

# Architectural Support for Low Overhead Memory Safety Checks

Mohamed Tarek Ibn Ziad, Miguel Arroyo, Evgeny Manzhosov,  
Ryan Piersma and Simha Sethumadhavan



COMPUTER SCIENCE

Columbia University

06/16/2021



# Memory Safety is a serious problem!

Computing Sep 6

...

## Apple says China's Uighur Muslims were targeted in the recent iPhone hacking campaign

The tech giant gave a rare statement that bristled at Google's analysis of the novel hacking operation.



# Memory Safety is a serious problem!

Computing Sep 6

...

## Apple says China's Uighur Muslims were targeted in the recent iPhone hacking campaign

The tech giant gave a rare statement that bristled at Google's analysis of the novel hacking operation.

EDITOR'S PICK | 42,742 views | Nov 21, 2018, 07:00am

**Exclusive: Saudi Dissidents Hit With Stealth iPhone Spyware Before Khashoggi's Murder**



# Memory Safety is a serious problem!

Computing Sep 6

...

## Apple says China's Uighur Muslims were targeted in the recent iPhone hacking campaign

The tech giant gave a rare statement that bristled at Google's analysis of the novel hacking operation.

*The New York Times*

EDITOR'S PICK | 42,742 views | Nov 21, 2018, 07:00am

*WhatsApp Rushes to Fix Security Flaw Exposed in Hacking of Lawyer's Phone*

**Exclusive: Saudi Dissidents Hit With Stealth iPhone Spyware Before Khashoggi's Murder**



**It's easy to make mistakes**




# It's easy to make mistakes



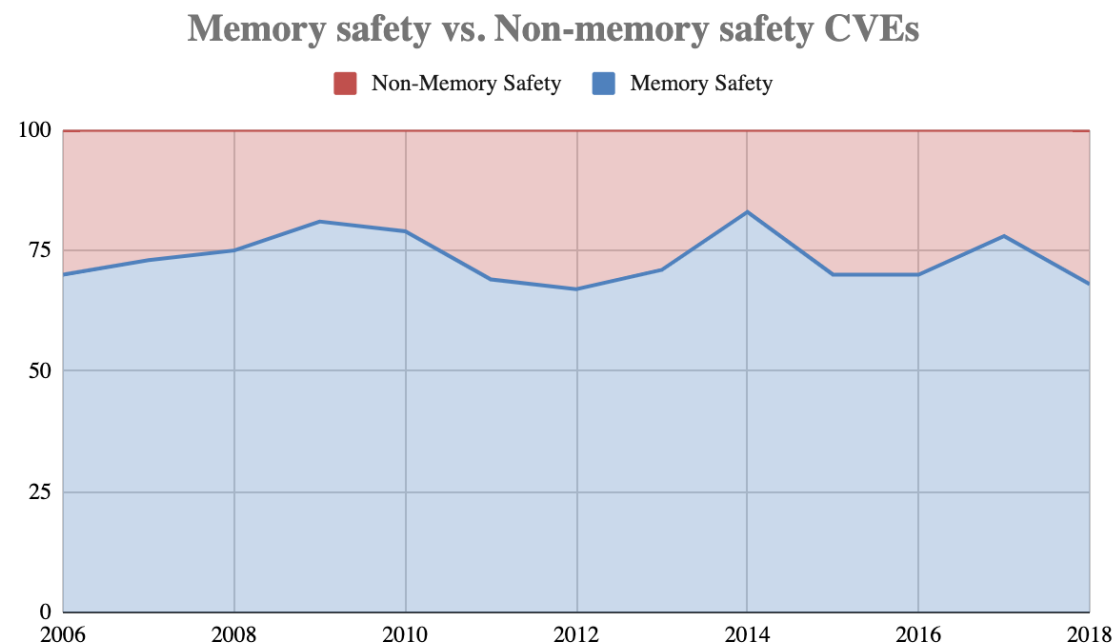
SEGFALT!

# Prevalence of Memory Safety Vulns

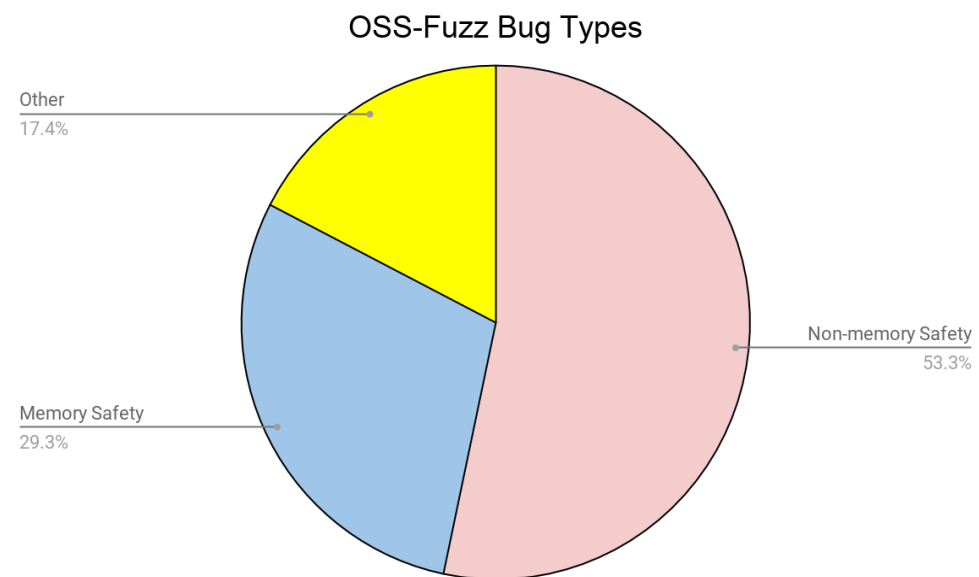


Microsoft Product CVEs

# Prevalence of Memory Safety Vulns



Microsoft Product CVEs



Google OSS-Fuzz bugs from 2016-2018.



TEEN COMIX



BRANSON '15





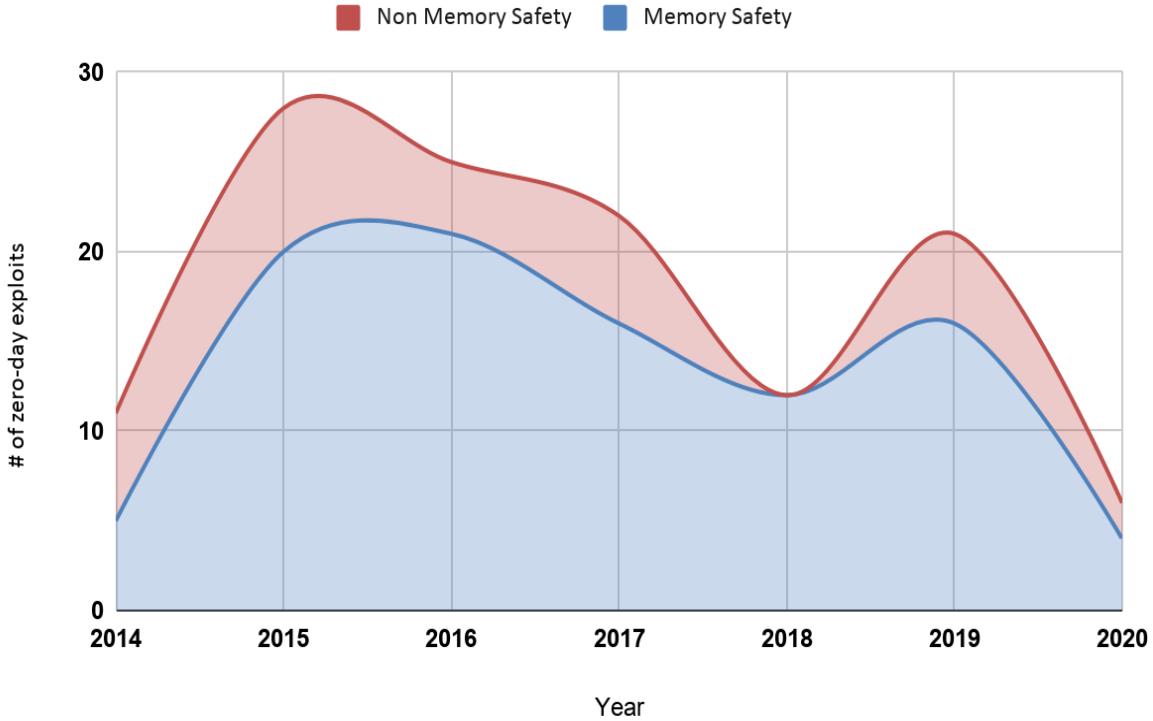
ATTACKERS



MEMORY SAFETY



# Attackers prefer Memory Safety Vulns




Zero-day “in the wild” exploits from 2014-2020

Source: Google Project Zero, Oday "In the Wild" spreadsheet. Last updated: April 2020

**ALL HOPE IS NOT  
LOST!**






**Modern  
software design  
is useful for  
security**




# Modern software design is useful for security



## Increasing adoption of binning allocators

- Maintains memory locality.
- Implicit lookup of allocation information.

# Modern software design is useful for security



## Increasing adoption of binning allocators

- Maintains memory locality.
- Implicit lookup of allocation information.

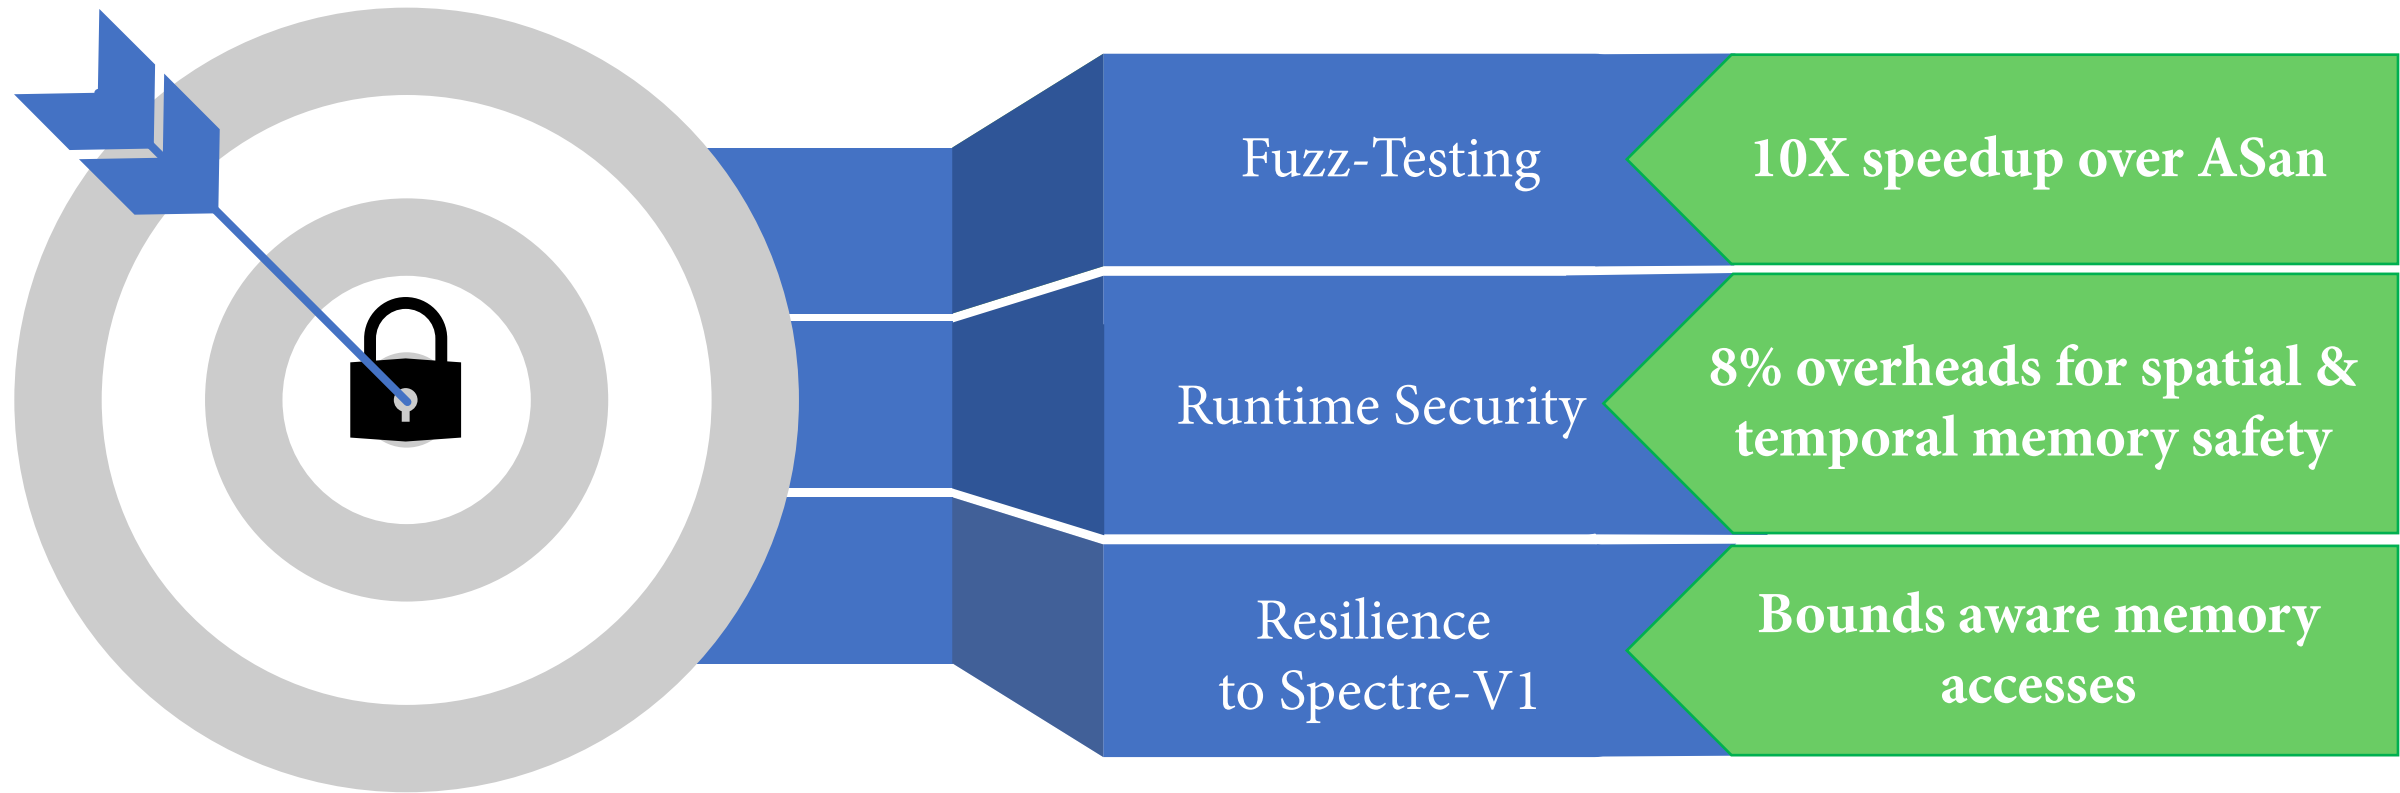


mi-malloc



tcMalloc

# The benefits of No-FAT







# **Binning Memory Allocators**

101



# Binning Memory Allocators

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```



...


*Virtual Memory*



# Binning Memory Allocators



```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```



...

*Virtual Memory*



# Binning Memory Allocators



```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```

Memory is requested by the allocator.



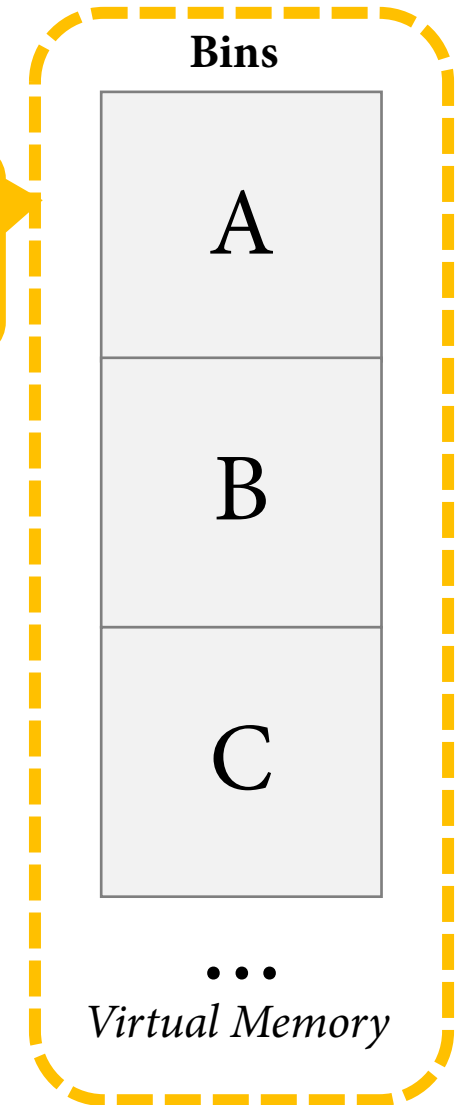
...

*Virtual Memory*

# Binning Memory Allocators

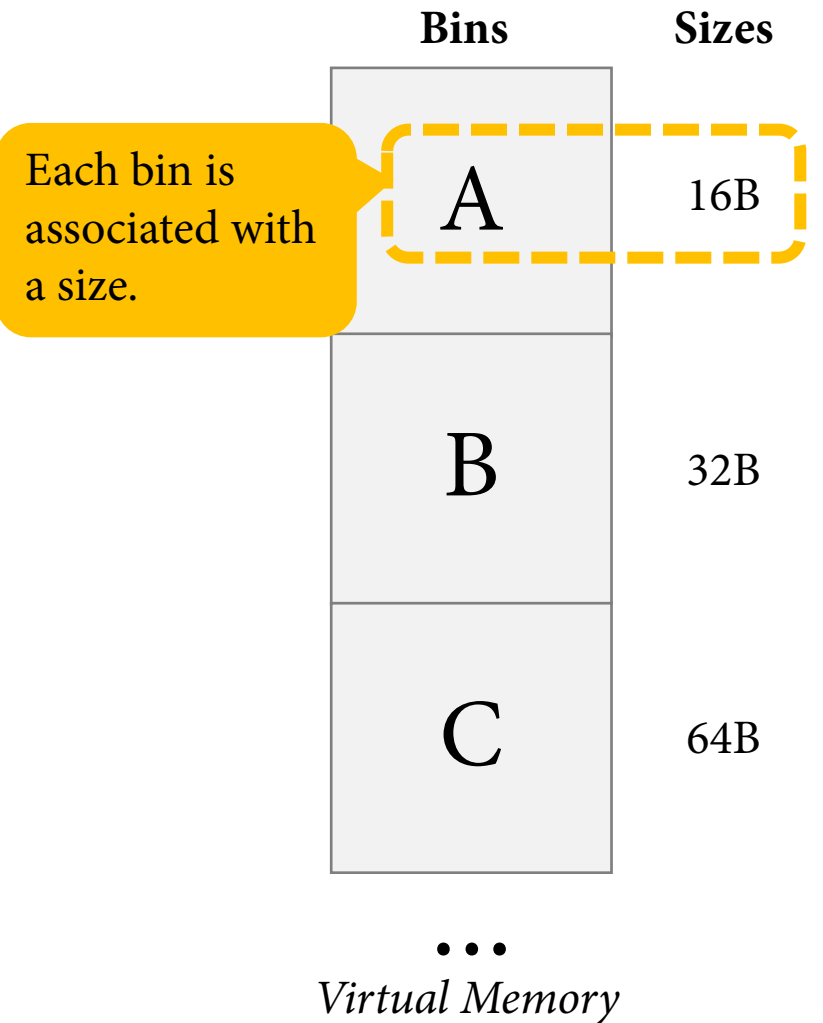
```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```

Memory is divided into bins.



# Binning Memory Allocators

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```






# Binning Memory Allocators




```
40. int main() {  
41.     char* ptr = 12B  
42.     ...  
50. }
```



# Binning Memory Allocators

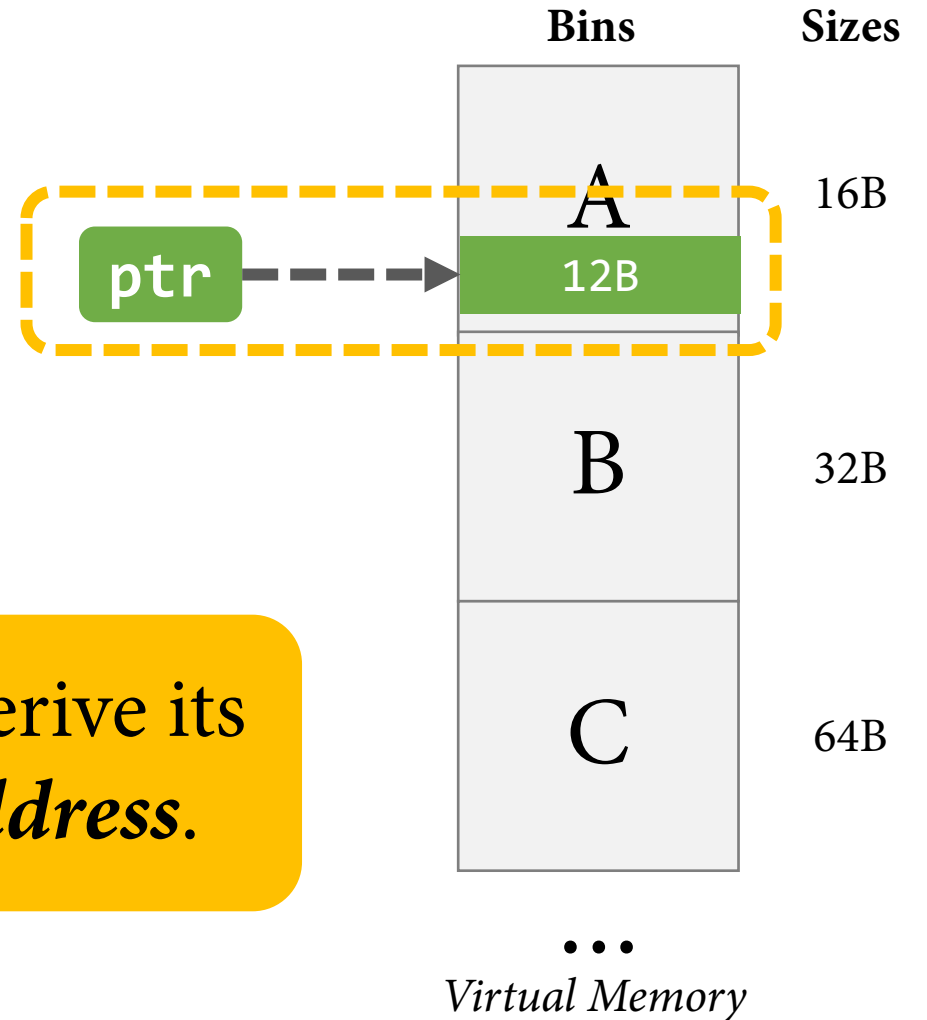
```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```






# Binning Memory Allocators

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ...  
50. }
```



Given **any** pointer, we can derive its *allocation size* and *base address*.




# From Bins to Security



# Spatial Memory Safety (Inter-Object)

The Problem

**Bin A**



Virtual Memory



# Spatial Memory Safety (Inter-Object)

The Problem

Adjacent objects can overflow into each other.

**Bin A**



Virtual Memory



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
50. }
```



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A'; ➡ store ptr[1], 'A'  
43.     ...  
50. }
```

# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A'; → s_store ptr[1], 'A', ptrtrusted base  
43.     ...  
50. }
```

We add one extra operand for loads/stores.



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12); → ptr_trusted_base  
42.     ptr[1] = 'A';           s_store ptr[1], 'A', ptr_trusted_base  
43.     ...  
50. }
```



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
50. }
```

`s_store ptr[1], 'A'` `ptr` trusted base

The compiler propagates the allocation base address.



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';           s_store ptr[1], 'A', ptr_trusted base  
43.     ...  
50. }
```




# Spatial Memory Safety (Inter-Object)

```
s_store ptr[1], 'A', ptr_trusted base
```



# Spatial Memory Safety (Inter-Object)





# Spatial Memory Safety (Inter-Object)

```
s_store ptr[1], 'A', ptr_trusted base
```

$$\text{offset} = \text{ptr}[1] - \text{ptr}_{\text{trusted base}}$$



# Spatial Memory Safety (Inter-Object)

```
s_store ptr[1], 'A', ptr_trusted base
```

**offset** = `ptr[1]` - `ptr_trusted base`

**size** = `getSize(ptr_trusted base)`



# Spatial Memory Safety (Inter-Object)

```
s_store ptr[1], 'A', ptr_trusted_base
```

**offset** = `ptr[1]` - `ptr_trusted_base`

**size** = `getSize(ptr_trusted_base)`

**Bounds Check**

**offset** < **size** ?



# Spatial Memory Safety (Inter-Object)

The allocation size information is made available to the hardware to verify memory accesses.

```
size = getSize(ptrtrusted base)
```

Bounds Check

```
offset < size ?
```





# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);           ptr_trusted_base  
42.     ptr[1] = 'A';                   s_store ptr[1], 'A', ptr_trusted_base  
43.     ...  
50. }
```



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
49.     foo(ptr);  
50. }
```

`ptr_trusted_base`

`s_store ptr[1], 'A', ptr_trusted_base`


Let's pass the pointer to another context (e.g., foo).

# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
49.     foo(ptr);  
50. }  
51. void Foo (char* xptr){  
52.     ...  
53.     xptr[7] = 'B';  
54.     ...  
60. }
```

`ptrtrusted_base`

`s_store ptr[1], 'A', ptrtrusted_base`

An orange arrow originates from the parameter `xptr` in the `Foo` function (line 51) and points back to the `ptr` variable in the `main` function (line 41), illustrating that `xptr` is a pointer to the memory location of `ptr`.



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);           ptr_trusted_base  
42.     ptr[1] = 'A';                    s_store ptr[1], 'A', ptr_trusted_base  
43.     ...  
49.     foo(ptr);  
50. }  
51. void Foo (char* xptr){  
52.     ...  
53.     xptr[7] = 'B'; → s_store xptr[7], 'A', xptr_trusted_base  
54.     ...  
60. }
```

# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
49.     foo(ptr);  
50. }  
51. void Foo (char* xptr){  
52.     ...  
53.     xptr[7] = 'B';  
54.     ...  
60. }
```

`ptrtrusted_base`

`s_store ptr[1], 'A', ptrtrusted_base`

`s_store xptr[7], 'A', xptrtrusted_base`

How do we get this?

# Spatial Memory Safety (Inter-Object)

```
40. int main() {
41.     char* ptr = malloc(12);           ptrtrusted_base
42.     ptr[1] = 'A';                     s_store ptr[1], 'A', ptrtrusted_base
43.     ...
49.     foo(ptr);
50. }
51. void Foo (char* xptr){
52.     ...                               xptrtrusted_base ← compBase(xptr[7])
53.     xptr[7] = 'B';                     s_store xptr[7], 'A', xptrtrusted_base
54.     ...
60. }
```



# Spatial Memory Safety (Inter-Object)

```
xptrtrusted base ← compBase(xptr[7])
```



# Spatial Memory Safety (Inter-Object)

```
xptrtrusted base ← compBase(xptr[7])
```

**Bin** = xptr >> log<sub>2</sub>(S) where S is the size of the bins.





# Spatial Memory Safety (Inter-Object)

```
xptrtrusted base ← compBase(xptr[7])
```

**Bin** = `xptr` >> `log2(S)` where S is the size of the bins.

**size** = `getSize(Bin)`

# Spatial Memory Safety (Inter-Object)

```
xptrtrusted_base ← compBase(xptr[7])
```

**Bin** = `xptr` >> `log2(S)` where S is the size of the bins.

**size** = `getSize(Bin)`

`xptrtrusted_base` = `[xptr × (1 / size)] × size`

# Spatial Memory Safety (Inter-Object)

```
xptrtrusted_base ← compBase(xptr[7])
```

**Bin** = `xptr` >> `log2(S)` where S is the size of the bins.

**size** = `getSize(Bin)`

**xptr<sub>trusted\_base</sub>** = `[xptr × (1 / size)] × size`

Base pointer is **implicitly** derived!



# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ...  
49.     foo(ptr);  
50. }
```

`ptrtrusted_base`

`s_store ptr[1], 'A', ptrtrusted_base`

# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ptr = ptr + 100;  
44.     ...  
49.     foo(ptr);  
50. }
```

`ptrtrusted_base`

`s_store ptr[1], 'A', ptrtrusted_base`

Pointer arithmetic can push the pointer out-of-bounds before calling foo!

# Spatial Memory Safety (Inter-Object)

```
40. int main() {  
41.     char* ptr = malloc(12);  
42.     ptr[1] = 'A';  
43.     ptr = ptr + 100;  
44.     verifyBounds ptr, ptr  
45.     ...  
49.     foo(ptr);  
50. }
```

`ptrtrusted_base`

`s_store ptr[1], 'A', ptrtrusted_base`

`verifyBounds ptr, ptrtrusted_base`

Verify the bounds of all pointers that escape to memory (or another function).



# **Spatial Memory Safety (Intra-Object)**

The Problem



# Spatial Memory Safety (Intra-Object)

## The Problem

```
typedef struct {  
    char a;  
    double b;  
    char c[3];  
    void (*fp)();  
} A_t;
```

Adjacent fields can be overflowed into.



# Spatial Memory Safety (Intra-Object)

```
typedef struct {  
    char a;  
    double b;  
    char c[3];  
    void (*fp)();  
} A_t;
```



```
typedef struct {  
    char a;  
    double b;  
    A_t_c *c_ptr;  
    void (*fp)();  
} A_t;
```



```
typedef struct {  
    char c[3];  
} A_t_c;
```

The **Buf2Ptr** transformation promotes intra-allocation buffers to standalone allocations.




# Temporal Memory Safety



# Temporal Memory Safety


The Problem





# Temporal Memory Safety


## The Problem





# Temporal Memory Safety

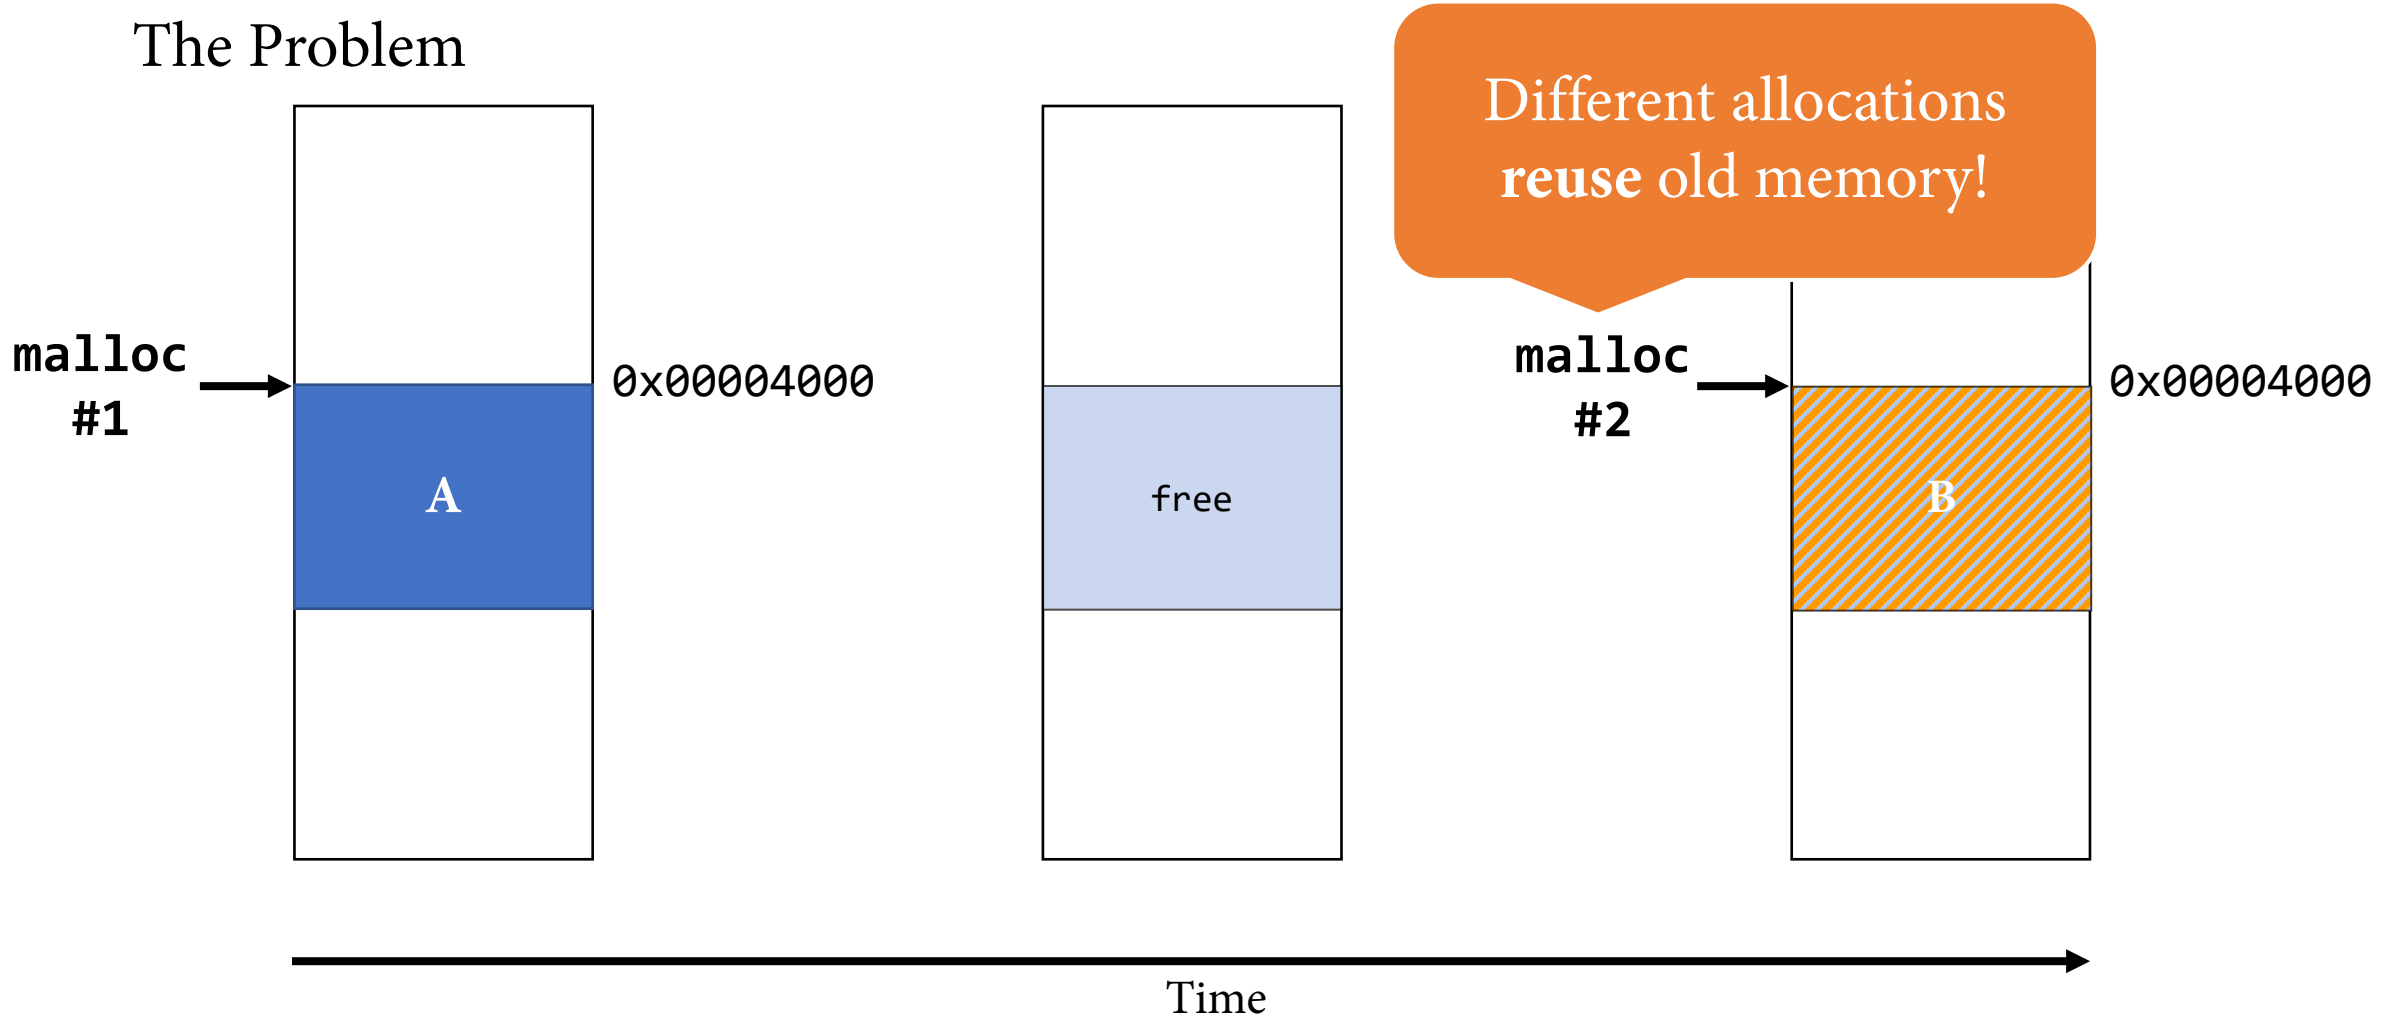
## The Problem





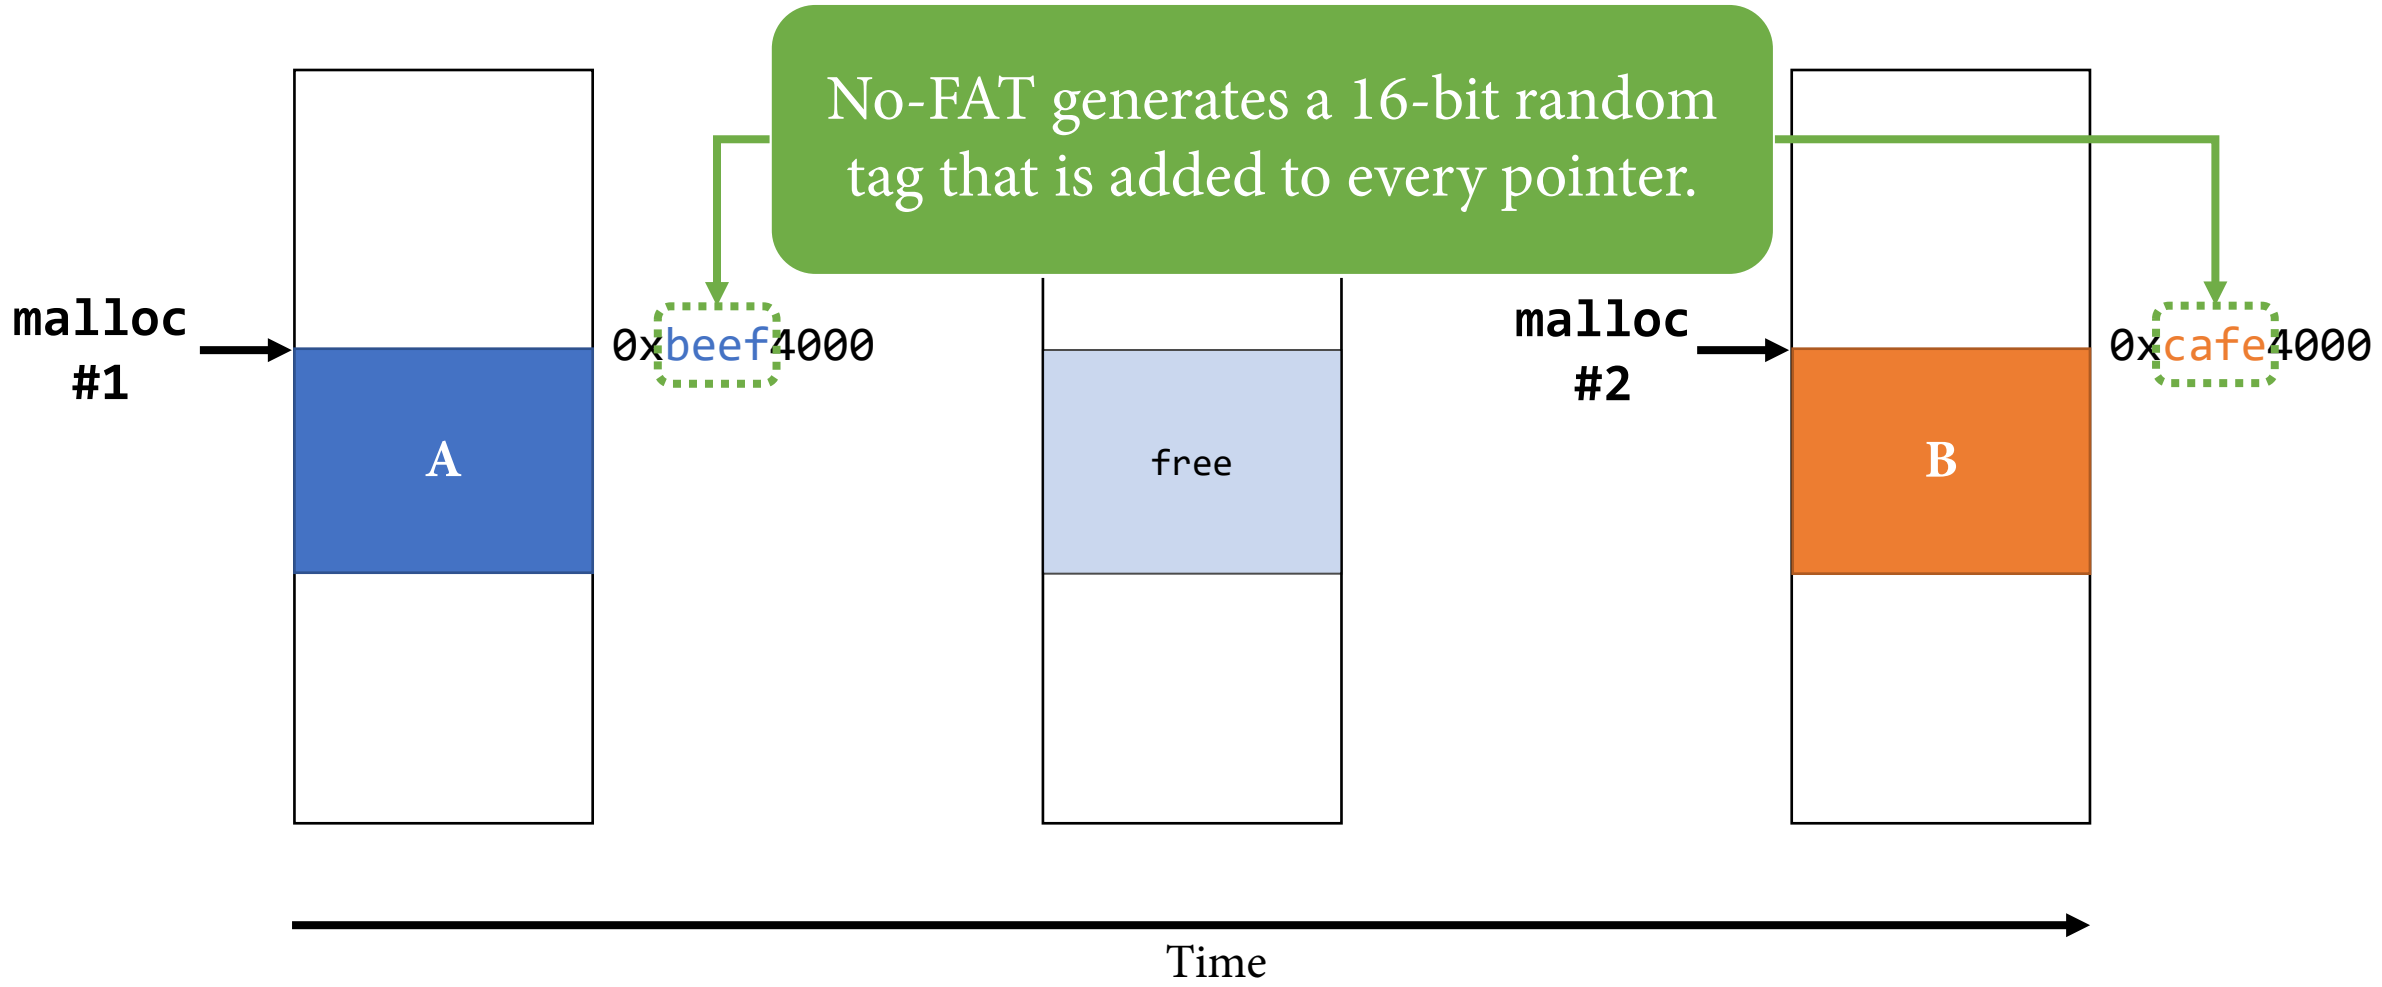
# Temporal Memory Safety


## The Problem





# Temporal Memory Safety





# ISA Extensions





# ISA Extensions

- 1 `s_store Addr, Dest, BaseAddr`
- 2 `s_load Addr, Src, BaseAddr`



# ISA Extensions


- 1 `s_store Addr, Dest, BaseAddr`
- 2 `s_load Addr, Src, BaseAddr`
- 3 `verifyBounds Addr, BaseAddr`

# ISA Extensions

1 `s_store Addr, Dest, BaseAddr`

2 `s_load Addr, Src, BaseAddr`

3 `verifyBounds Addr, BaseAddr`



Exceptions are thrown in the case that the target memory address does not match BaseAddr.

# ISA Extensions

1 `s_store Addr, Dest, BaseAddr`


2 `s_load Addr, Src, BaseAddr`

3 `verifyBounds Addr, BaseAddr`

4 `compBase Addr, Dest`




Exceptions are thrown in the case that the target memory address does not match BaseAddr.



# Microarchitectural Overview




# Microarchitectural Overview




# Microarchitectural Overview

Traditional memory safety techniques frequently access per-pointer or per-allocation metadata.



# Microarchitectural Overview


Traditional memory safety techniques frequently access per-pointer or per-allocation metadata.






# Microarchitectural Overview

**NO** changes to the memory subsystems!






# Microarchitectural Overview




# Microarchitectural Overview







# Microarchitectural Overview





# Microarchitectural Overview





# Resilience to Common Exploits



# Resilience to Common Exploits

**1** Buffer Over-/Under-flows  
Cannot corrupt memory.






# Resilience to Common Exploits

1

Buffer  
Over-/Under-flows  
Cannot corrupt  
memory.



Original:





# Resilience to Common Exploits

**1** Buffer Over-/Under-flows  
Cannot corrupt memory.





# Resilience to Common Exploits

**1** **Buffer Over-/Under-flows**  
Cannot corrupt memory.

**2** **Use-after-free**  
Each allocation instance is tagged randomly.



**Tag is propagated with the allocation base address.**

# Resilience to Common Exploits

1

Buffer  
Over-/Under-flows  
Cannot corrupt  
memory.

2

Use-after-free  
Each allocation  
instance is tagged  
randomly.

3

Spectre-V1

```
// mispredicted branch
if (i < sizeof(a)) {

    secret = a[i];

    // secret is leaked
    val = b[64 * secret];
}
```

# Resilience to Common Exploits

1

**Buffer  
Over-/Under-flows**  
Cannot corrupt  
memory.

2

**Use-after-free**  
Each allocation  
instance is tagged  
randomly.

3

**Spectre-V1**  
Speculative loads are  
aware of the legitimate  
allocation-bounds.

```
// mispredicted branch
if (i < sizeof(a)) {

    secret = a[i];

    // secret is leaked
    val = b[64 * secret];
}
```

# Resilience to Common Exploits

1

**Buffer  
Over-/Under-flows**  
Cannot corrupt  
memory.

2

**Use-after-free**  
Each allocation  
instance is tagged  
randomly.

3

**Spectre-V1**  
Speculative loads are  
aware of the legitimate  
allocation-bounds.

```
// mispredicted branch
if (i < sizeof(a)) {
    secret = a[i];

    // secret is leaked
    val = b[64 * secret];
}
```

- Speculative out-of-bounds loads are not allowed to change the cache state or forward values to dependent instructions.

# Resilience to Common Exploits

1

**Buffer  
Over-/Under-flows**  
Cannot corrupt  
memory.

2

**Use-after-free**  
Each allocation  
instance is tagged  
randomly.


3

**Spectre-V1**  
Speculative loads are  
aware of the legitimate  
allocation-bounds.

```
// mispredicted branch
if (i < sizeof(a)) {
    secret = a[i];
    // secret is leaked
    val = b[64 * secret];
}
```

Detected!


- Speculative out-of-bounds loads are not allowed to change the cache state or forward values to dependent instructions.



**Performance**



# Performance Overheads




## Hardware Modifications

Our measurements show minimal latency/area/power overheads.



# Performance Overheads



## Hardware Modifications


Our measurements show minimal latency/area/power overheads.

00010010  
101001101  
00010010  
111001001  
00010010

## Software Modifications

- Our special loads/stores do not change the binary size.

# Performance Overheads



## Hardware Modifications


Our measurements show minimal latency/area/power overheads.

00010010  
101001101  
00010010  
111001001  
00010010

## Software Modifications

- Our special loads/stores do not change the binary size.
- We verify pointer bounds before storing them to memory.

# Performance Overheads



## Hardware Modifications

Our measurements show minimal latency/area/power overheads.


```
00010010
101001101
00010010
111001001
00010010
```

## Software Modifications

- Our special loads/stores do not change the binary size.
- We verify pointer bounds before storing them to memory.
- We compute the allocation base address of arbitrary pointers when they are loaded from memory.



# Performance Results (x86\_64)



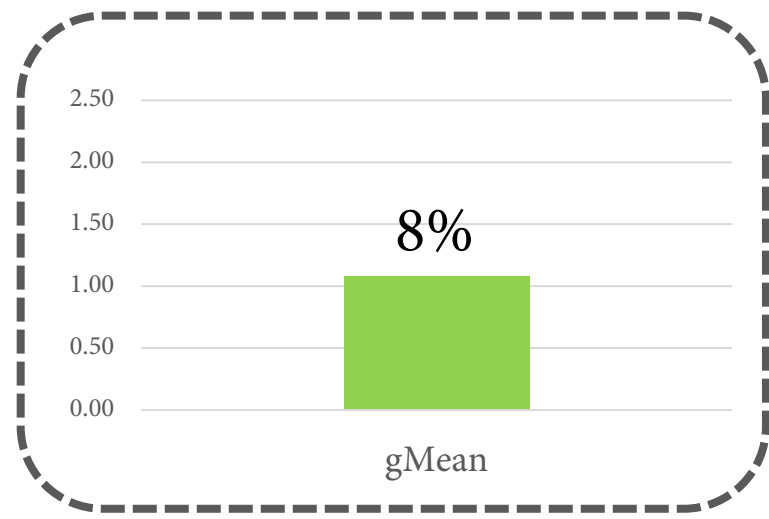
## Experimental Setup

We emulate No-FAT on x86\_64 by modifying LLVM to emit new instructions.


- `CompBase` is emulated using two multiplications followed by a `store`.
- `VerifyBounds` is emulated using dummy `stores`.



# Performance Results (x86\_64)

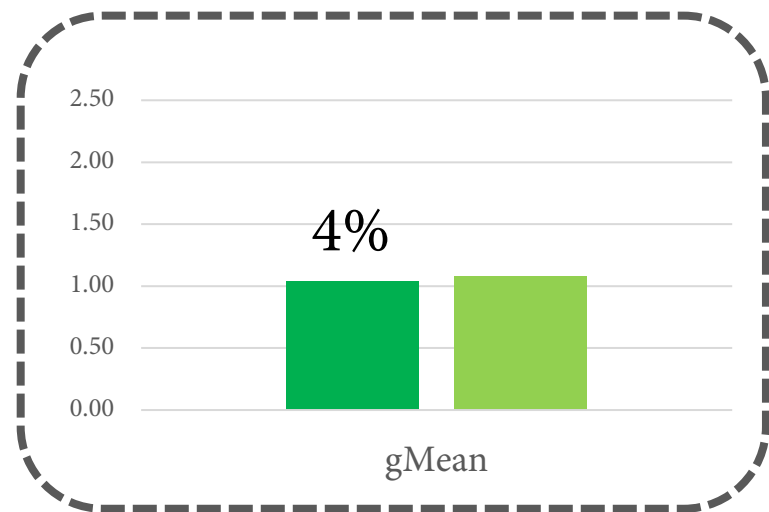


■ No-FAT






# Performance Results (x86\_64)





■ Binning-Malloc      ■ No-FAT







# Performance Results (x86\_64)







# Performance Results (x86\_64)





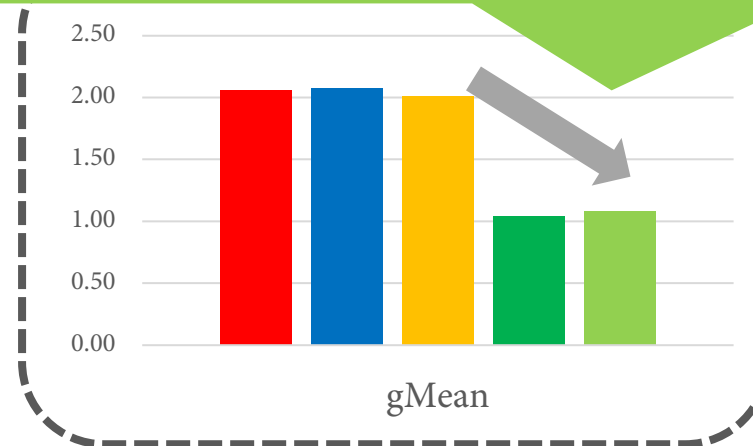



# Performance Results (x86\_64)



# Performance Results (x86\_64)

We reduce the average runtime overheads of full memory safety **from 100% to 8%!**





# Related Work



# Related Work

## Metadata

## Concerns

Memory Tagging

N-bits per pointer & allocation

Spatial & temporal safety  
limited by tag width



# Related Work

## Metadata

## Concerns

Memory Tagging

N-bits per pointer & allocation

Spatial & temporal safety limited by tag width

Tripwires

N-bits per allocation

Susceptible to non-adjacent overflows



# Related Work

## Metadata

## Concerns

Memory Tagging

N-bits per pointer & allocation

Spatial & temporal safety limited by tag width

Tripwires

N-bits per allocation

Susceptible to non-adjacent overflows

Explicit Base & Bounds

N-bits per pointer or allocation


Breaks compatibility with the rest of the system (eg. unprotected libraries).

# Related Work

	Metadata	Concerns
Memory Tagging	N-bits per pointer & allocation	Spatial & temporal safety limited by tag width
Tripwires	N-bits per allocation	Susceptible to non-adjacent overflows
Explicit Base & Bounds	N-bits per pointer or allocation	Breaks compatibility with the rest of the system (eg. unprotected libraries).
<b>No-FAT</b>	<b>Fixed (1K) bits per process</b>	<b>Requires binning allocator</b>



# Takeaways

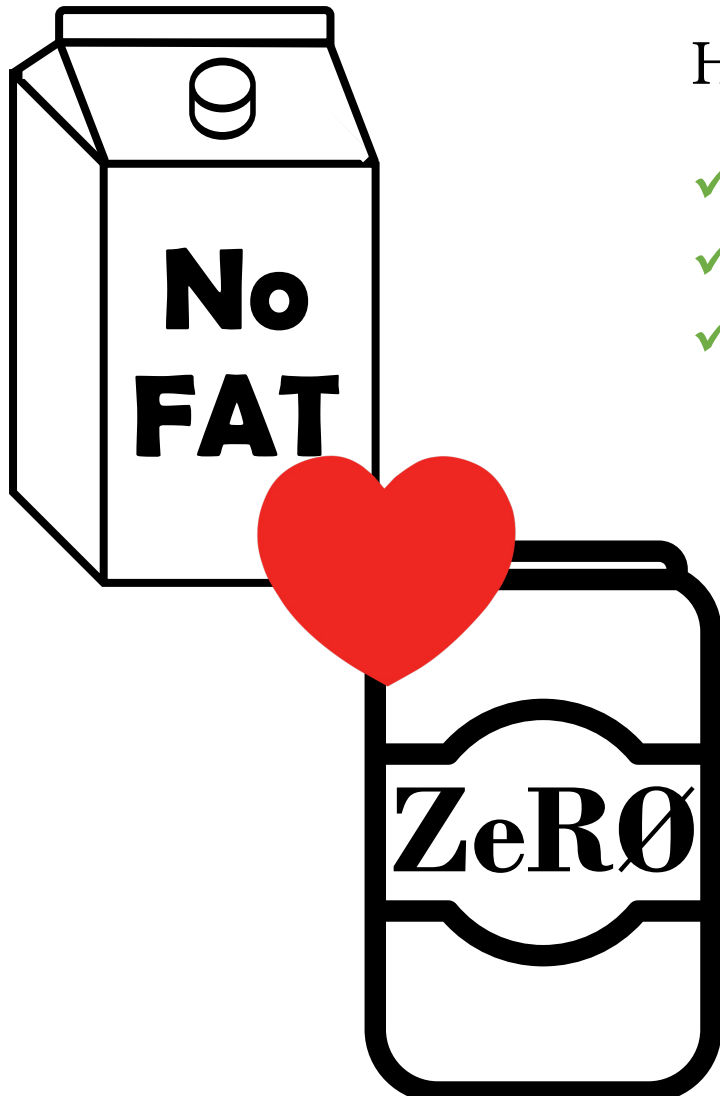


Having no metadata

- ✓ **Improves Fuzzing**
- ✓ **Improves Runtime Security**
- ✓ **Improves Resilience to Spectre-V1**



# Takeaways



Having no metadata

- ✓ Improves Fuzzing
- ✓ Improves Runtime Security
- ✓ Improves Resilience to Spectre-V1


For applications where an 8% overhead is too much, checkout [ZeRØ](https://isca21.arroyo.me).



<https://isca21.arroyo.me>



# Takeaways



Having no metadata

- ✓ **Improves Fuzzing**
- ✓ **Improves Runtime Security**
- ✓ **Improves Resilience to Spectre-V1**

The benefits of having allocation sizes as an architectural feature can go well beyond memory safety!