

RecIPE: Revisiting the Evaluation of Memory Error Defenses


Yuancheng Jiang, Roland H.C. Yap, Zhenkai Liang, Hubert Rosier

AsiaCCS 2022




Memory Errors

- Memory errors usually form critical vulnerabilities in low-level languages



Memory Errors

- Memory errors usually form critical vulnerabilities in low-level languages




- The primary source of vulnerabilities
 - 70% severe security bugs in Chrome

Out of the 58 in-the-wild 0-days in 2021, 67% were memory corruption vulnerabilities.

Memory Errors


- Memory errors usually form critical vulnerabilities in low-level languages



- The primary source of vulnerabilities
 - 70% severe security bugs in Chrome


Out of the 58 in-the-wild 0-days in 2021, 67% were memory corruption vulnerabilities.

- Defenses against memory errors



Memory Errors


- Memory errors usually form critical vulnerabilities in low-level languages



- The primary source of vulnerabilities
 - 70% severe security bugs in Chrome

Out of the 58 in-the-wild 0-days in 2021, 67% were memory corruption vulnerabilities.

- Defenses against memory errors




```
#0 0x10813bc85 in main clang-asan.c:6
#1 0x10813bc85 in main clang-asan.c:6
#2 0x7fffa6c46254 in start (libdyld.dylib+0x5254)
SUMMARY: AddressSanitizer: heap-buffer-overflow clang-asan.c:10 in main
Shadow bytes around the buggy address:
0x1c0400001d90: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001da0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001db0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dc0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dd0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001de0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001df0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001e00: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
```

Address Sanitizer

Memory Errors


- Memory errors usually form critical vulnerabilities in low-level languages



- The primary source of vulnerabilities
 - 70% severe security bugs in Chrome

Out of the 58 in-the-wild 0-days in 2021, 67% were memory corruption vulnerabilities.

- Defenses against memory errors



```
0x60200000ef3d is located 0 bytes to the right of 13-byte region [0x0000000000000000]
located by thread T0 here:
#0 0x10818ebf0 in wrap_malloc (libclang_rt.asan_oss_dynamic.dylib+0x408f9)
#1 0x10813bc85 in main clang-asan.c:6
#2 0x7fffa6c46254 in start (libdyld.dylib+0x5254)


SUMMARY: AddressSanitizer: heap-buffer-overflow clang-asan.c:10 in main
Shadow bytes around the buggy address:
0x1c0400001d90: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001da0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001db0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dc0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dd0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001de0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001df0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001e00: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
```

Address Sanitizer

- Tradeoffs due to constraints from overhead and compatibility.


Memory Errors

- Memory errors usually form critical vulnerabilities in low-level languages



- The primary source of vulnerabilities
 - 70% severe security bugs in Chrome

- Defenses against memory errors



```
0x6020000ef3d is located 0 bytes to the right of 13-byte region [0x00000000,0x00000013)
allocated by thread T0 here:
#0 0x10818ebf0 in wrap_malloc (libclang_rt.asan_osx_dynamic.dylib+0x409f)
#1 0x10813bc85 in main clang-asan.c:6
#2 0x7ffff6c46254 in start (libdyld.dylib+0x5254)

SUMMARY: AddressSanitizer: heap-buffer-overflow clang-asan.c:10 in main
shadow bytes around the buggy address:
0x1c0400001d90: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001da0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001db0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dc0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001dd0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001de0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001df0: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
0x1c0400001e00: fa fa fa fa fa fa fa fa fa fa fa fa fa fa fa
```

Address Sanitizer

- Tradeoffs due to constraints from overhead and compatibility.

Question: How to accurately evaluate the **security effectiveness** of defenses?

Existing Evaluation Methods

Theoretical Validation

Experimental Validation

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Diagram illustrating the components of the equation:

- $P(B|A)$: Probability of B occurring given evidence A has already occurred
- $P(A)$: Probability of A occurring
- $P(A|B)$: Probability of A occurring given evidence B has already occurred
- $P(B)$: Probability of B occurring

Experimental Validation

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Diagram illustrating the formula for conditional probability $P(A|B)$. The formula is $P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$. Annotations with arrows point to each part of the formula:

- $P(B|A)$: Probability of B occurring given evidence A has already occurred
- $P(A)$: Probability of A occurring
- $P(B)$: Probability of B occurring
- $P(A|B)$: Probability of A occurring given evidence B has already occurred

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

☹️ GAP between T. and P.

Experimental Validation

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Probability of B occurring given evidence A has already occurred

Probability of A occurring

Probability of A occurring given evidence B has already occurred

Probability of B occurring

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

☹️ GAP between T. and P.

Experimental Validation

CVEs



Real-world Case Studies

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Probability of B occurring given evidence A has already occurred

Probability of A occurring

Probability of A occurring given evidence B has already occurred

Probability of B occurring

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

☹️ **GAP** between T. and P.

Experimental Validation

CVEs



Real-world Case Studies



Limited Scope
Biased Choice

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Probability of B occurring given evidence A has already occurred

Probability of A occurring

Probability of A occurring given evidence B has already occurred

Probability of B occurring

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

☹️ **GAP** between T. and P.

Experimental Validation

CVEs



Real-world Case Studies



Limited Scope
Biased Choice

Benchmarks



Synthesized test cases

Existing Evaluation Methods

Theoretical Validation

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Probability of B occurring given evidence A has already occurred

Probability of A occurring

Probability of A occurring given evidence B has already occurred

Probability of B occurring

- Showing Probabilities
 - Hash Collision
 - Bypassing ASLR

☹️ **GAP** between T. and P.

Experimental Validation

CVEs



Real-world Case Studies



Limited Scope
Biased Choice

Benchmarks



Synthesized test cases



Higher Coverage
Better Accuracy

Runtime Intrusion Prevention Evaluator(RIPE)

- Runtime Intrusion Prevention Evaluator(RIPE), ACSAC'11
 - RIPE generates hundreds of test cases via five dimensions

Runtime Intrusion Prevention Evaluator(RIPE)

- Runtime Intrusion Prevention Evaluator(RIPE), ACSAC'11
 - RIPE generates hundreds of test cases via five dimensions
- RIPE is widely used in evaluating memory error defenses' effectiveness
 - LBC, COTS CFI, Griffin, Fine-CFI, etc.

Runtime Intrusion Prevention Evaluator(RIPE)


- Runtime Intrusion Prevention Evaluator(RIPE), ACSAC'11
 - RIPE generates hundreds of test cases via five dimensions
 - RIPE is widely used in evaluating memory error defenses' effectiveness
 - LBC, COTS CFI, Griffin, Fine-CFI, etc.
- > "the exploits are all very **similar**; ... the sole purpose of overflowing buffers; ... it makes **strong assumptions** about the compiler and the operating systems.", High system-code security with low overhead, Wagner, SP'15

Runtime Intrusion Prevention Evaluator(RIPE)

- Runtime Intrusion Prevention Evaluator(RIPE), ACSAC'11
 - RIPE generates hundreds of test cases via five dimensions
 - RIPE is widely used in evaluating memory error defenses' effectiveness
 - LBC, COTS CFI, Griffin, Fine-CFI, etc.
- > "the RIPE testbed only considers a **limited number** of buffer overflow vulnerabilities. There could be many other types of vulnerabilities and exploit skills in practice", Finding Cracks in Shields: On the Security of Control Flow Integrity Mechanisms, Yuan, CCS'20

RIPE(2011) Performance in 2021


Effectiveness = Attacks/Exploits Prevented Rate



RIPE(2011) Performance in 2021

Effectiveness = Attacks/Exploits Prevented Rate

☹️ Inaccurate Baseline Result - 51% (many RIPE exploits fail)




RIPE(2011) Performance in 2021

Effectiveness = Attacks/Exploits Prevented Rate

☹️ Inaccurate Baseline Result - 51% (many RIPE exploits fail)

=> Overrated Effectivenesses on defenses - over 90%



RecIPE

An **extensible** and **comprehensive** successor to RIPE

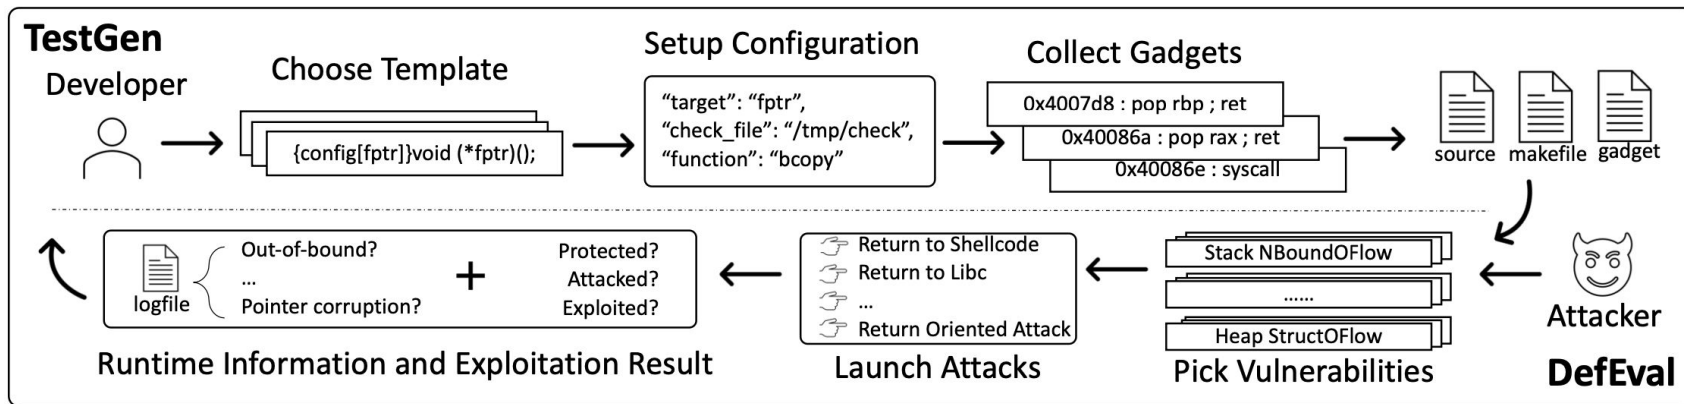
Design Goals

- 😊 Easy to extend and customize
- 😊 Diverse testcases and higher coverage
- 😊 Accurate and comprehensive measurement
- 😊 Support both 32 and 64 bits architecture


Benchmark Overview

Two Components:



- TestGen: Generating Testcases to Individual Folders
- DefEval: Mimicing Attackers; Giving Analysis Results




TestGen




TestGen




TestGen



Configurable Templates => Extensibility & Customizability



TestGen




Configurable Templates => Extensibility & Customizability

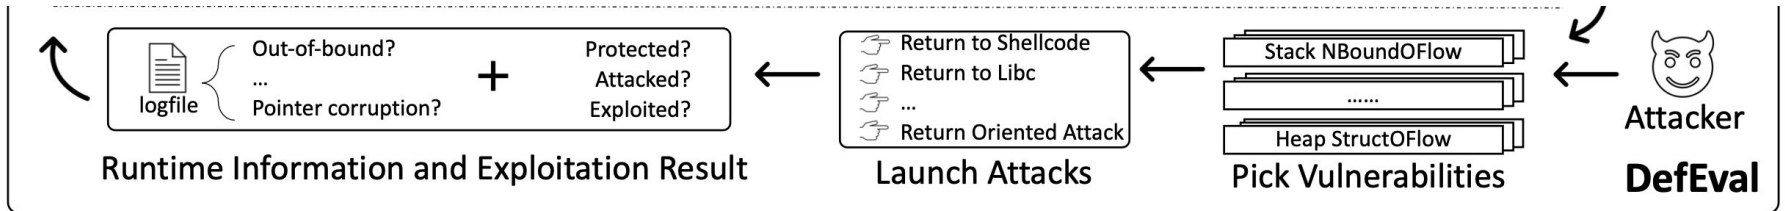


Challenges:

- To be compatible with various attributes
- To switch context under various scenarios



DefEval



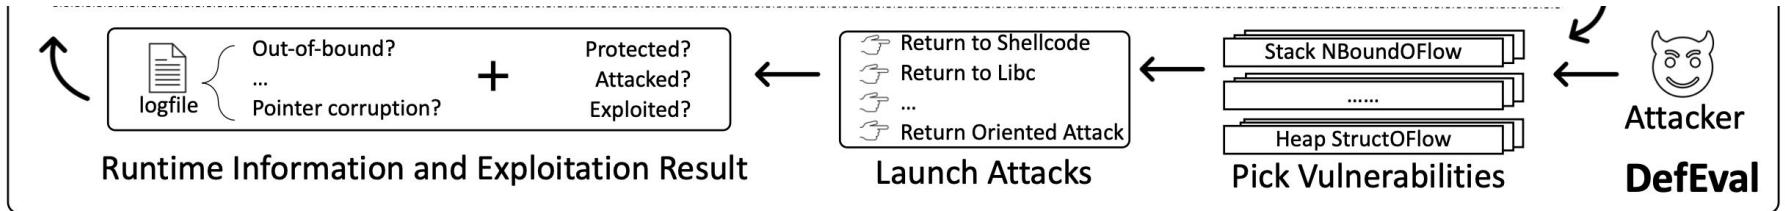
DefEval



Attacker

Equipped with exploits:

- return-to-shellcode
- return-to-libc
- ROP, SROP



DefEval



Attacker


Equipped with exploits:

- return-to-shellcode
- return-to-libc
- ROP, SROP



Analyzer

- Runtime Log => Attacked?
- File Existence => Exploited?



DefEval



Attacker

Equipped with exploits:


- return-to-shellcode
- return-to-libc
- ROP, SROP




Analyzer

- Runtime Log => Attacked?
- File Existence => Exploited?

Attack? + Exploit?
=> Effectiveness Results




Effectiveness Evaluation



Effectiveness Evaluation


😊 Accurate Baseline – No Protection 0%



Effectiveness Evaluation


😊 Accurate Baseline — No Protection 0%

😊 Moderate Score — ASan 50%
=> Reserve space for stronger defense




Effectiveness towards various defenses

- Clearly present pros and cons of each defense




Effectiveness towards various defenses

- Clearly present pros and cons of each defense
- Help to understand limitations:




Effectiveness towards various defenses

- Clearly present pros and cons of each defense
- Help to understand limitations:
 - ASan does not protect non-linear OOB




Effectiveness towards various defenses

- Clearly present pros and cons of each defense
- Help to understand limitations:
 - ASan does not protect non-linear OOB
 - LowFat does not protect Sub-Object Overflow




Some Surprises

- Special Hardenings on Special Functions




Some Surprises

- Special Hardenings on Special Functions




Some Surprises

- Special Hardenings on Special Functions
 - memcpy => "memcpy_safe"




Some Surprises

- Special Hardenings on Special Functions
 - memcpy => "memcpy_safe"
 - homebrew memcpy => get "instrumented"




Some Surprises

- Special Hardenings on Special Functions
 - memcpy => "memcpy_safe"
 - homebrew memcpy => get "instrumented"
 - bcopy => not widely considered



Some Surprises

- Special Hardenings on Special Functions
 - memcpy => "memcpy_safe"
 - homebrew memcpy => get "instrumented"
 - bcopy => not widely considered
 - read => "syscall"



More Detailed Results

Table 1: Evaluating security flags and sanitizers from region and technique attributes

Region		Stack												Heap						Global					
Technique		BoundOfFlow		NBoundOfFlow		SubOfFlow		NBoundOfFlow		OOBPtrHijack		PtrHijack		BoundOfFlow		NBoundOfFlow		SubOfFlow		BoundOfFlow		NBoundOfFlow		SubOfFlow	
Defense	Compiler	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP	ATK	EXP
Security Flags																									
Baseline	gcc	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
FullRELRO	gcc	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
DEP	gcc	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
PIE	gcc	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
FORTIFY	gcc	● ¹	● ¹	● ¹	● ¹	● ²	● ²	● ²	● ²	● ¹	● ¹	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²
	clang	● ¹	● ¹	● ¹	● ¹	● ²	● ²	● ²	● ²	● ¹	● ¹	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²	● ²
StackPtrAll	gcc	● ¹	● ¹	● ¹	● ¹	●	●	●	●	● ¹	● ¹	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	● ¹	● ¹	● ¹	● ¹	●	●	●	●	● ¹	● ¹	●	●	●	●	●	●	●	●	●	●	●	●	●	●
Default	gcc	● ¹	● ¹	● ¹	● ¹	●	●	●	●	● ¹	● ¹	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	clang	● ¹	● ¹	● ¹	● ¹	●	●	●	●	● ¹	● ¹	●	●	●	●	●	●	●	●	●	●	●	●	●	●
Best-Flags	gcc	● ^{1,2}	● ^{1,2}	● ¹	● ¹	●	●	●	●	● ^{1,8}	● ^{1,8}	● ^{1,6,8}	● ^{1,6,8}	●	●	●	●	●	●	●	●	●	●	●	●
	clang	● ^{1,2}	● ^{1,2}	● ^{1*}	● ^{1*}	●	●	●	●	● ^{1,8}	● ^{1,8}	● ^{1,6,8}	● ^{1,6,8}	●	●	●	●	●	●	●	●	●	●	●	●
SafeStack	clang	● ⁵	● ⁵	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
Sanitizers																									
ASan	gcc	● ^{1,2,4}	● ^{1,2,4}	● ^{1,4}	● ^{1,4}	●	●	●	●	● ^{1,2,4}	● ^{1,2,4}	● ^{1,4}	● ^{1,4}	●	●	●	●	●	●	● ^{1,4}	● ^{1,4}	●	●	●	●
	clang	● ^{1,2,4}	● ^{1,2,4}	● ^{1,4}	● ^{1,4}	●	●	●	●	● ^{1,2,4}	● ^{1,2,4}	● ^{1,4}	● ^{1,4}	●	●	●	●	●	●	● ^{1,4}	● ^{1,4}	●	●	●	●
ASanPlus	gcc	● ^{1,2,4,6}	● ^{1,2,4,6}	● ^{1,4}	● ^{1,4}	●	●	●	●	● ^{1,2,4}	● ^{1,2,4}	● ^{1,6,8}	● ^{1,6,8}	●	●	●	●	●	●	● ^{1,4}	● ^{1,4}	●	●	●	●
	clang	● ^{1,2,4,6}	● ^{1,2,4,6}	● ^{1,4}	● ^{1,4}	●	●	●	●	● ^{1,2,4}	● ^{1,2,4}	● ^{1,6,8}	● ^{1,6,8}	●	●	●	●	●	●	● ^{1,4}	● ^{1,4}	●	●	●	●
Valgrind	clang	●	●	●	●	●	●	●	●	● ⁹	● ⁹	● ⁹	● ⁹	●	●	●	●	●	●	●	●	●	●	●	●
SoftBound	clang	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	●	●	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	
LowFat	clang	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	●	●	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	
EffectiveSan	clang	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ¹	● ¹	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	
ESanPlus	clang	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ¹	● ¹	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	● ^{1,3,4}	

Symbol	Meaning
○	All ATK/EXP Fail
◐	Half ATK/EXP Fail
●	All ATK/EXP Succeed
Super	Defense Strategy
○ ¹	Memory Layout Change
○ ^{1*}	CLang Memory Layout
○ ²	Guard/Redzone checks
○ ³	Bound Checks
○ ⁴	Compiler Hardening
○ ⁵	Separate Stack
○ ⁶	Memory Permissions
○ ⁷	Dynamic Instrumentation
○ ⁸	Randomization
○ ⁹	Allocator Change
Sub	Defense Weakness
○ ¹	Missed Check
○ ²	Analysis Limitation
Suffix	Other Reasons
○ [*]	ATK/EXP Not Robust

Summary

- We propose a newly designed benchmark, RecIPE, for evaluating memory error defenses' effectiveness.
- RecIPE is extensible, comprehensive, and accurate.
- RecIPE helps to understand the security coverage of memory error defense and even implementation details.
- RecIPE is available at <https://github.com/YuanchengJiang/recipe-benchmark>

Thanks

