	Nodes	RegSet	RegMod	MemRd	MemWr	Call	Cond	Total
ProFTPD	CVE-2006-5815	27,087	40,143	387	1,592	199	77	3,029	45,427
nginx	CVE-2013-2028	24,169	31,497	1,168	1,522	279	35	3375	37,876
sudo	CVE-2012-0809	3,399	5,162	26	157	18	45	307	5715
orzhttpd	BID 41956	1,345	2,317	9	39	8	11	89	2473
wuftp	CVE-2000-0573	8,899	14,101	62	274	11	94	921	15,463
nullhttpd	CVE-2002-1496	1,488	2,327	77	54	7	19	125	2,609
opensshd	CVE-2001-0144	6,688	8,800	98	214	19	63	558	9,752
wireshark	CVE-2014-2299	74,186	124,053	639	1,736	193	100	4555	131276
apache	CVE-2006-3747	18,790	33,615	212	490	66	127	1,768	36,278
smbclient	CVE-2009-1886	166,081	265,980	1,481	6,791	951	119	28,705	304,027

RegSet: Register Assignment Gadgets
RegMod: Register Modification Gadgets
MemRd: Memory Read Gadgets
MemWr: Memory Write Gadgets
Call: Function/System Call Gadgets
Cond: Conditional Statement Gadgets
Total: Total number of Functional Gadgets

SPL payloads

Payload	Description
<i>regset4</i>	Initialize 4 registers with arbitrary values
<i>regref4</i>	Initialize 4 registers with pointers to arbitrary memory
<i>regset5</i>	Initialize 5 registers with arbitrary values
<i>regref5</i>	Initialize 5 registers with pointers to arbitrary memory
<i>regmod</i>	Initialize a register with an arbitrary value and modify it
<i>memrd</i>	Read from arbitrary memory
<i>memwr</i>	Write to arbitrary memory
<i>print</i>	Display a message to stdout using write
<i>execve</i>	Spawn a shell through execve
<i>abloop</i>	Perform an arbitrarily long bounded loop utilizing regmod
<i>inloop</i>	Perform an infinite loop that sets a register in its body
<i>ifelse</i>	An if-else condition based on a register comparison
<i>loop</i>	Conditional loop with register modification