

Bounds Checking for the Masses

Towards Practical Exploit Mitigation in the Data-Only Attack Era



Adriaan Jacobs

Ericsson, 2025/02/25

FBI

NEWS

The M

30 Year

U.S. declares worth Ko

In a statement Tuesday

19 Dec

Chrome, Firefox and other br
WebP vulnerability

Google LLC, the Mozilla Foundation
to fix a zero-day vulnerability affecting the

3 weeks ago

Discovered in polkit

The Guardian
Heartbleed: catastrophic

Code error means
'heartbeat' function

DAN GOODIN

SC Magazine

Google patches new zero-day actively exploited in the Chrome browser

Google's most recent patch for Chrome is the fifth actively exploited zero-day targeted by threat actors this year in the popular browser.

3 days ago

Critical libwebp Vulnerability Under Active Exploitation - Gets Maximum CVSS Score

Google has assigned a new CVE identifier for a critical security flaw in the libwebp image library for rendering images in the WebP format...

5 days ago

Half a million widely trusted websites vulnerable to Heartbleed bug
A serious overrun vulnerability in the OpenSSL crypto of SSL web servers which use certificates

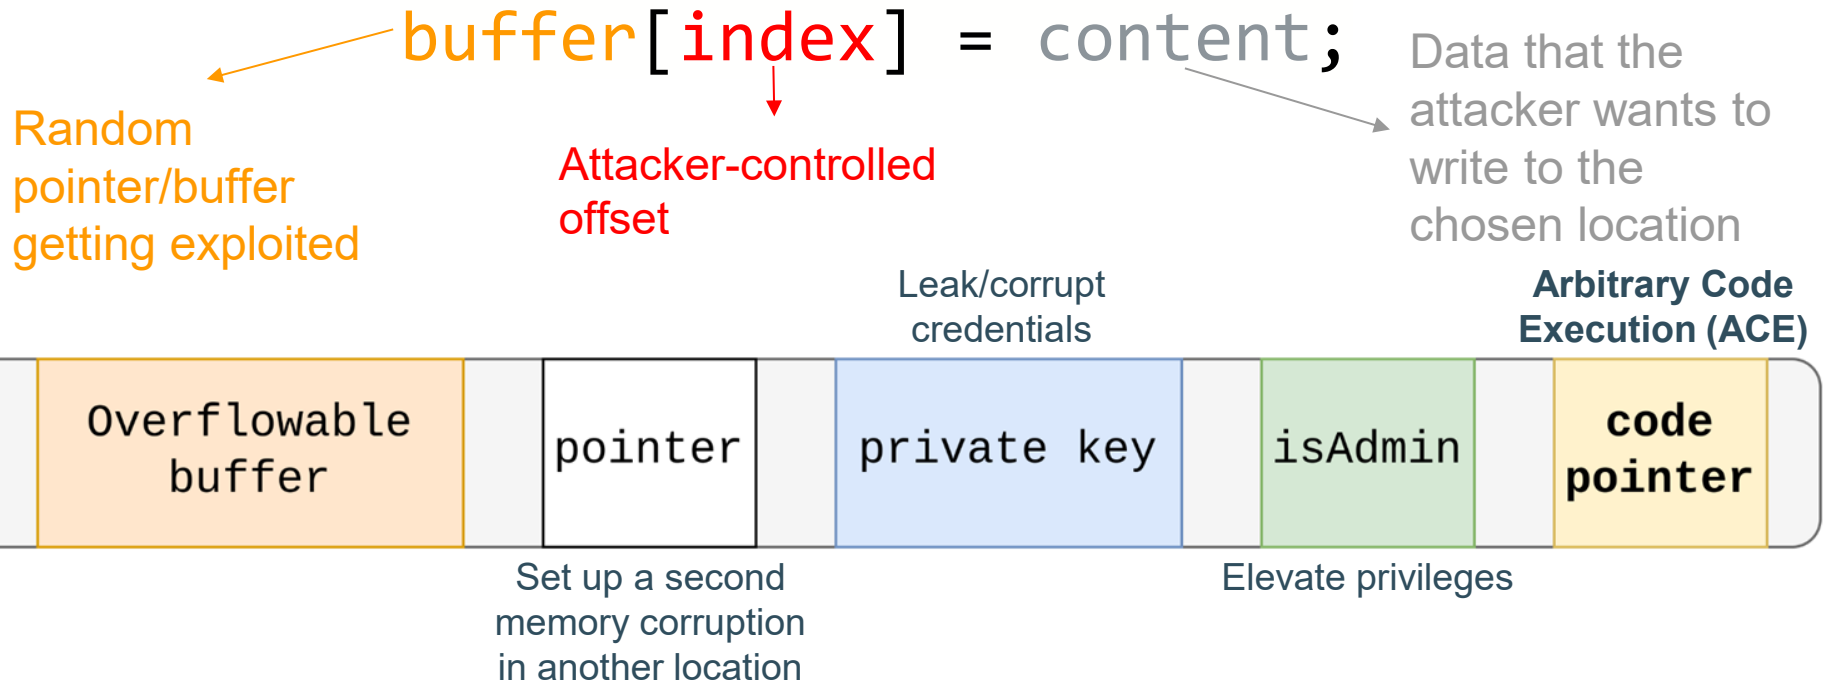


nerability

1034



Memory Errors: Limitless Fun!



An Incomplete Timeline

31st Annual Network and Distributed System Security Symposium



Hosted by: Internet Society

What's the Big Deal?

1. **Universal:** *All computer systems* run memory-unsafe code
2. **Guaranteed:** *All memory-unsafe code* has memory errors
3. **Replicated:** Everyone runs *the same unsafe binaries*
4. **Impactful:** Memory errors lead to *expressive exploits* (e.g. code execution)

ALL MODERN DIGITAL INFRASTRUCTURE

Microsoft: 70 percent of all security bugs are memory safety issues

Percentage of memory safety issues has been hovering at 70 percent for the past 12 years.

2021 in Memory Unsafety - Apple's Operating Systems

"Memory unsafety continues to dominate the total percentage of security bugs on Apple's platforms."

CYBERSECURITY & INFRASTRUCTURE SECURITY AGENCY AMERICA'S CYBER DEFENSE AGENCY

The Urgent Need for Memory Safety in Software Products

Chrome: 70% of all security bugs are memory safety issues

Google software engineers are looking into ways of eliminating memory management-related bugs from Chrome.

<https://xkcd.com/2347/>

What's the Big Deal?

1. **Un** Initial Infection Vector

- me For the fifth year in a row, exploits were the most frequently observed initial infection vector in Mandiant incident response investigations. For intrusions in which an initial infection vector was identified, 33% began with exploitation of a vulnerability. This is a decline from 2023, during which exploits represented the initial intrusion vector for 38% of intrusions, but nearly identical to the share of exploits in 2022, 32%.
2. **Gu**
- me
3. **Re**
- bir
4. **Im**
- ex

Initial Infection Vector, 2024




Didn't We Solve This Problem?

Memory-Safe Software

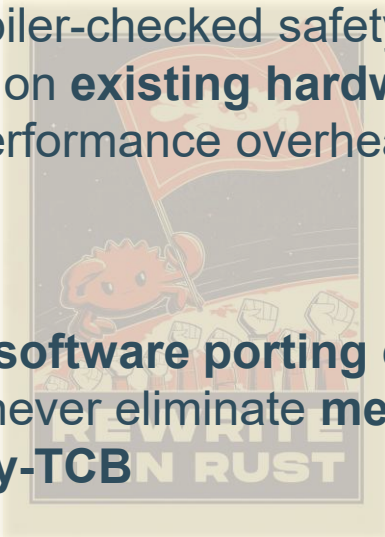


Memory-Safe Hardware



Didn't We Solve This Problem?

Memory-Safe Software

- ✓ Compiler-checked safety
 - ✓ Runs on **existing hardware**
 - ✓ No performance overhead
-
- ✗ High **software porting** cost
 - ✗ Can never eliminate **memory-safety-TCB**
- 

Memory-Safe Hardware

- ✓ Hardware-enforced safety
 - ✓ Low software porting cost
 - ✓ **Memory-safety-TCB** eliminated
-
- ✗ High **hardware replacement** cost
 - ✗ Non-zero **performance overhead**
- 

Didn't We Solve This Problem?

Memory-Safe Software

- ✓ Compiler-checked safety
- ✓ Runs on **existing hardware**
- ✓ No performance overhead

- ✗ High **software porting** cost
- ✗ Can never eliminate **memory-safety-TCB**

Memory-Safe Hardware

- ✓ Hardware-enforced safety
- ✓ Low software porting cost
- ✓ **Memory-safety-TCB** eliminated

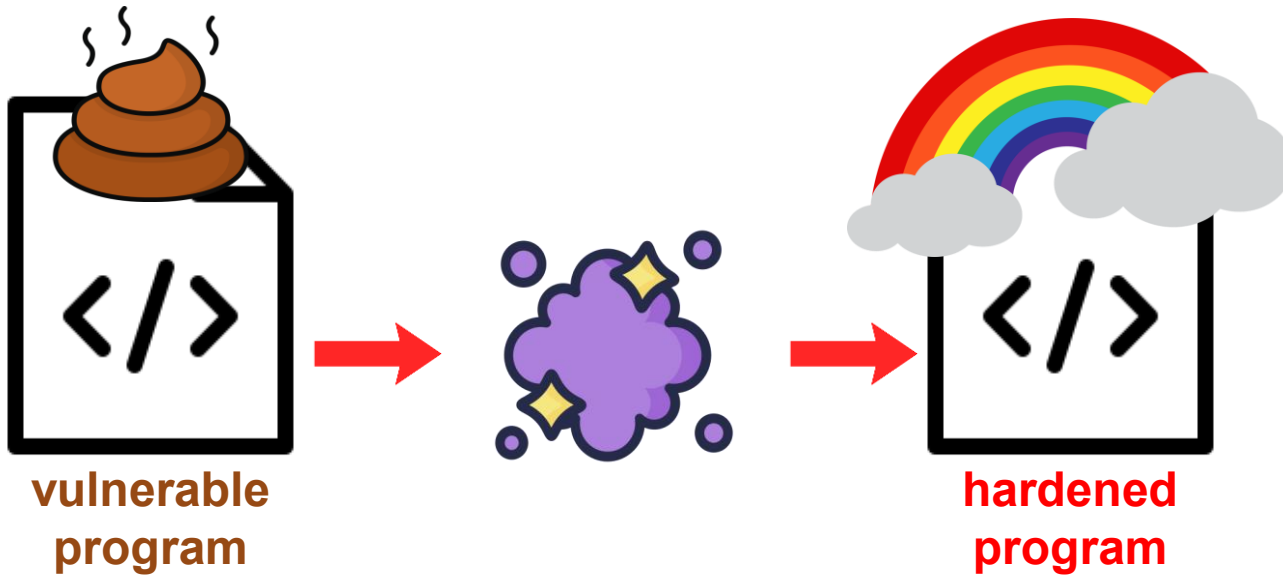
- ✗ High **hardware replacement** cost
- ✗ Non-zero **performance overhead**



**EXPLOIT
MITIGATIONS**

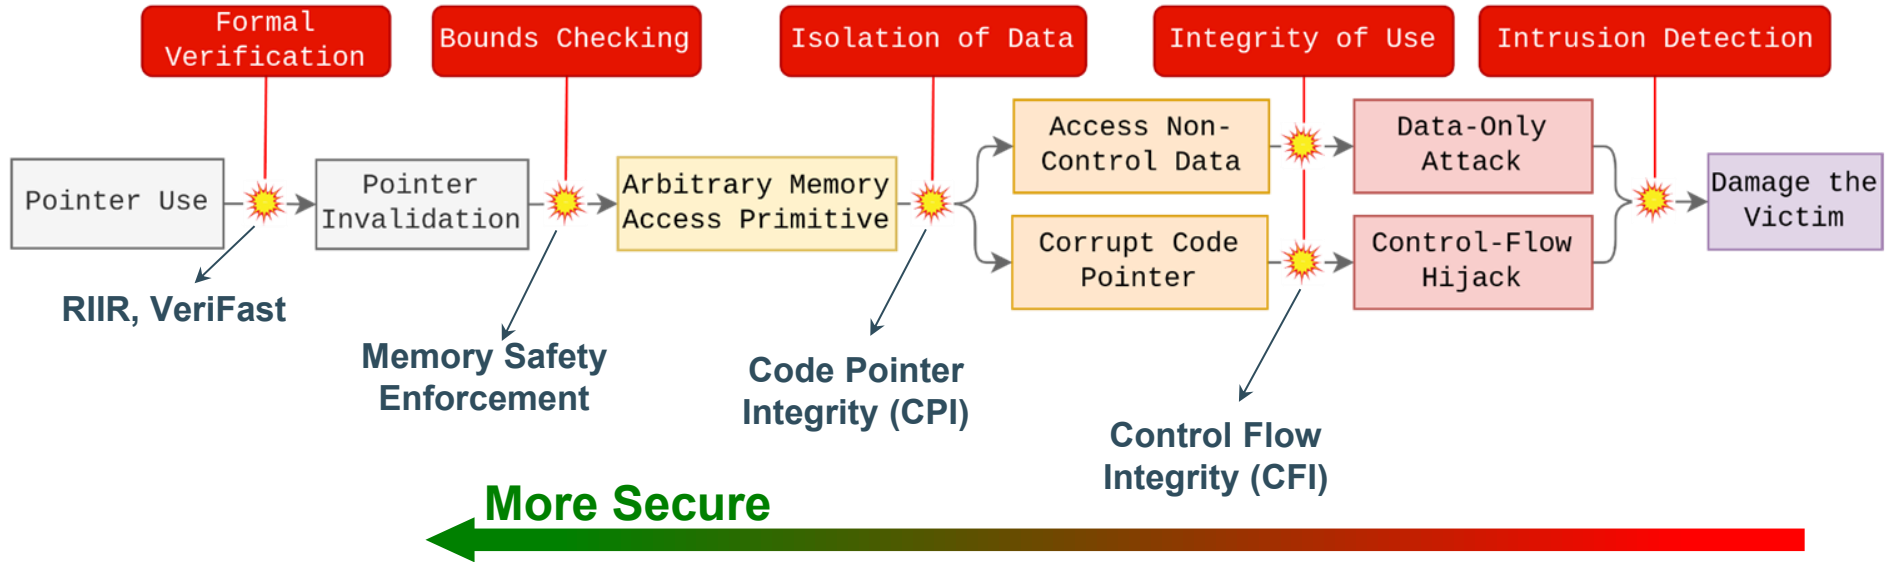


Rapid Adoption: Exploit Mitigations



Exploit Mitigations: What Are Our Options?

More Practical



Attackers Are Only Getting More Creative

Control Flow Hijacking



```
while (connect_limit-->0) {  
    pkt = read_from_network(buf);  
    srv->log(pkt);  
    if (pkt->type == NONE)  
        return;  
    if (pkt->type == STREAM)  
        send(srv->fd, srv->stream);  
    else {  
        srv->typ = *type;  
        srv->total += *size;  
    }  
}
```

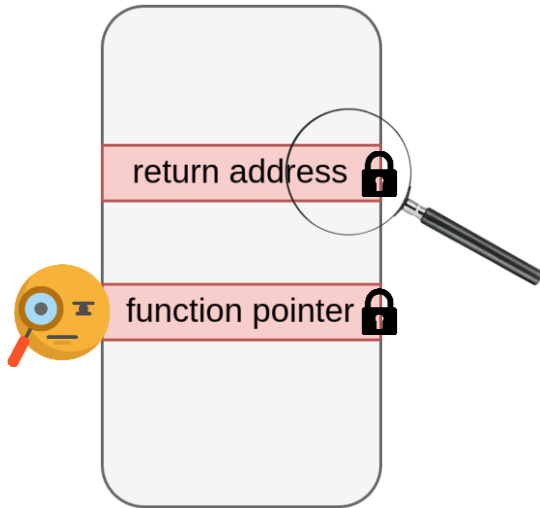


Data-Only Attacks

```
while (connect_limit-->0) {  
    pkt = read_from_network(buf);  
    srv->log(pkt);  
    if (pkt->type == NONE)  
        return;  
    if (pkt->type == STREAM)  
        send(srv->fd, srv->stream);  
    else {  
        srv->typ = *type;  
        srv->total += *size;  
    }  
}
```

Attackers Are Only Getting More Creative

Control Flow Hijacking

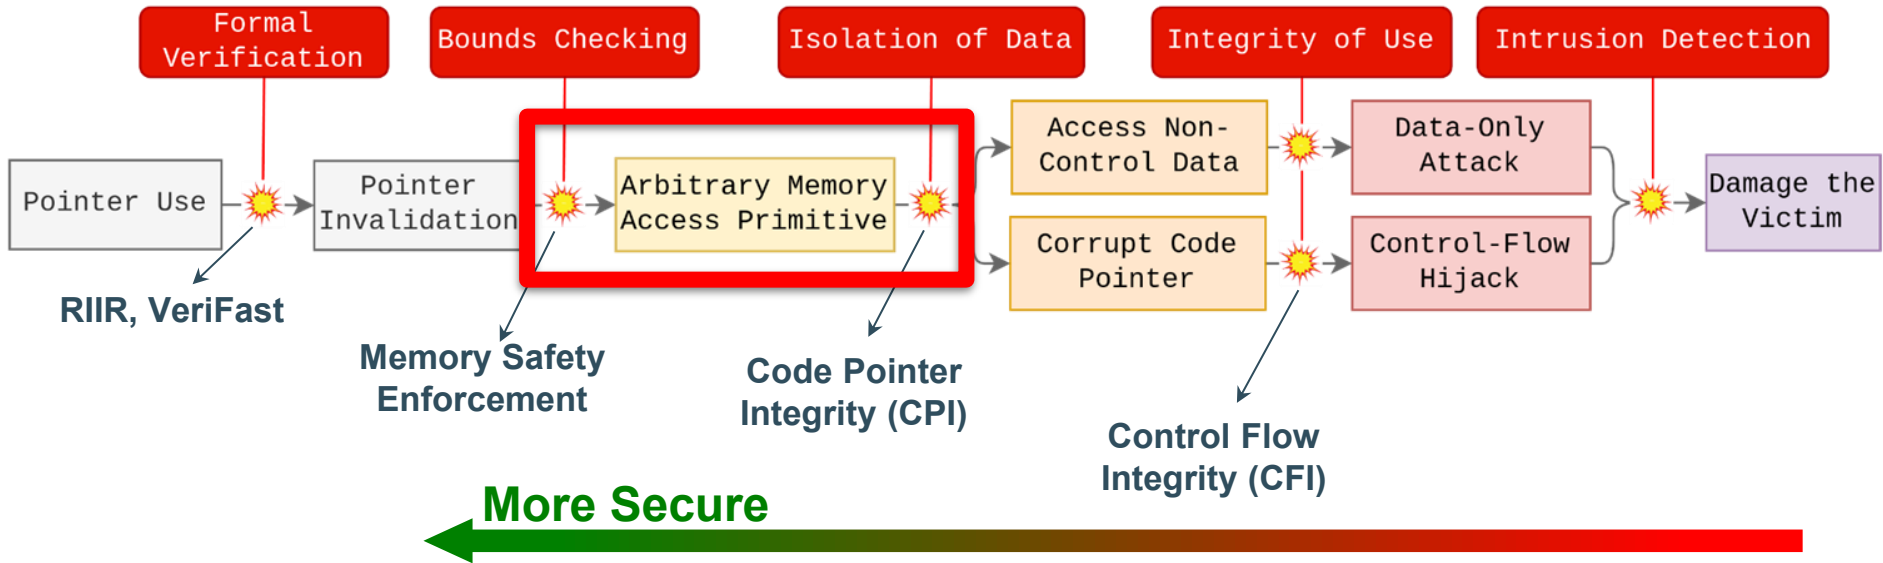


Data-Only Attacks



Data-Only Attack Mitigations: What Are Our Options?


More Practical



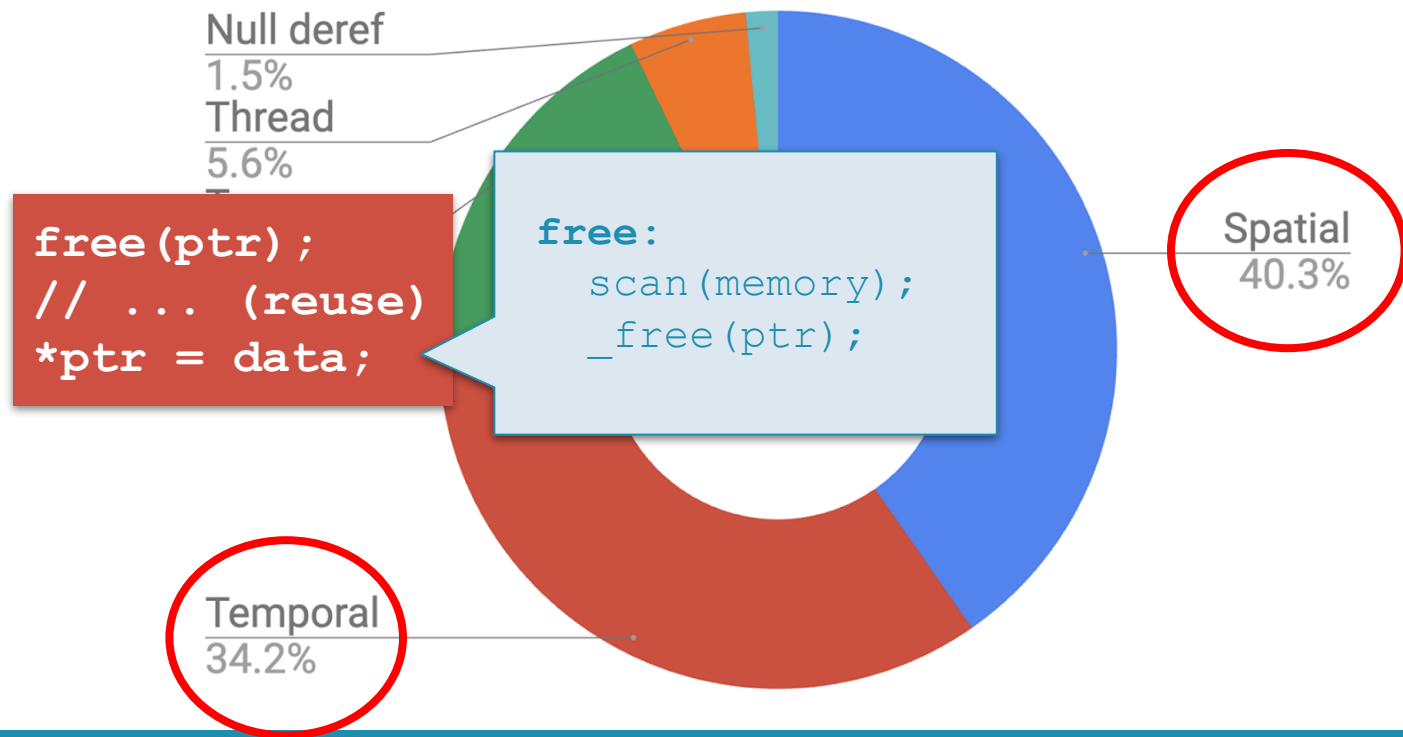
More Secure



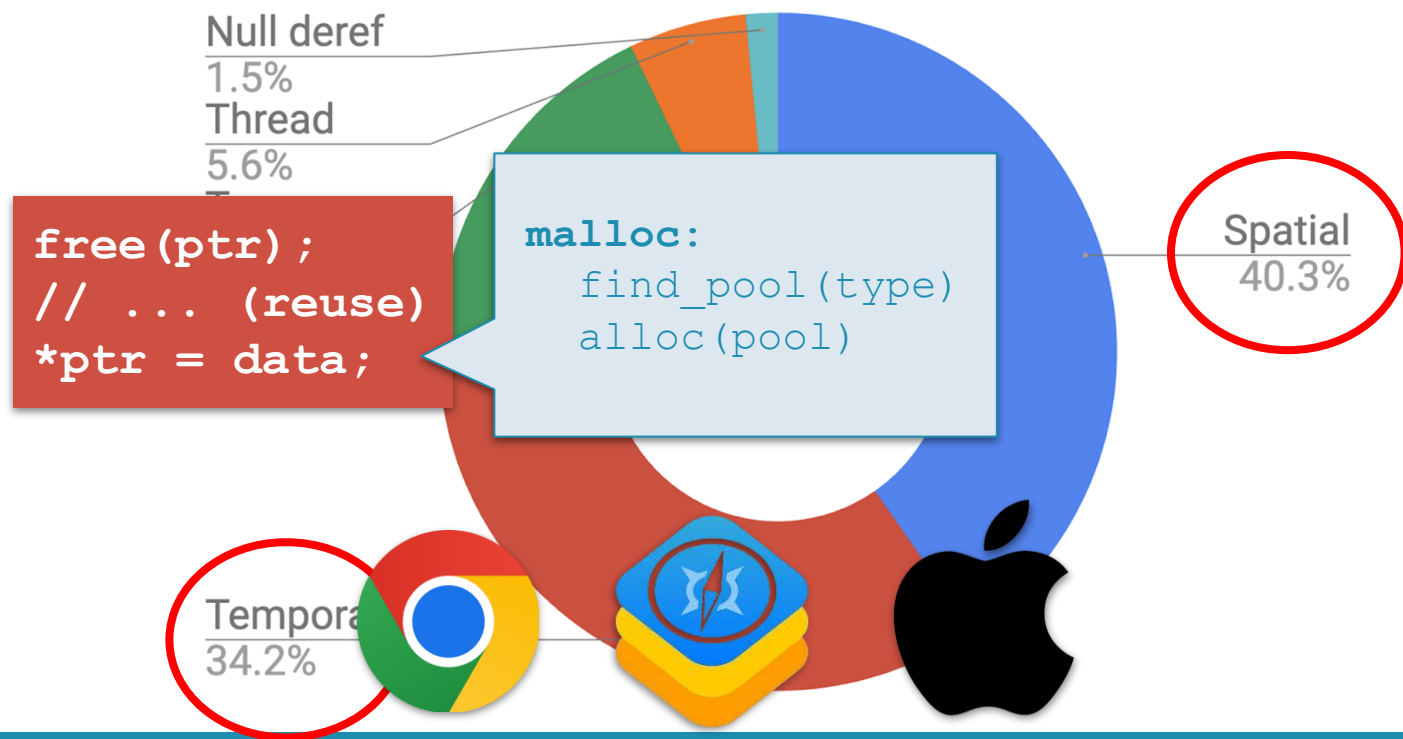
Memory Safety 0-days Exploited in the Wild




Memory Safety 0-days Exploited in the Wild



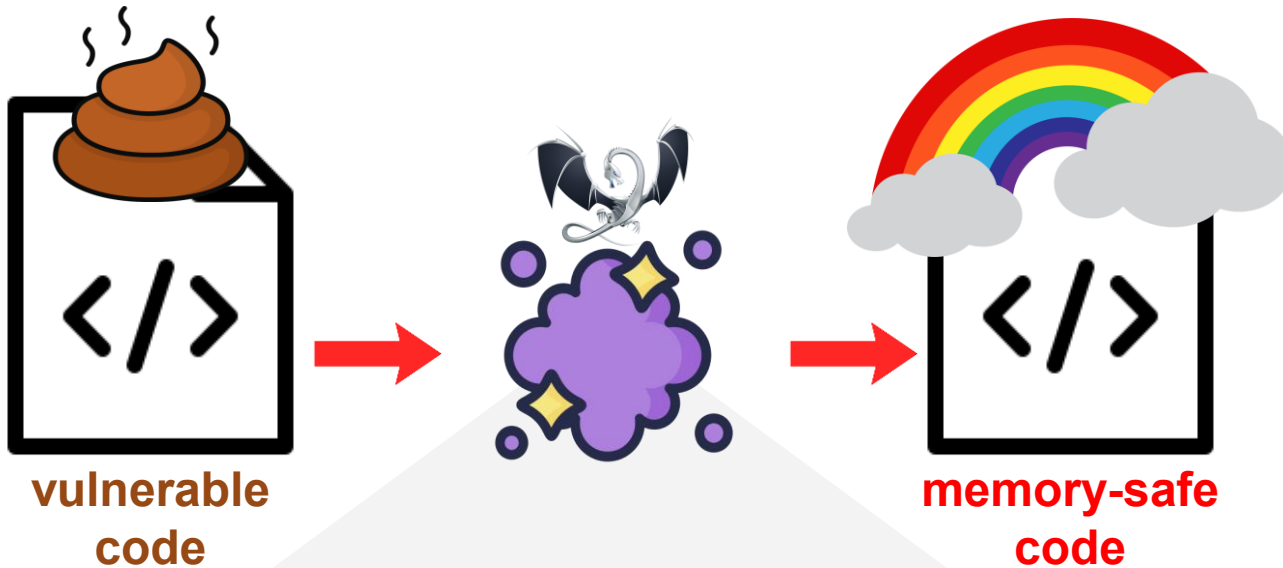
Memory Safety 0-days Exploited in the Wild



Memory Safety 0-days Exploited in the Wild



How Bounds Checkers Work ...



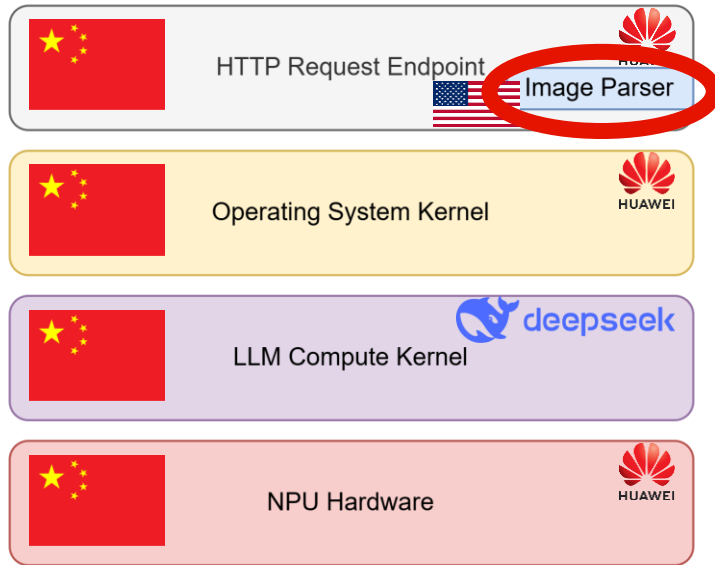
```
if (i < 0 || i >= size)  
    trap();  
buf[i] = data;
```

How Bounds Checkers Work ... They Don't


Approach	Type	Hardware Assistance	Pointer Instrumentation				Target Memory			Ext. code Interop	Benign Obj ptrs	Compatibility		Run-Time Overhead
			Deref.	Arith.	Escape	Copy	Stack	Heap	Global			Maximal Address Space (2 ⁿ)	Max. # Live Objects	
Fat Pointers [15–18]	P	-	✓	-	-	-	✓	✓	✓	●	●	▲	▲	●
SoftBound [19]	P	-	✓	-	✓	✓	✓	✓	●	●	▲	▲	▲	●
Intel MPX [39, 122]	P	Intel MPX	✓	-	✓	-	✓	✓	✓	●	●	▲	▲	●
MESH [41]	P	-	✓	✗	-	-	✗	✓	✓	●	●	48	2 ¹⁶	●
PACMem [20]	P	ARM PA	✓	-	-	-	✓	✓	✓	●	●	39	2 ²⁴	●
CUP [27]/Midfat [40]	P	-	✓	-	✓	-	✓	✓	✓	●	●	▲	2 ¹¹	4 GB
CGuard [30]	P	-	✓	✓	-	-	✓	✓	✓	●	●	48	▲	●
Delta Pointers [28]	P	-	✓	-	-	-	✓	✓	✓	●	●	32	▲	●
J&K [14]	O	-	✓	✓	-	-	✓	✓	✓	●	●	▲	▲	●
CRED [24]	O	-	✓	✓	-	-	✓	✓	✓	●	●	▲	▲	●
PariCheck [34]	O	-	✓	✓	-	-	✓	✓	✓	●	●	▲	▲	●
Baggy Bounds ^{64-bit} [23]	O	-	-	✓	-	-	✓	✓	✓	●	●	58	▲	●
LowFat [21, 38, 93]	O	-	-	✓	-	-	✓	✓	✓	●	●	▲	▲	●
CAMP [23]	O	-	-	✓	-	-	✓	✓	✓	○ [†]	○	▲	▲	●
ShadowBound [56]	O	-	-	-	-	-	✗	✓	✓	●	○	▲	▲	●
TailCheck [31]	R	-	✓	✓	✓	-	✓	✓	✓	●	●	48	▲	▲
StickyTags [31]	R	ARM MTE	✓	✗	-	-	✓	✓	✓	●	●	▲	▲	▲
ZOMETag [11]	P+O+R	ARM MTE	✓	✓	-	-	✓	✓	✓	●	● [‡]	48	16 × 2 ⁴⁸⁻³² = 2 ²⁰	4 GB
FRAMER [21]	P+O	-	✓	✓	✓	-	✓	✓	✓	●	● [‡]	▲	▲	●

Legend: P Pointer-based, O Object-based, R Red-zone-based (aka Tripwire), ✓ Has this feature, ✗ Lacks this feature, ✓ Has this feature (but is fundamentally flawed), ✗ Lacks this feature (but no fundamental issue in the design), - Feature is not required, ● Full Compatibility, ○ Partial Compatibility, ○ No Compatibility, ● <30% (average) overhead, ● <50% (average) overhead, ● >50% (average) overhead, ▲ The maximum possible by the underlying OS/architecture, † Only off by One, ‡ Within limited range. If the defense includes more than just spatial checks, we model their spatial component alone.

Is Performance Really The (Only) Problem?




Harden
this part!



Existing Bounds Checkers Are Stuck

How to recover intended referent
during dereference?



Propagate bounds
information with all pointers
throughout the program.

POINTER-BASED

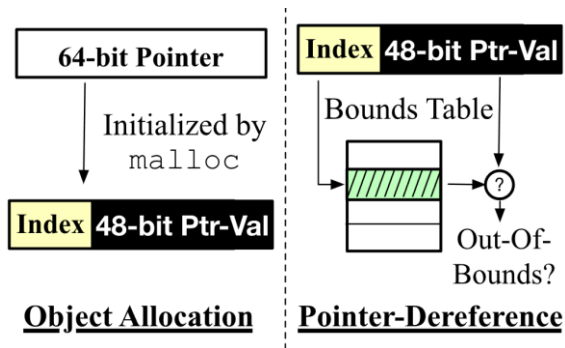
```
struct fat_ptr {  
    void* ptr;  
    void* upper_bound;  
    void* lower_bound;  
};
```

Existing Bounds Checkers Are Stuck

How to recover intended referent during dereference?

Propagate bounds information with all pointers throughout the program.

POINTER-BASED



Existing Bounds Checkers Are Stuck

How to recover intended referent during dereference?

Propagate bounds information with all pointers throughout the program.

POINTER-BASED

	ignored	canonical	effective address
ARM TBI	63 8-bit IGNORED	55 8-bit 0...0	47 48-bit addr
AMD UAI	63 7-bit IGNORED	56 10-bit 0...0	46 47-bit addr
Intel LAM U57	6362 0 6-bit IGNORED	56 10-bit 0...0	46 47-bit addr
Intel LAM U48	6362 0	15-bit IGNORED	4746 0 47-bit addr

Existing Bounds Checkers Are Stuck

How to recover intended referent during dereference?


Propagate bounds information with all pointers throughout the program.

POINTER-BASED

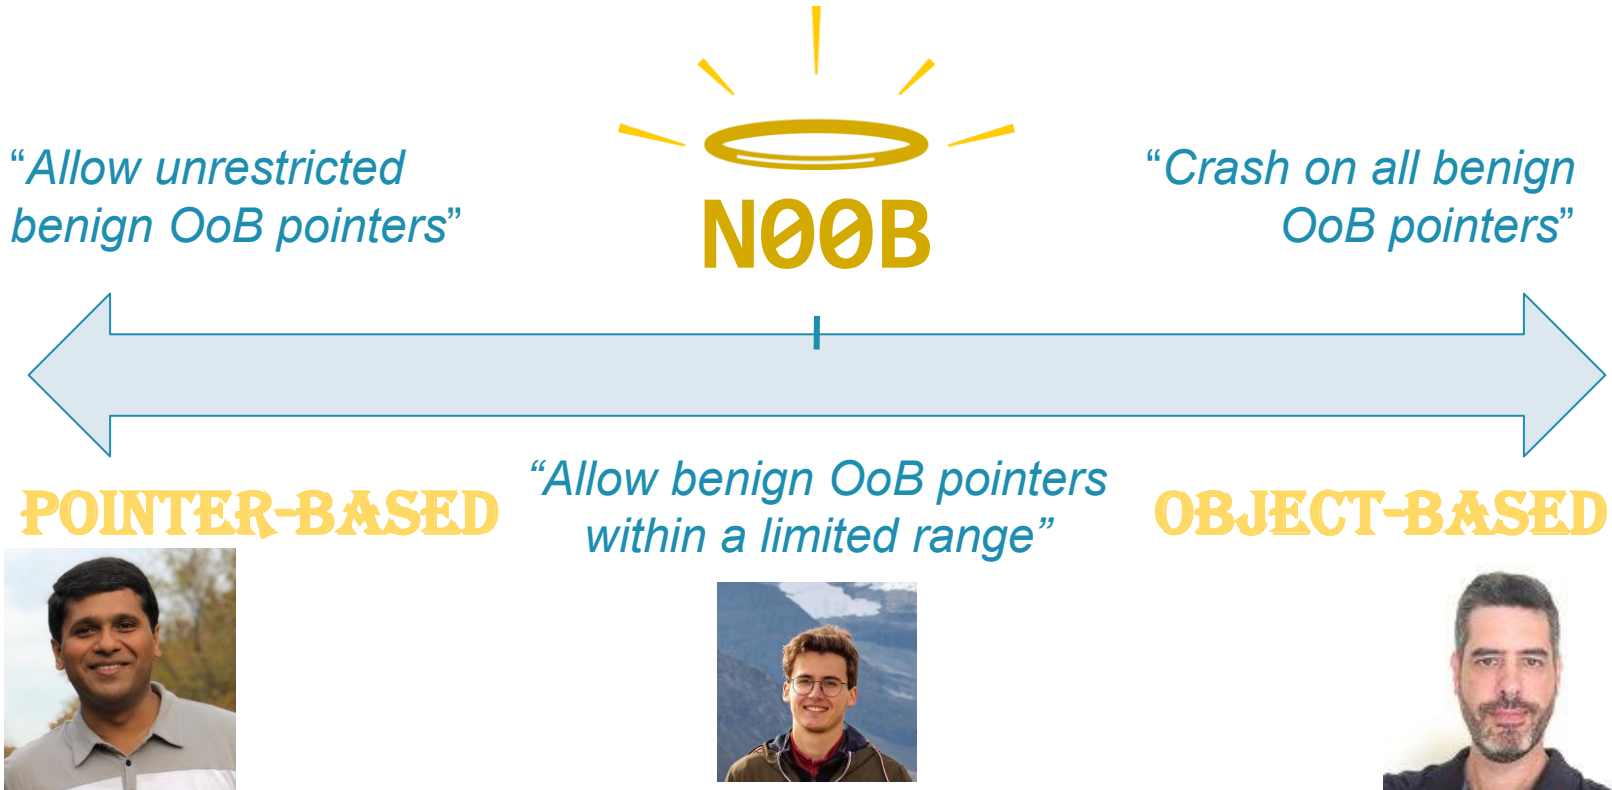
	ignored	canonical	effective address
ARM TBI	63 8-bit IGNORED	55 8-bit 0...0	47 48-bit addr
AMD UAI	63 7-bit IGNORED	56 10-bit 0...0	46 47-bit addr
Intel LAM U57	6362 6-bit IGNORED	56 10-bit 0...0	46 47-bit addr
Intel LAM U48	6362 15-bit IGNORED	0	4746 47-bit addr

Never let pointers go out of bounds in the first place.


OBJECT-BASED




N00B: Hybrid Bounds Checking



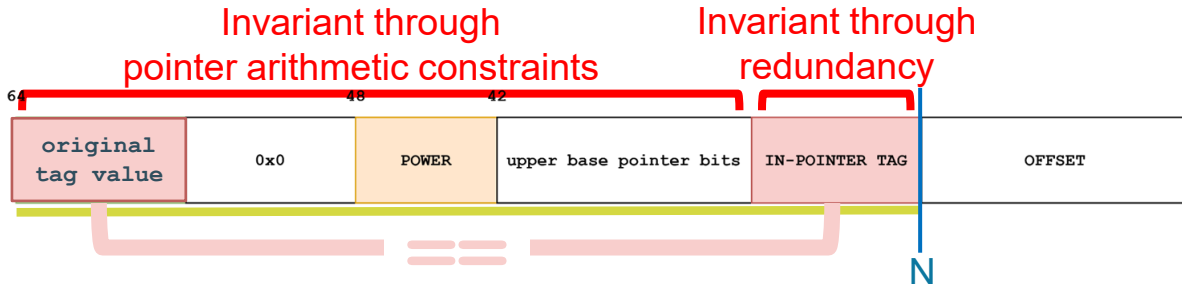
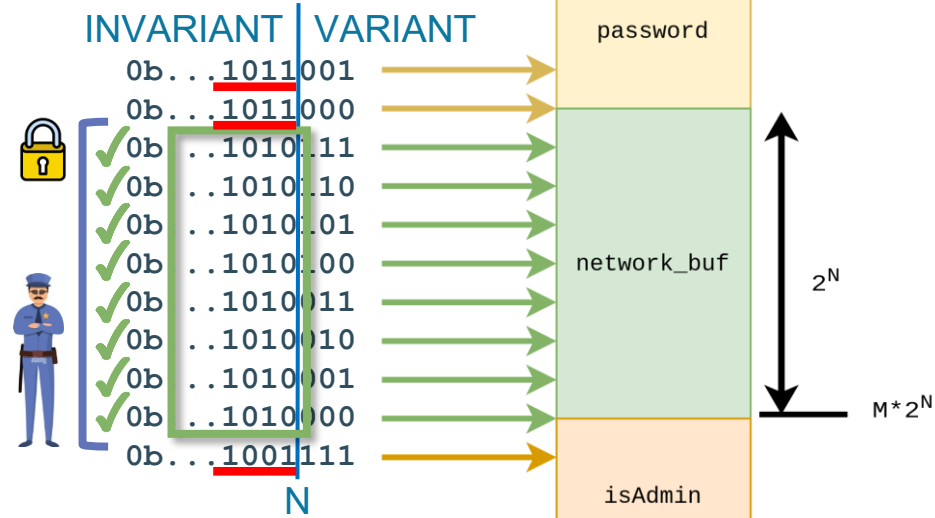
The N00B Idea



The N00B Idea

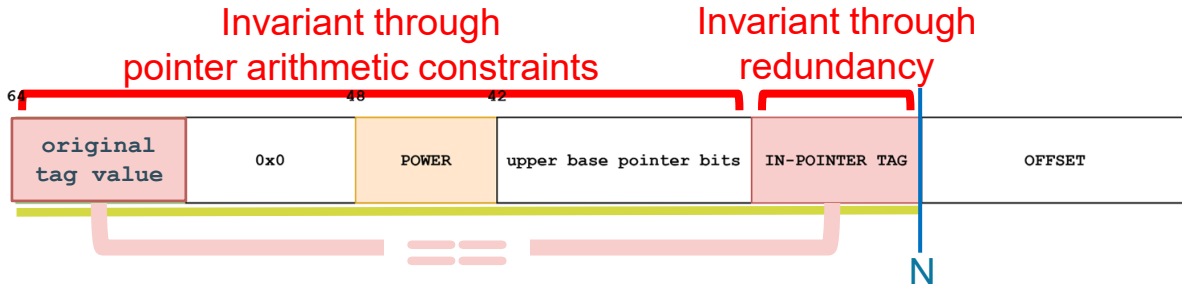
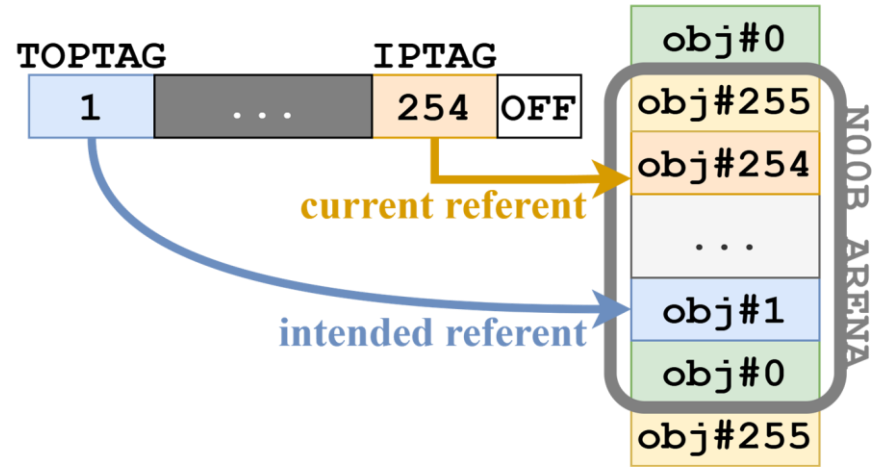


The N00B Idea



The N00B Idea

- **“Far” Overflows:** coarse-grained arithmetic checks (~ Low-Fat Pointers)
- **“Near” Overflows:** fine-grained tag checks (~ *Implicit* Memory Tagging)



Twice the Checks, Twice the Overhead?

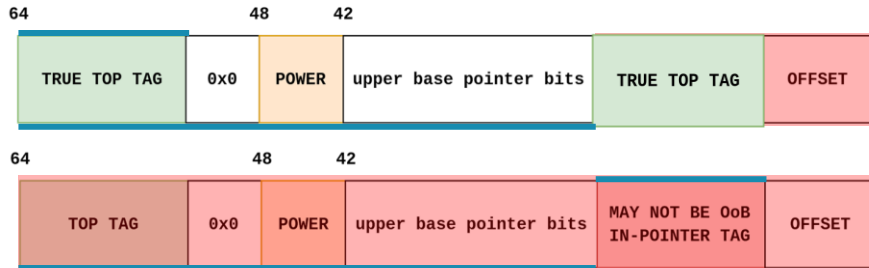
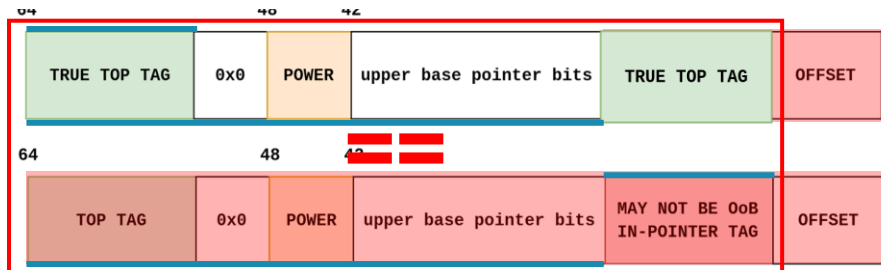



Table 4: Run-time overhead comparison of N00B versus the fastest current software-based bounds checkers.

Bounds Checker	SPEC CPU2006	SPECspeed 2017
CAMP [26]	76.1% [76]	48.6% [76]
ShadowBound [53]	20.1% [76]	5.72%
Low-Fat [24]	26-43%	27-42%
Delta Pointers [29]	35% [61]	-
CGuard [31]	-	42.1%
N00B	29-54%	29-54%

Twice the Checks, Twice the Overhead?



```

1 N = (base >> 42) & 0x7f;
2 referent = (base >> N) & ~TAG_MAX; /*arena base*/
3 toptag = base >> (64 - TW); /*toptag*/
4 referent |= toptag; /*orig obj*/
5 if (referent != (ptr >> N)) trap();

```

Bounds Checker	SPEC CPU2006	SPECspeed 2017
CAMP [26]	76.1% [76]	48.6% [76]
ShadowBound [53]	20.1% [76]	5.72%
Low-Fat [24]	26-43%	27-42%
Delta Pointers [29]	35% [61]	-
CGuard [31]	-	42.1%
N00B	29-54%	29-54%

Twice the Checks, Twice the Overhead?

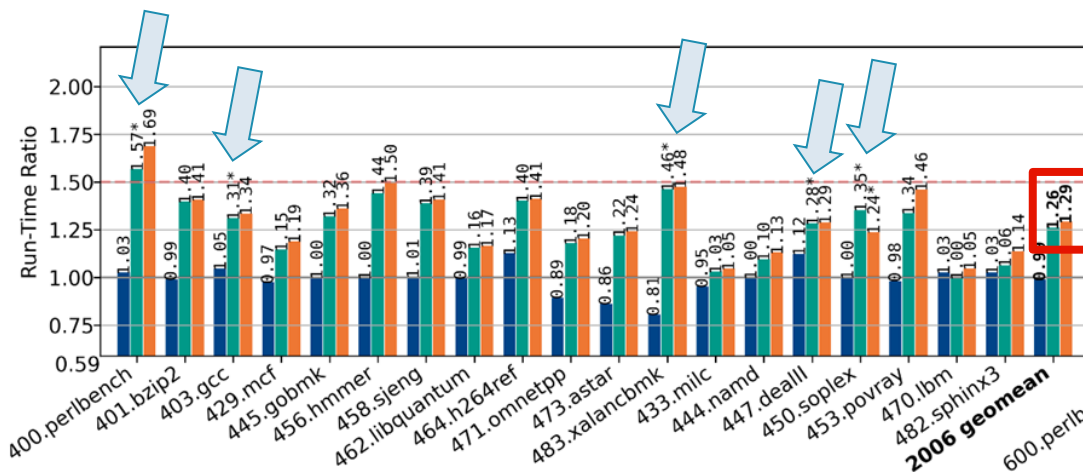
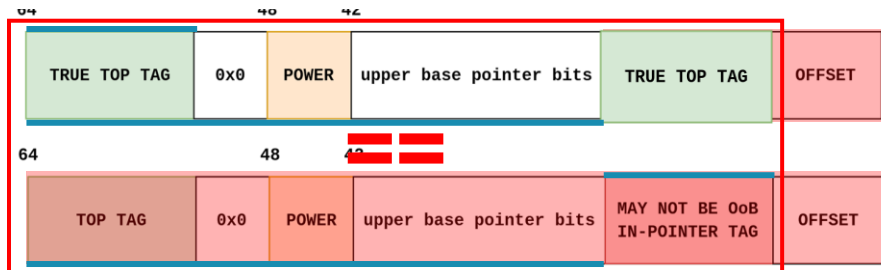


Table 3: Unique benign out-of-bounds cases in SPEC programs that Low-Fat Pointers crashes on.

Benchmark	Unique occurrences
400.perlbench	8
403.gcc	45
483.xalancbmk	49
447.dealII	11
600.perlbench_s	807
623.xalancbmk_s	1
641.leela_s	7



```

1 N = (base >> 42) & 0x7f;
2 referent = (base >> N) & ~TAG_MAX; /*arena base*/
3 toptag = base >> (64 - TW); /*toptag*/
4 referent |= toptag; /*orig obj*/
5 if (referent != (ptr >> N)) trap();

```

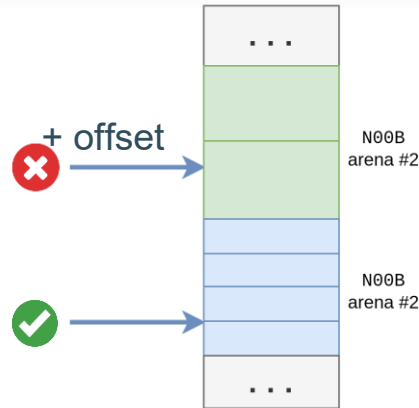
N00B Is Not Perfect...

Table 2. Applied N00B compatibility patches.

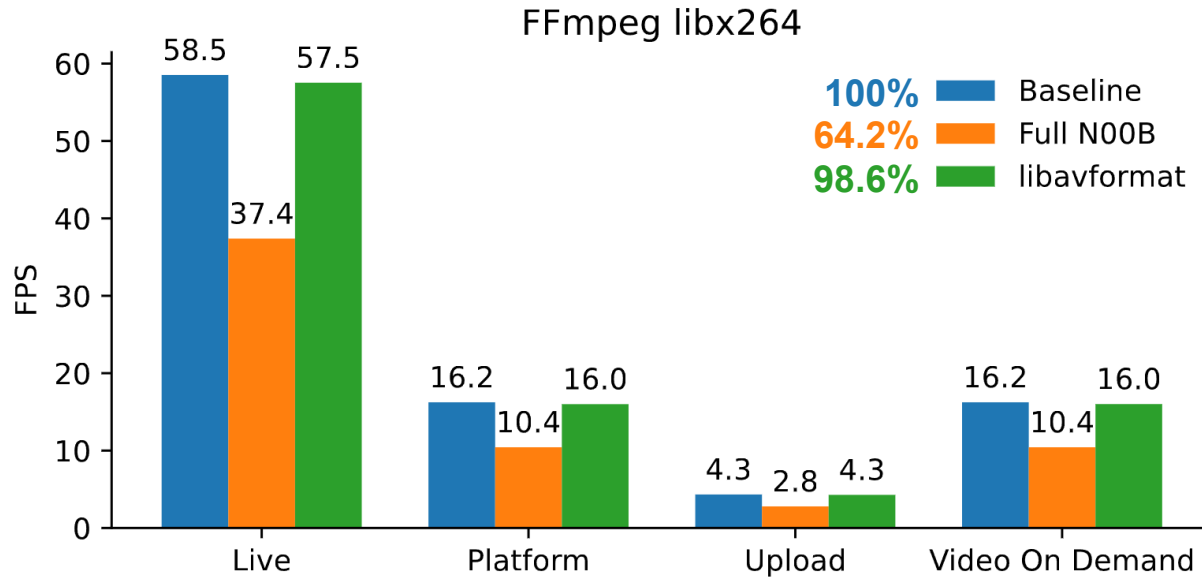
Program	Revision	Year	Message
403.gcc	r62672	2003	“Don’t use offset pointers.”
403.gcc	r89543	2004	“Avoid undefined pointer arithmetic on qty_table”
403.gcc	r79945	2004	“Don’t play queer offsetting games with reg_known_value”

N00B Is Not Perfect...

```
1 offset = realloc(array, grow) - array;  
2 for (i = 0; i < size; i++)  
3   array[i].ptr += offset;
```



N00B Is Not Perfect... But It *Works*



N00B: Bounds Checking for the Masses

Adriaan Jacobs Carlo Ramponi Jonas Roels
DistriNet, KU Leuven *University of Trento* *DistriNet, KU Leuven*

Bruno Crispo Silviu Vlasceanu Mahmoud Ammar Stijn Volckaert
University of Trento *Huawei Research* *Huawei Research* *DistriNet, KU Leuven*

Under submission to ACM CCS'26

So What's Industry Doing?



So What's Industry Doing?



```
int x = 0;
_Ptr<int> x = &x;
// no arithmetic on x allowed
```

```
int a[5] = { 0, 1, 2, 3, 4};
_Array_ptr<int> p : count(5) = a;
// p points to 5 elements.
```



C++ Safe Buffers

1. **STL hardening**
"terminate the program if you go out of bounds."
2. **Reject raw pointer arithmetic**



-fbounds-safety
Enforcing bounds safety for production C code



Most code **already contains bounds information**

- Specify through **annotations**

```
void foo(int *__counted_by(N) ptr, size_t N);
```

So What's Industry Doing?

1. Focus on **gradual adoption**: opt-in & interoperable

2. Willing to make **manual source code changes**

3. Maximize **static information & use type system**

```
int x = 0;
_Ptr<int> x = &x;
// no arithmetic on x allowed
```

```
int a[5] = { 0, 1, 2, 3, 4 };
_Array_ptr<int> p : counted(5);
// p points to 5 elements
```

```
void foo(int *__counted_by(N) ptr, size_t N);
```

Bounds Checking: Academics versus Industry

Industry



Slow to adopt



Fast to run

Academia




Fast to adopt



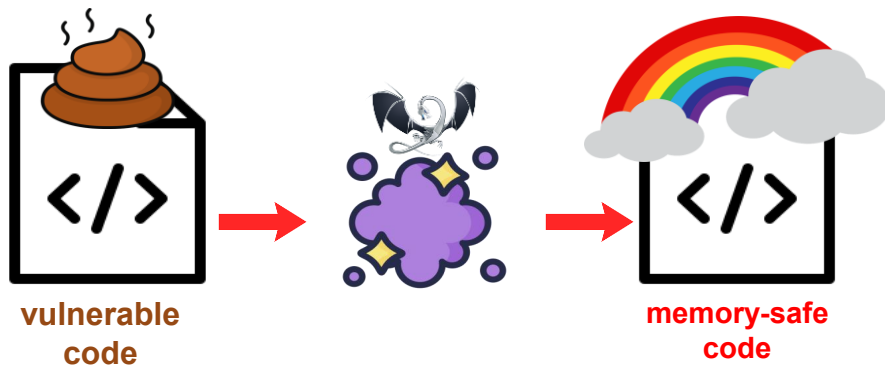
Slow to run

Type Information Is Key




```
struct node {  
    int data;  
    struct node* prev, next;  
};  
  
while ((node = node->next)) {  
    if (node->data > 5)  
        return node;  
}
```

Type Information Is Key



```
struct node {  
    int data;  
    struct node* __single prev, next;  
};  
  
while ((node = node->next)) {  
    if (node->data > 5)  
        return node;  
}
```

Type Information is Key



Questions?