


Precise and Scalable Detection of Double-Fetch Bugs in OS Kernels

Meng Xu, Chenxiong Qian, Kangjie Lu⁺, Michael Backes*, Taesoo Kim

*Georgia Tech | University of Minnesota⁺ | CISPA, Germany**

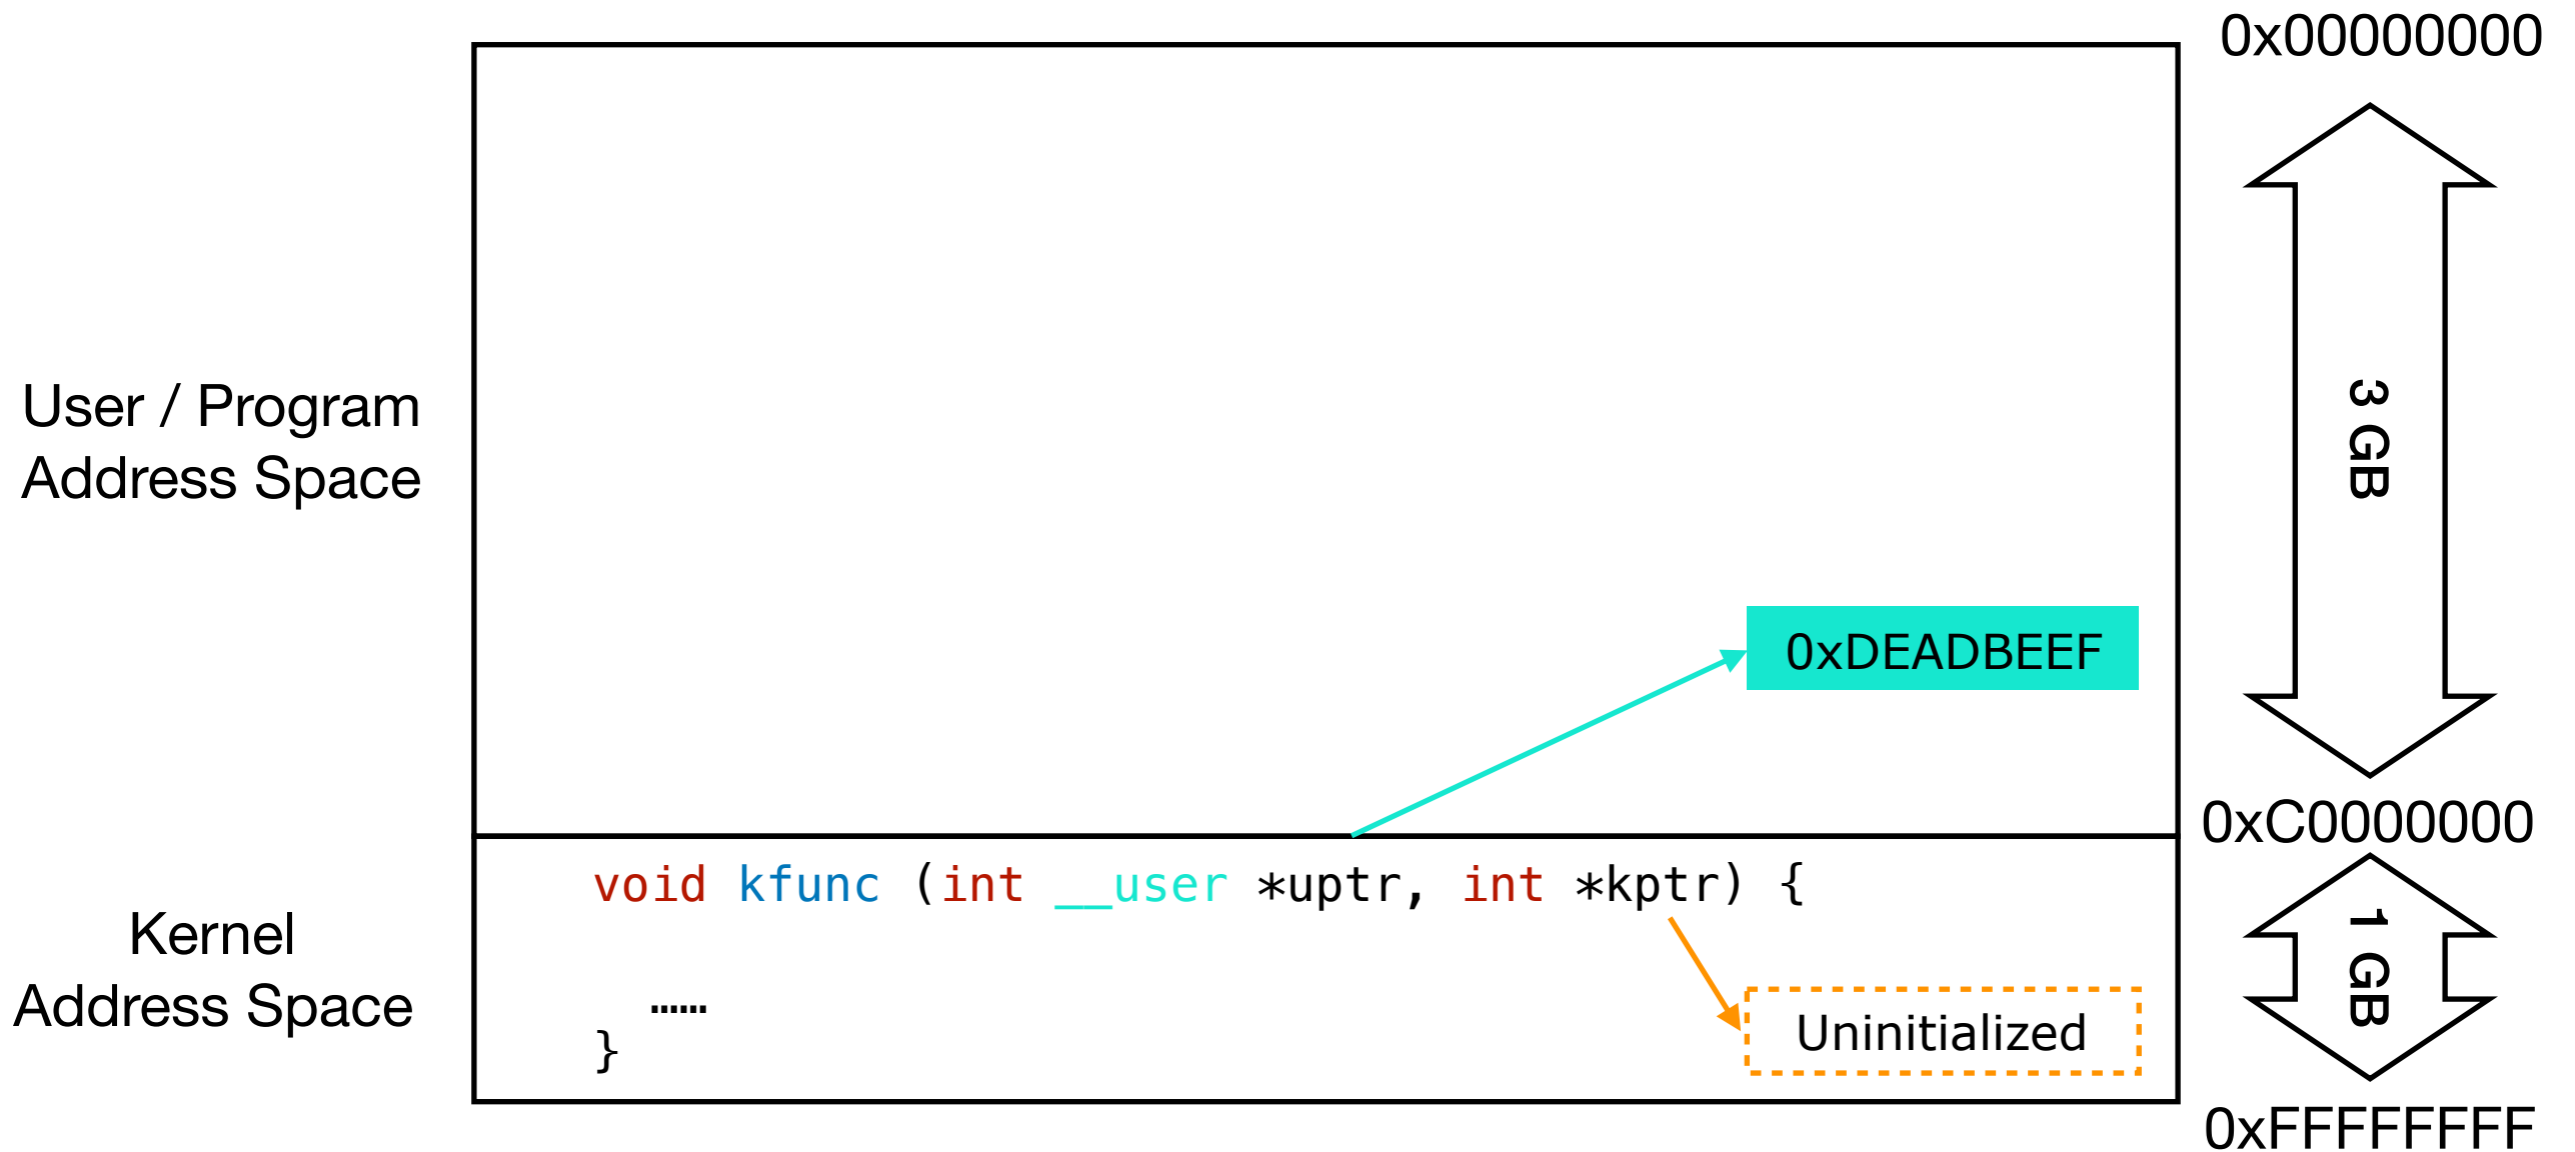
What is Double-Fetch?

Address Space Separation



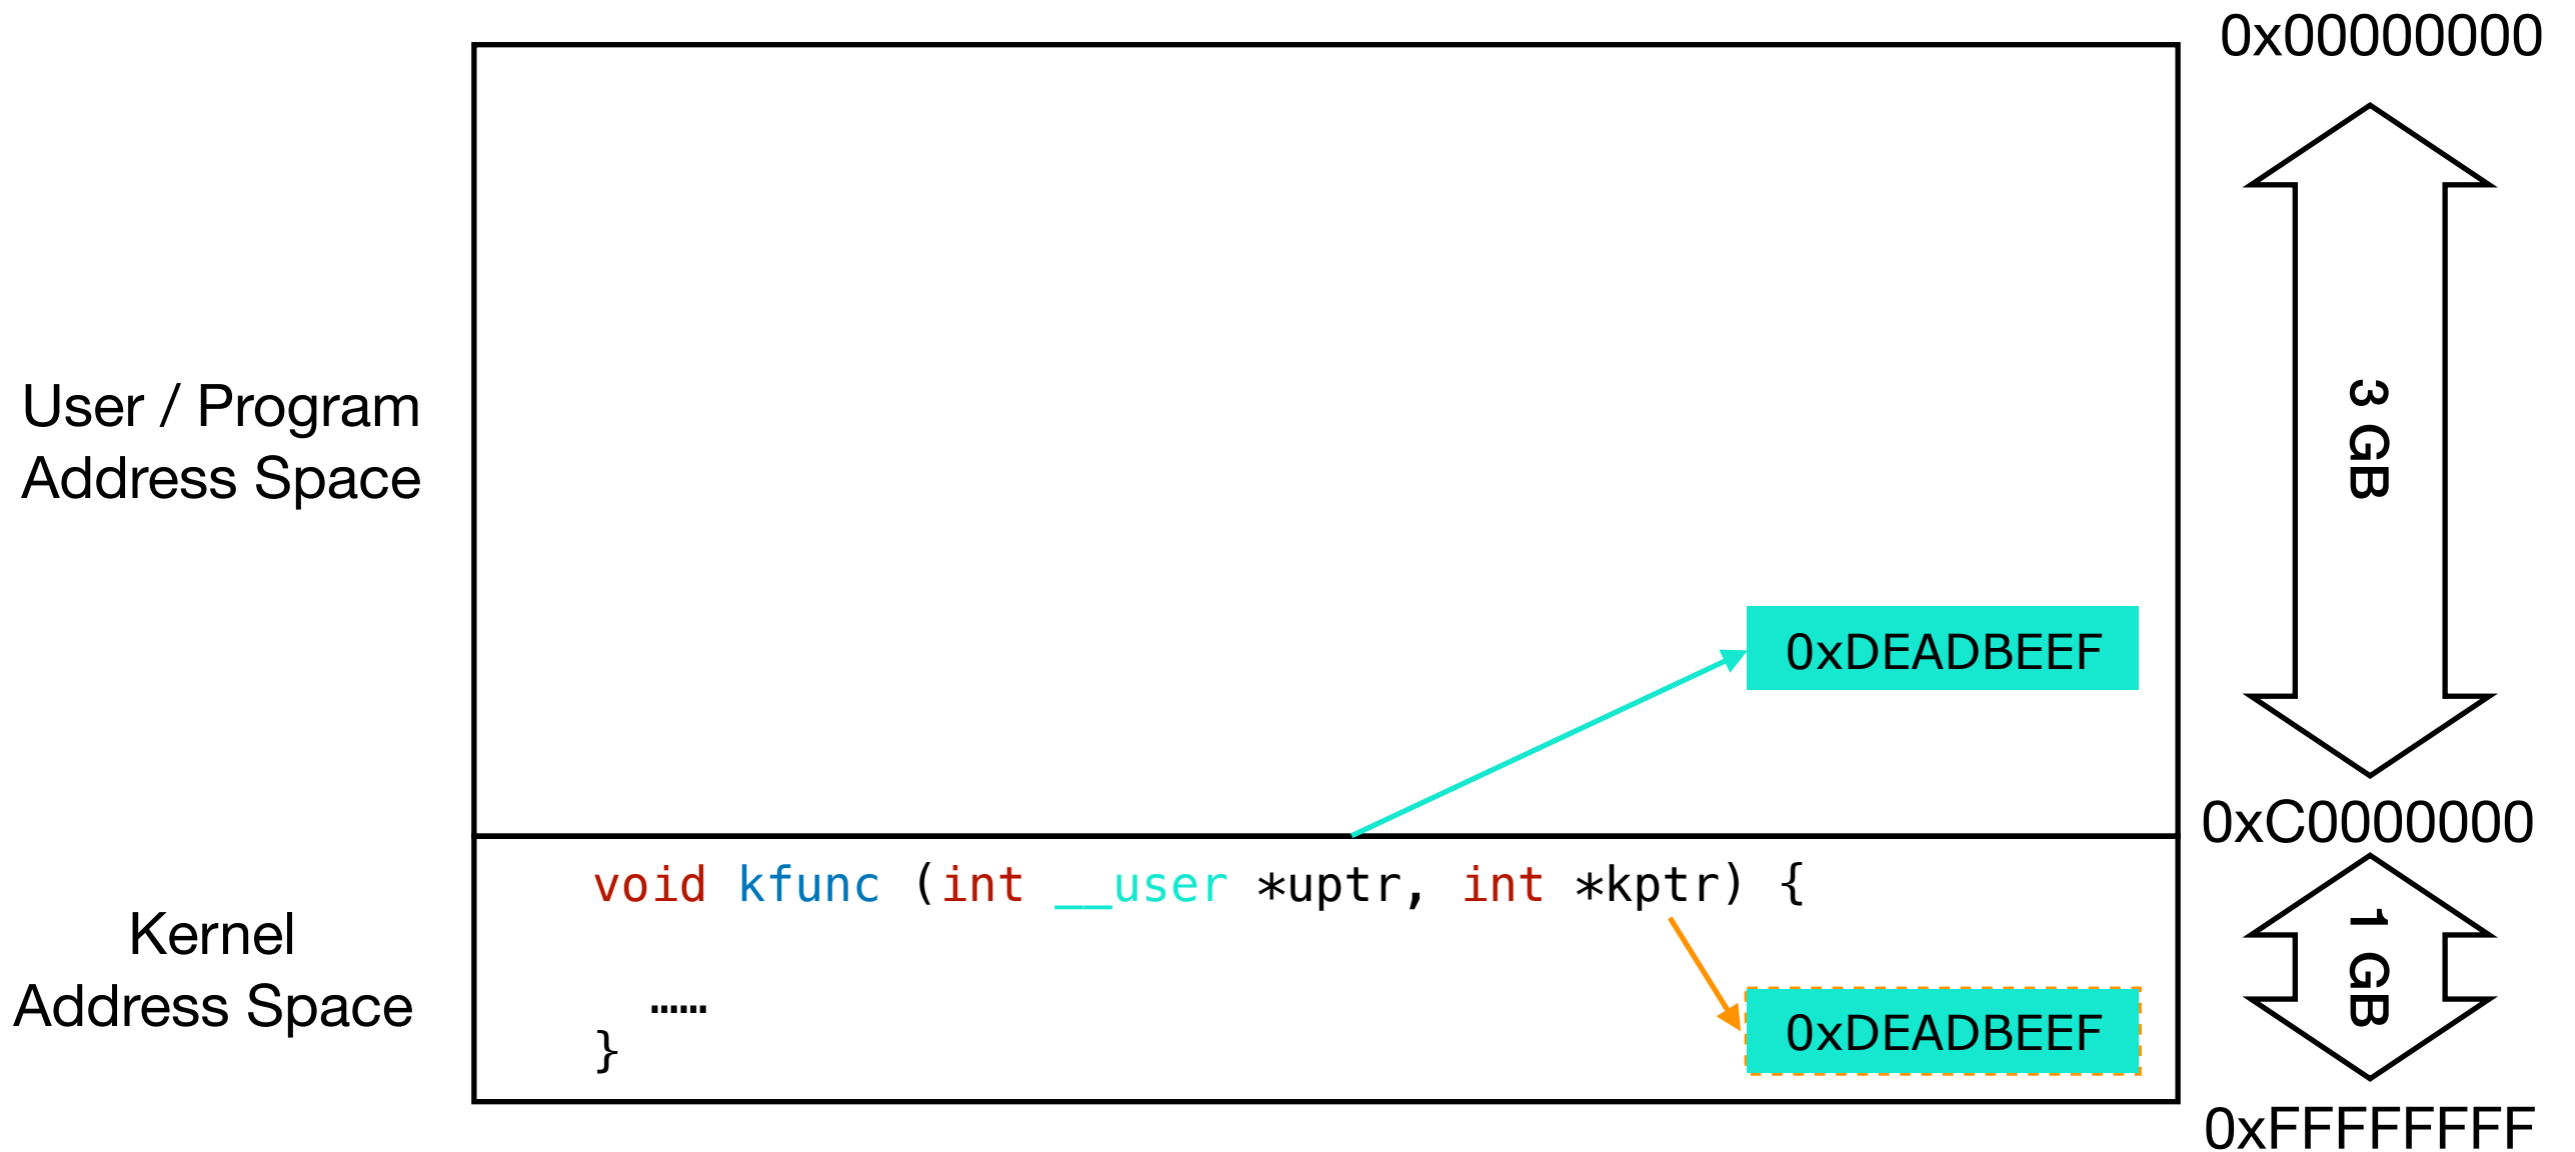
A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

No Dereference on Userspace Pointers



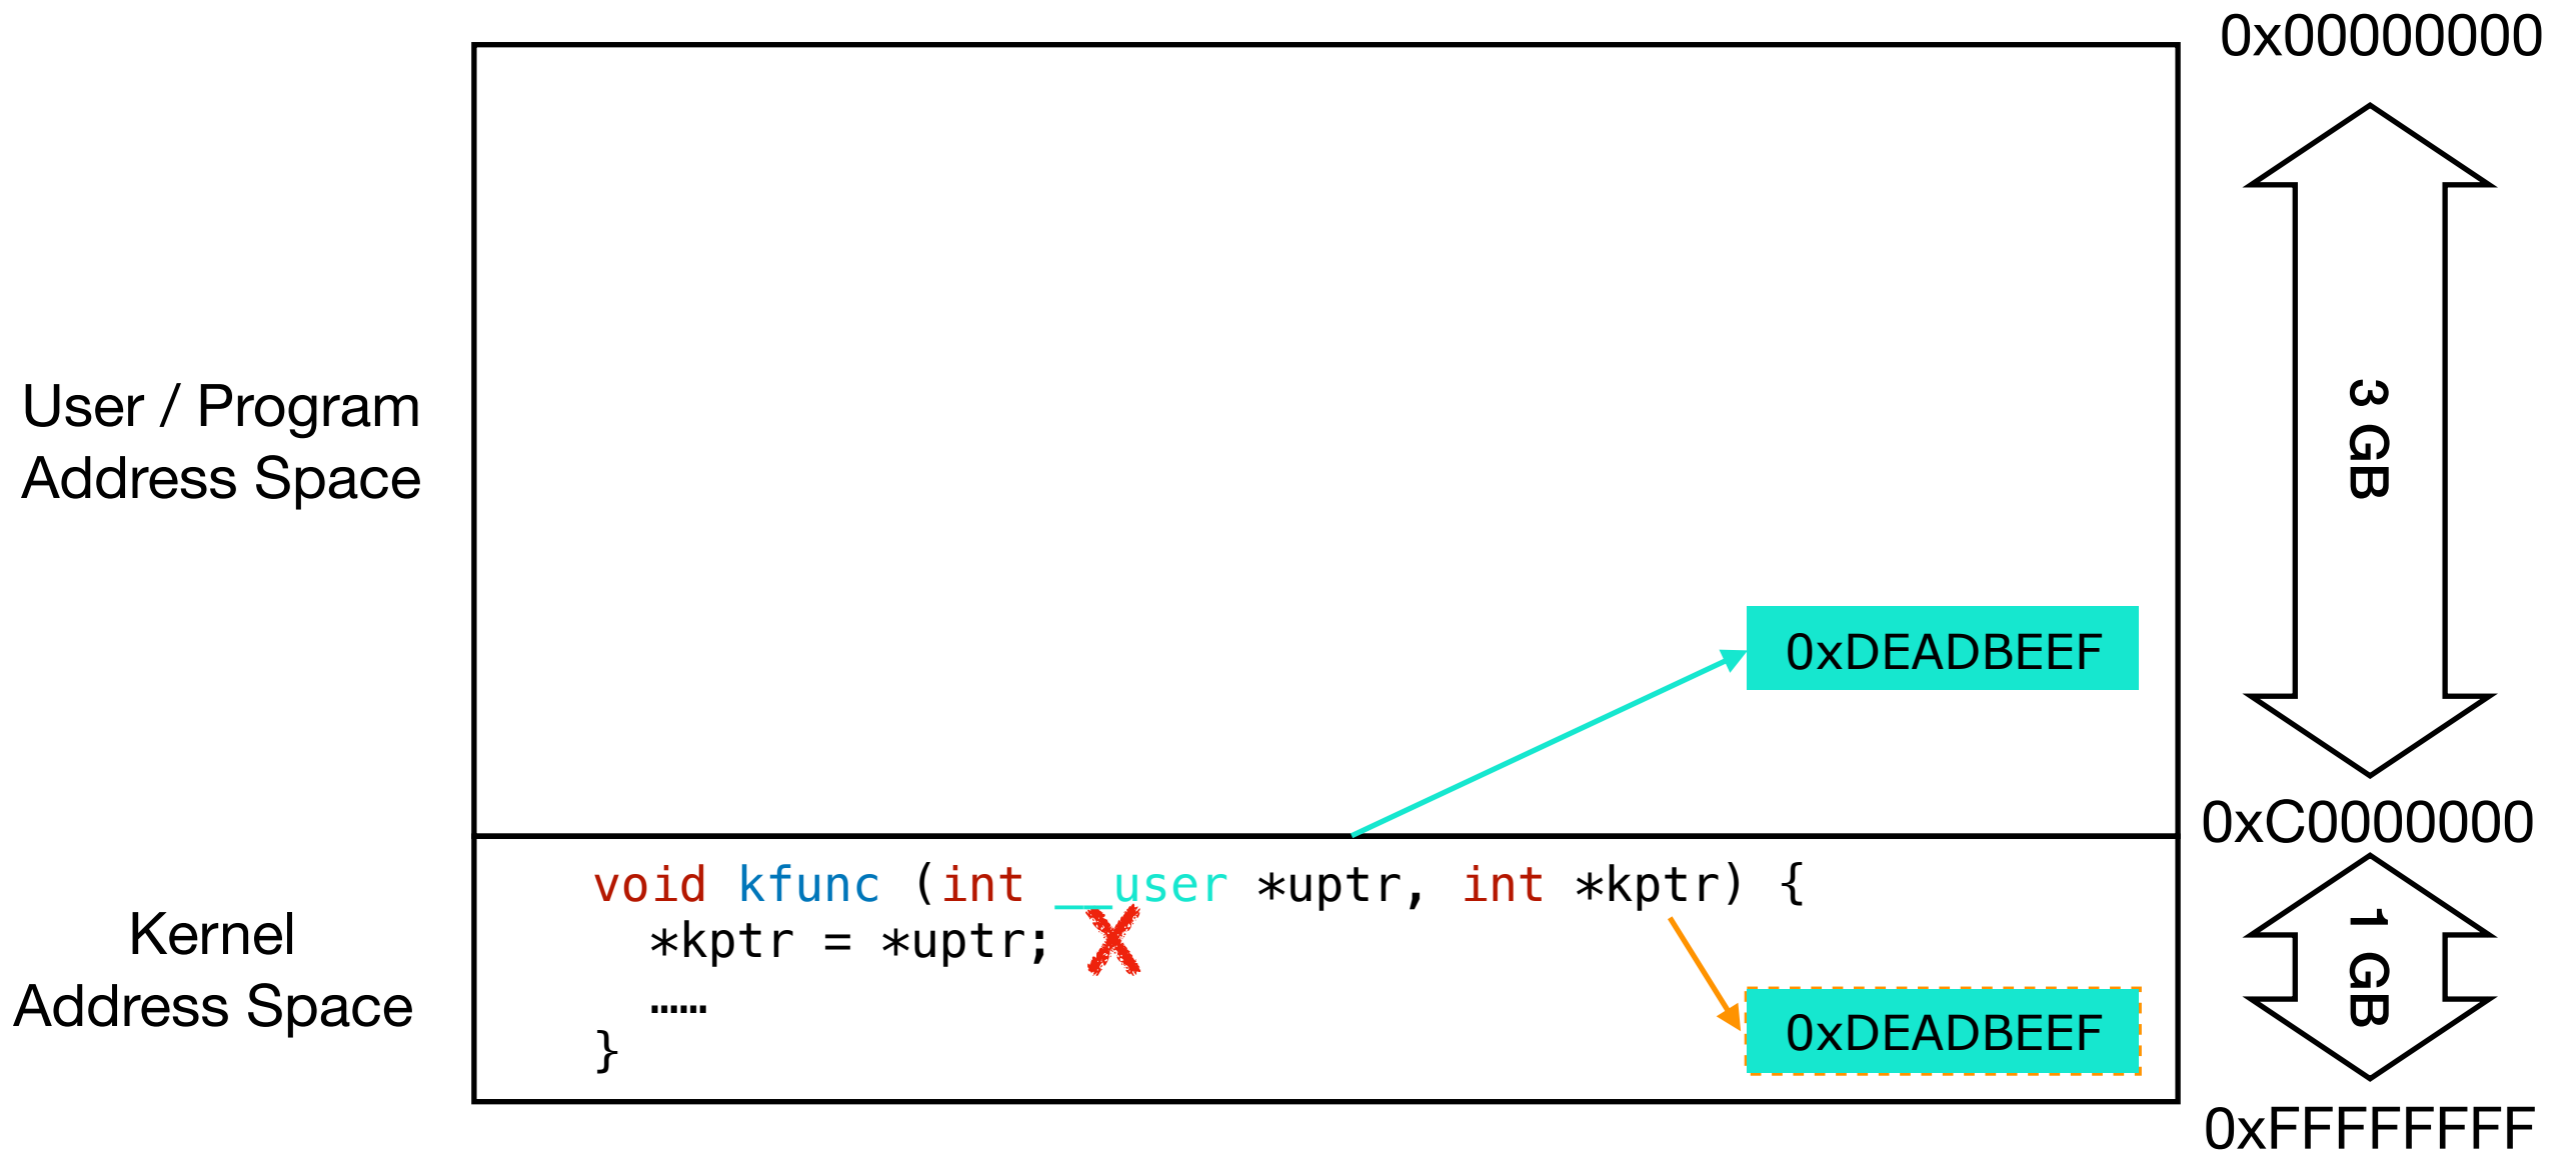
A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

No Dereference on Userspace Pointers



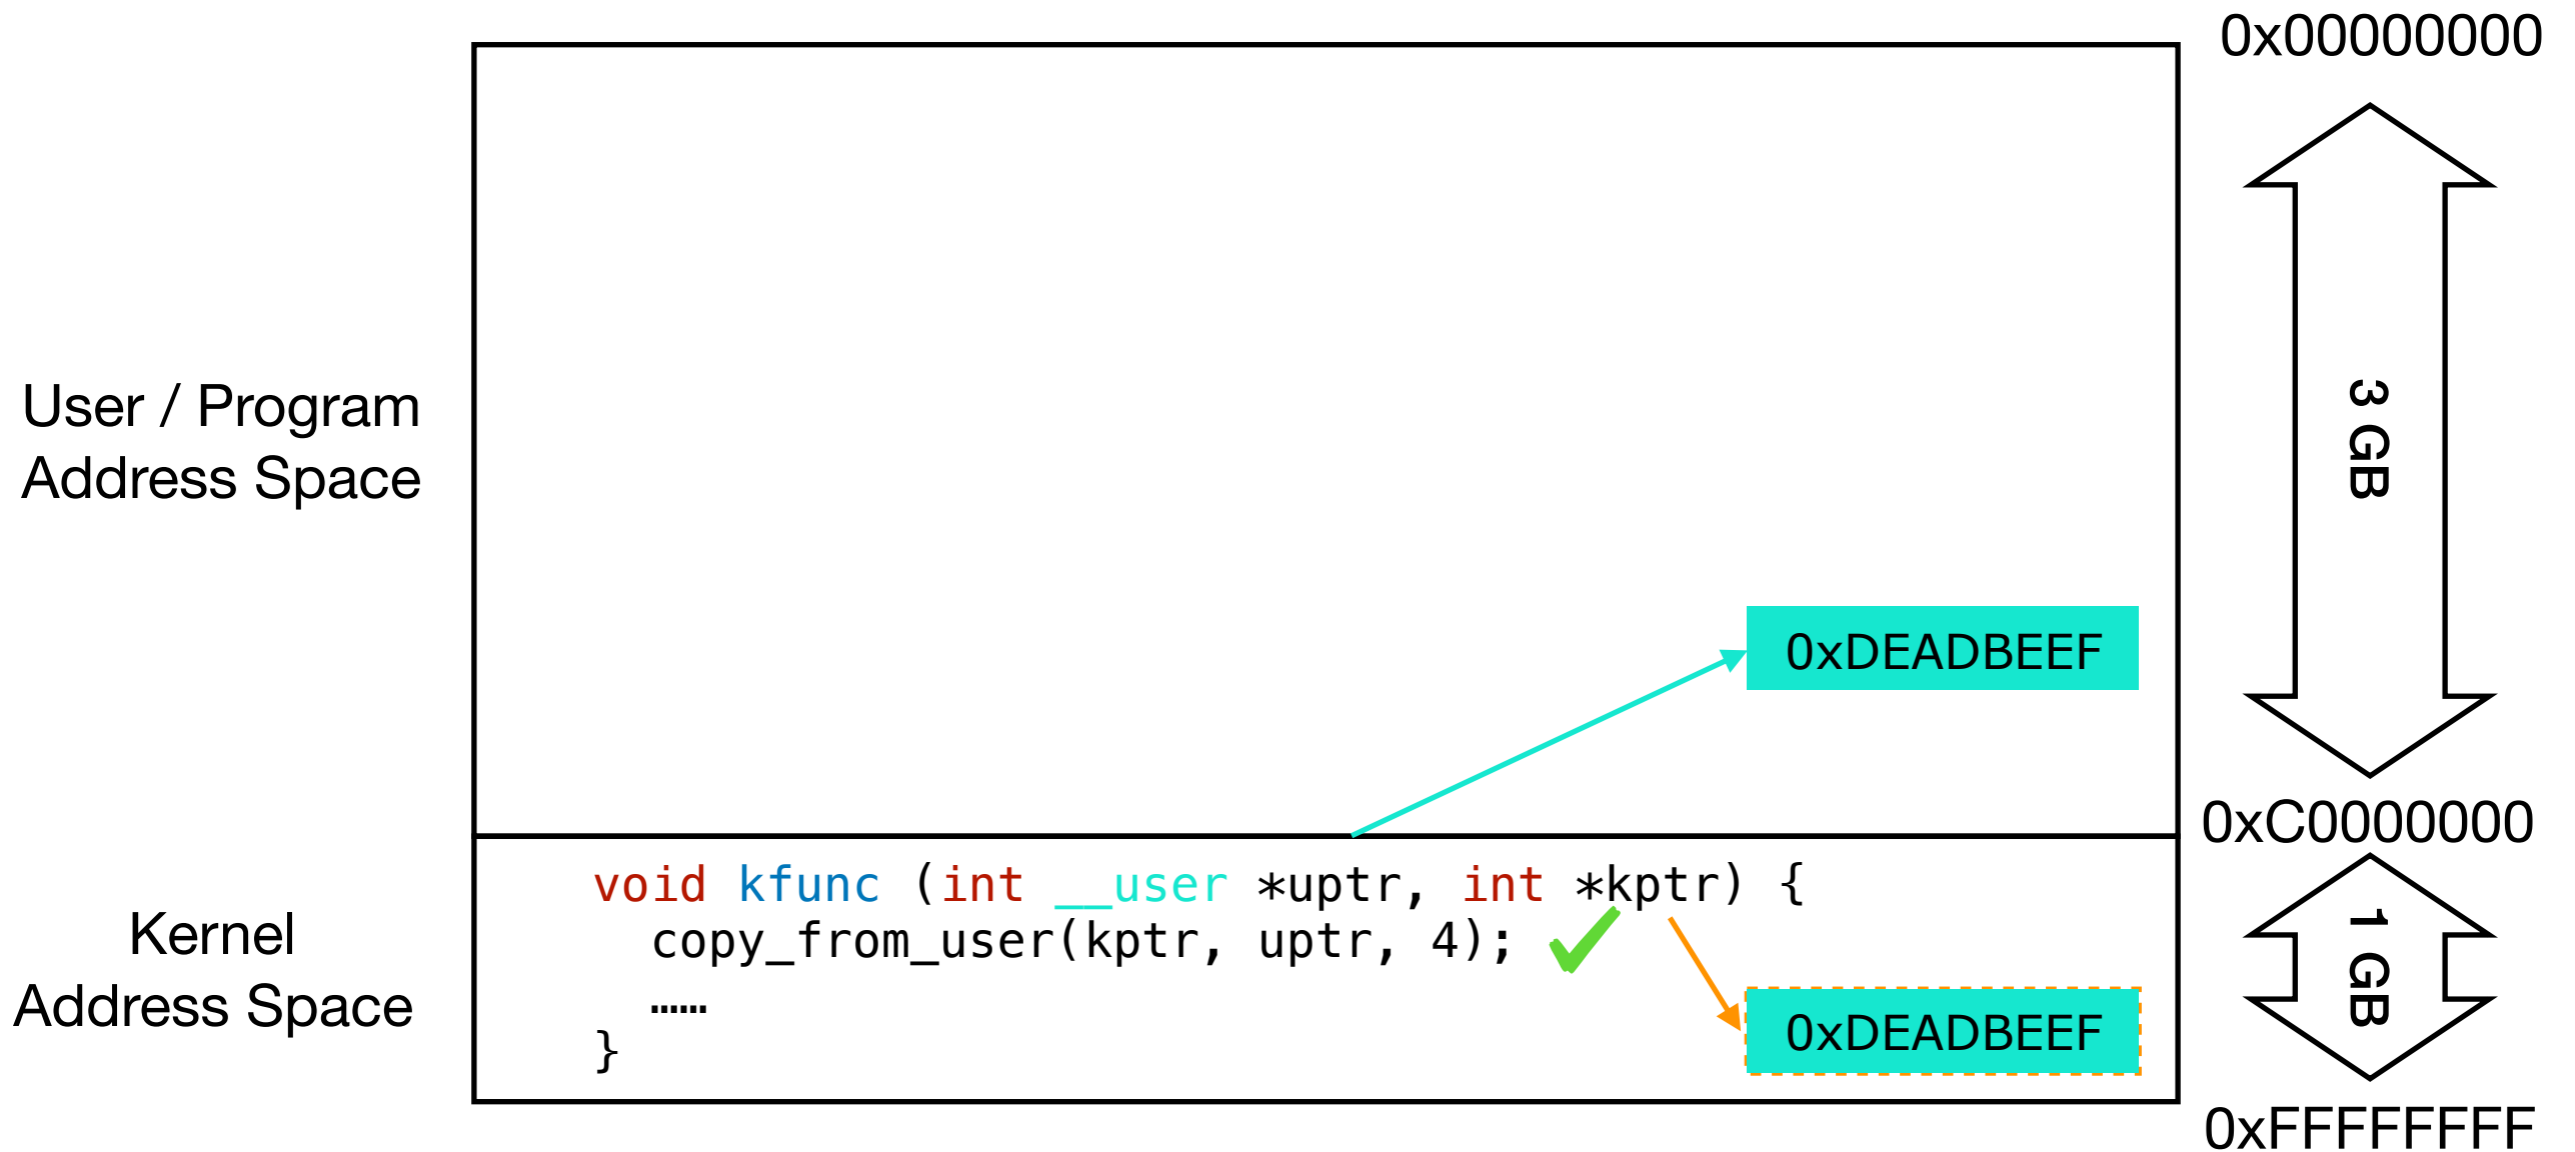
A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

No Dereference on Userspace Pointers



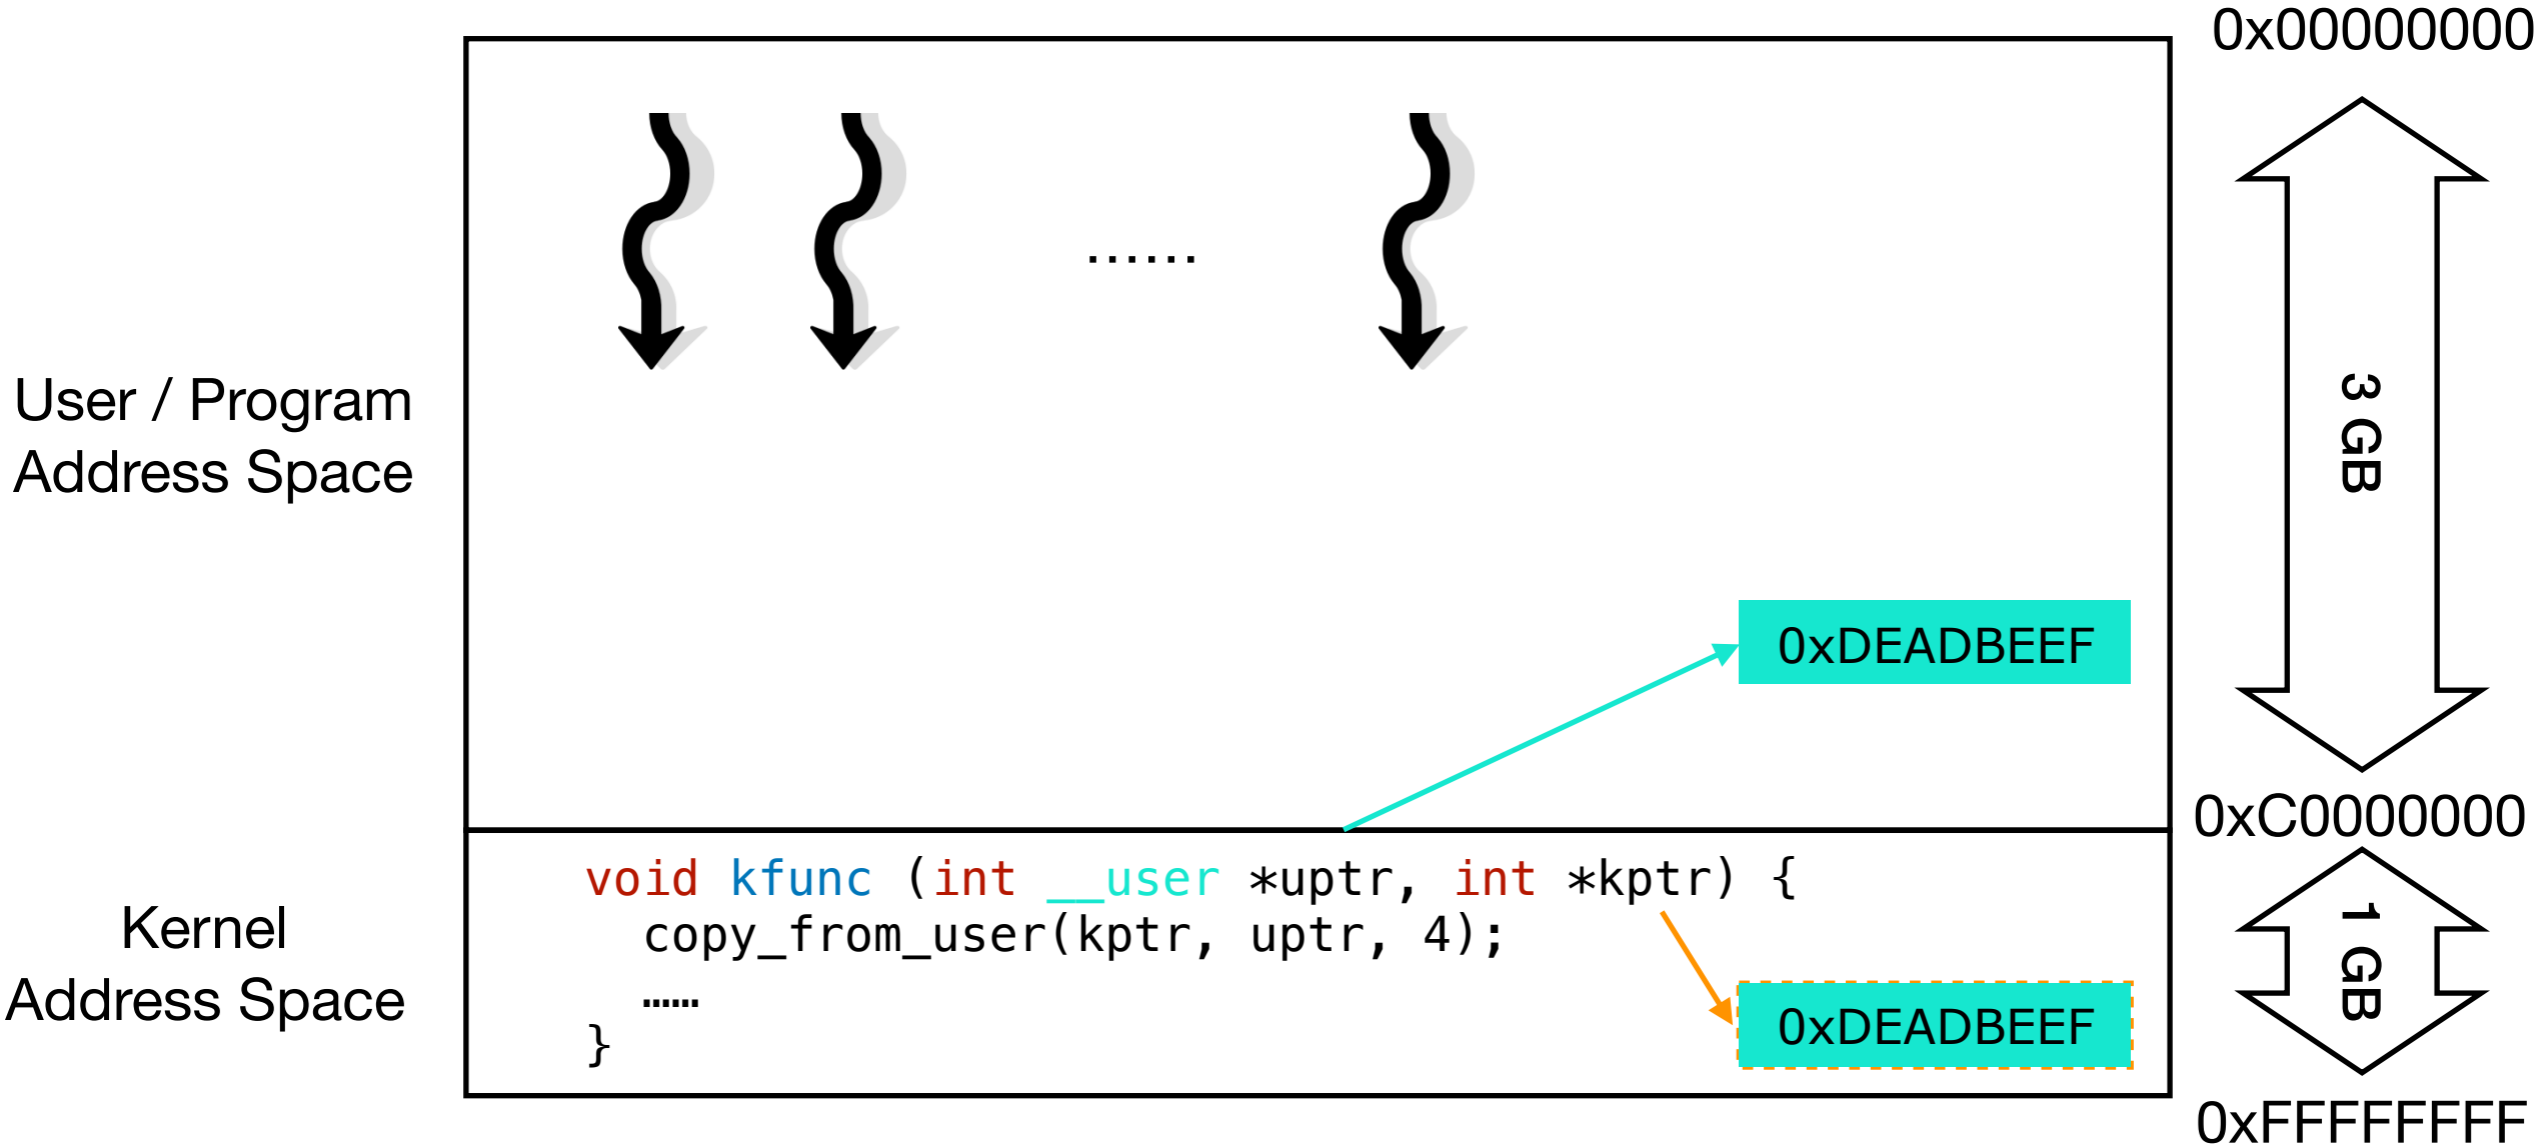
A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

No Dereference on Userspace Pointers



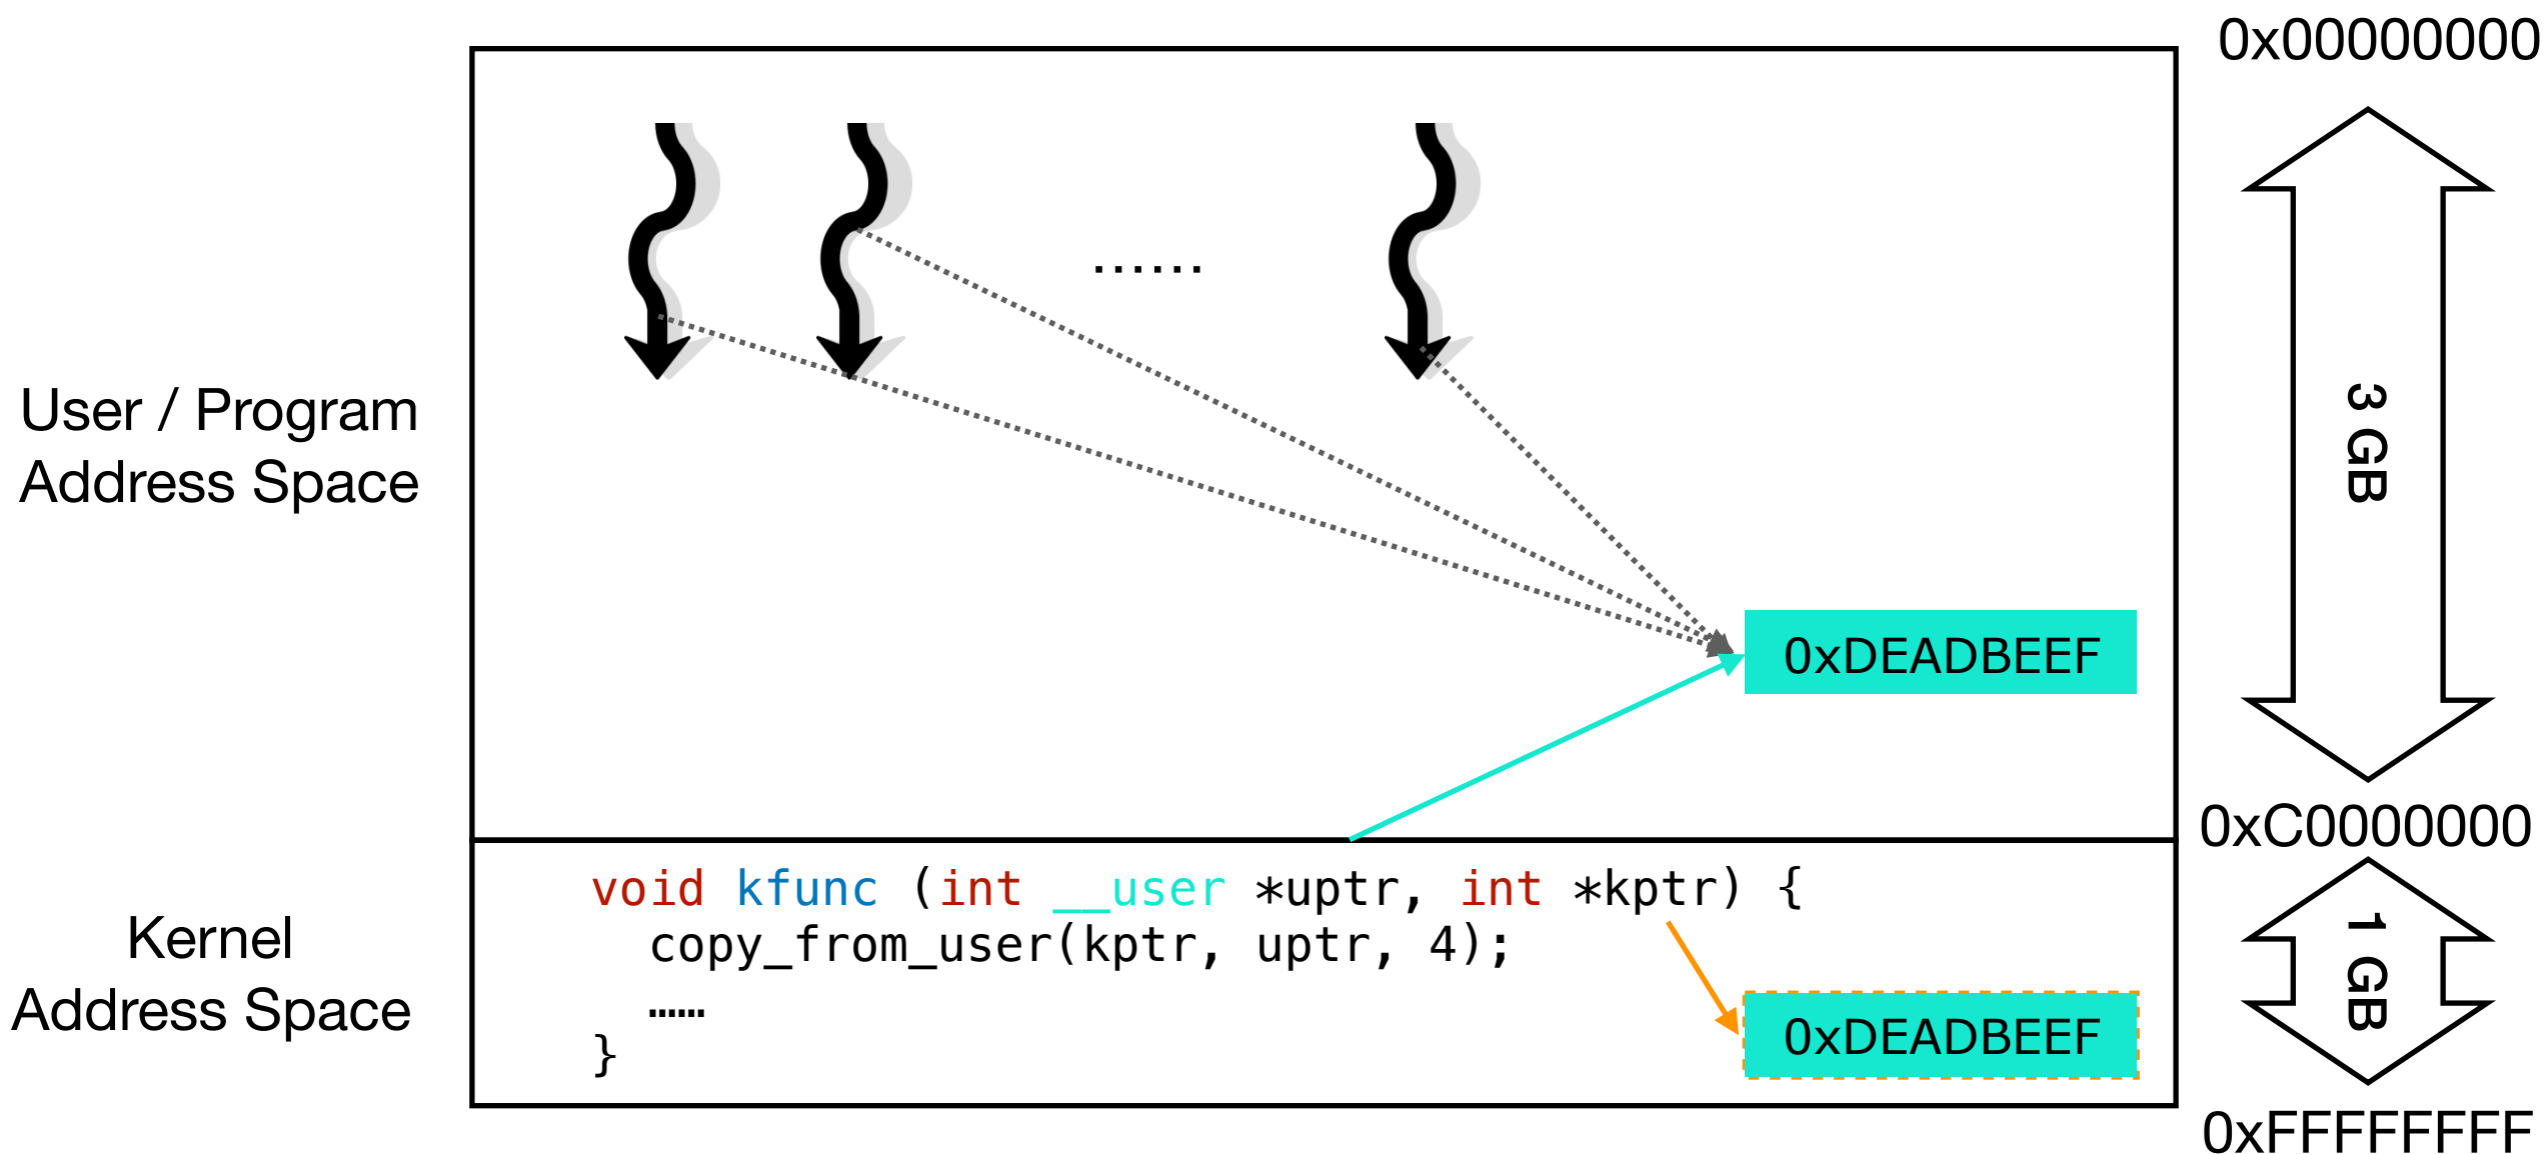
A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

Shared Userspace Pointer Across Threads



A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space


Shared Userspace Pointer Across Threads



A Typical Address Space Separation Scheme with a 32-bit Virtual Address Space

Why Double-Fetch?


```
1 static int perf_copy_attr_simplified  
2 (struct perf_event_attr __user *uattr,  
3  struct perf_event_attr *attr) {
```



Adapted from `perf_copy_attr` in file `kernel/events/core.c`

Why Double-Fetch?


```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
```



Adapted from `perf_copy_attr` in file `kernel/events/core.c`

Why Double-Fetch?

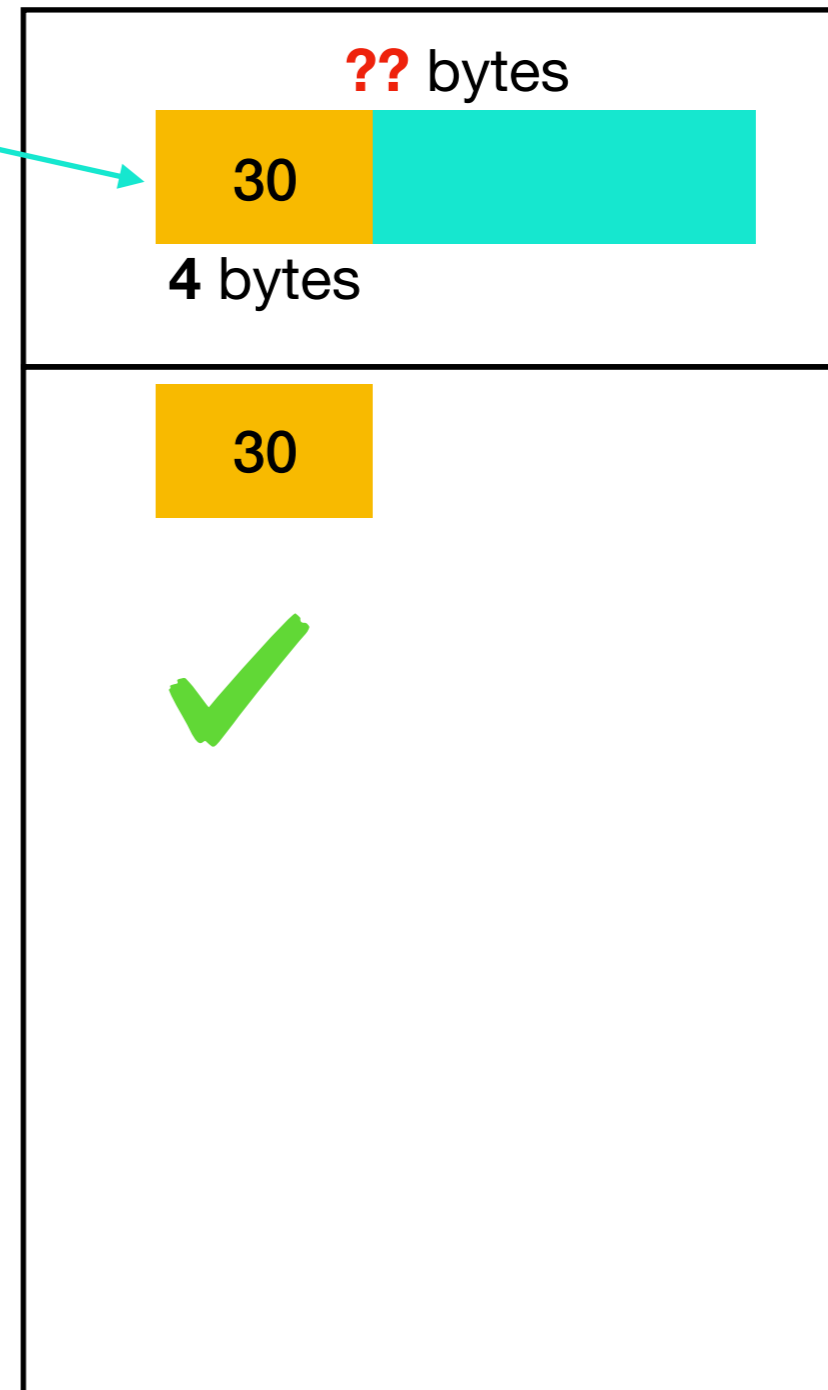
```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
```



Adapted from `perf_copy_attr` in file `kernel/events/core.c`

Why Double-Fetch?

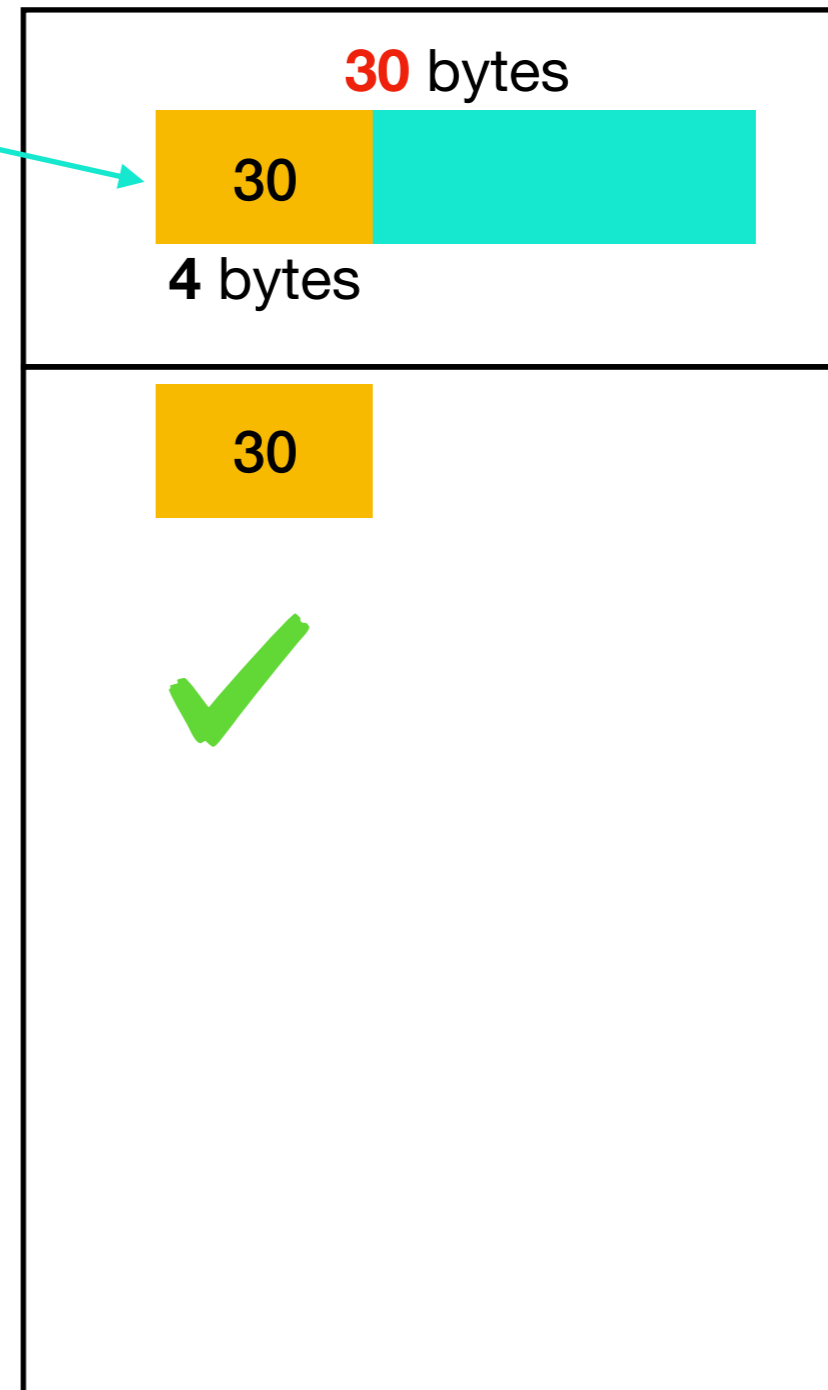
```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
```



Adapted from perf_copy_attr in file kernel/events/core.c

Why Double-Fetch?

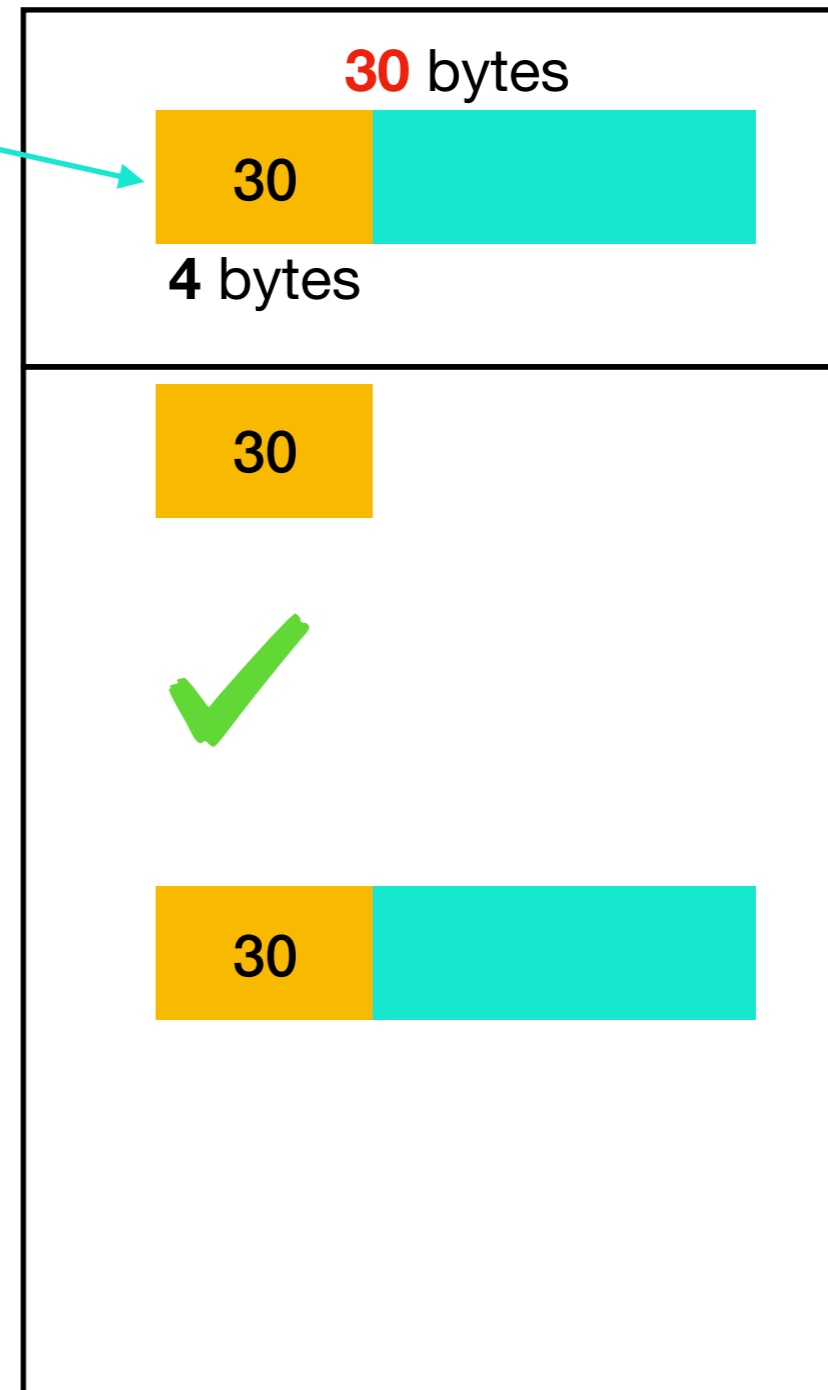
```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
```



Adapted from perf_copy_attr in file kernel/events/core.c

Why Double-Fetch?


```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
```



Adapted from perf_copy_attr in file kernel/events/core.c

Why Double-Fetch?

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
```




Adapted from perf_copy_attr in file kernel/events/core.c

What Goes Wrong in This Process?

Up-until First-Fetch

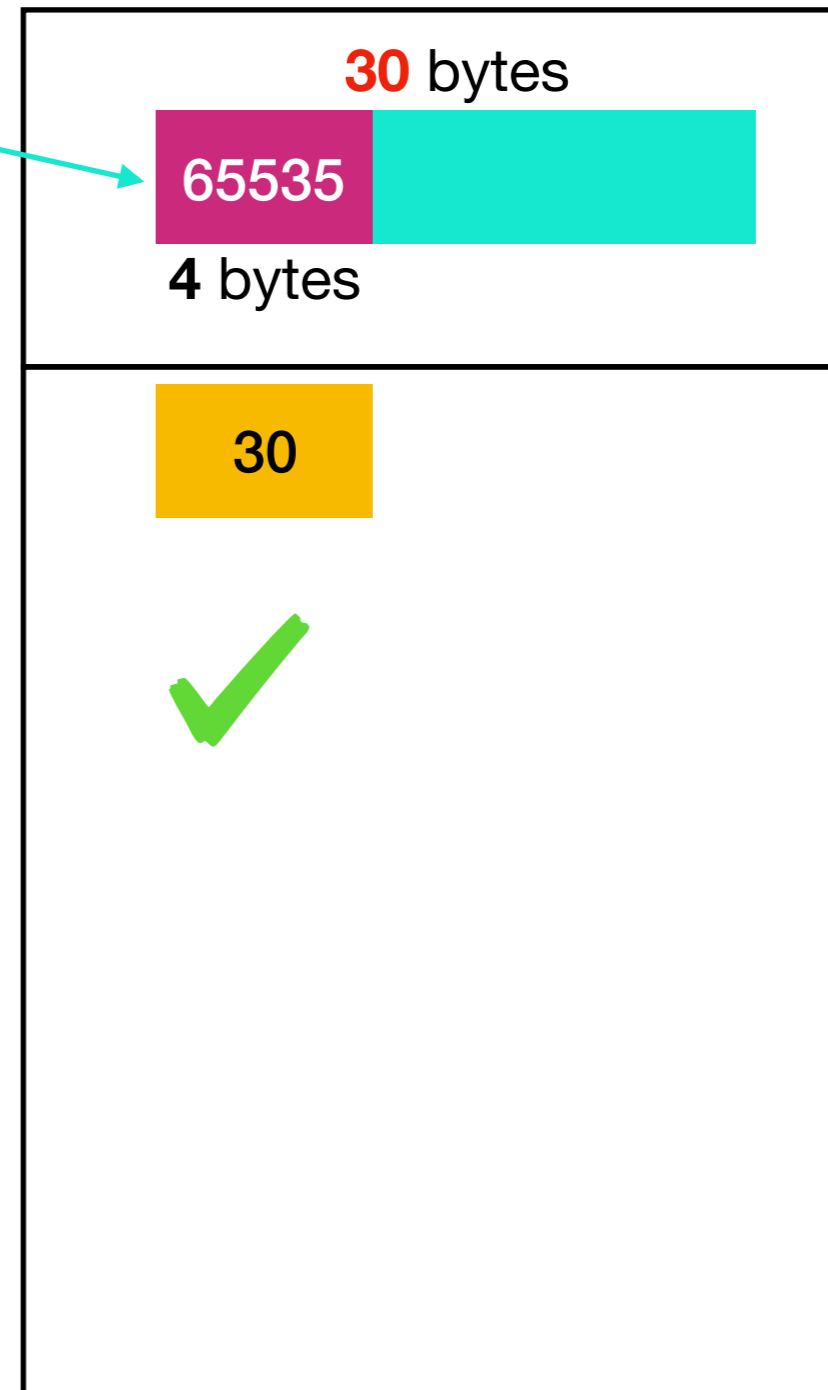
```
1 static int perf_copy_attr_simplified
2 (struct perf_event_attr __user *uattr,
3  struct perf_event_attr *attr) {
4
5  u32 size;
6
7  // first fetch
8  if (get_user(size, &uattr->size))
9      return -EFAULT;
```



Adapted from `perf_copy_attr` in file `kernel/events/core.c`

Wrong Assumption: Atomicity in Syscall


```
1 static int perf_copy_attr_simplified
2 (struct perf_event_attr __user *uattr,
3  struct perf_event_attr *attr) {
4
5  u32 size;
6
7  // first fetch
8  if (get_user(size, &uattr->size))
9      return -EFAULT;
10
11 // sanity checks
12 if (size > PAGE_SIZE ||
13     size < PERF_ATTR_SIZE_VER0)
14     return -EINVAL;
```



Adapted from perf_copy_attr in file kernel/events/core.c

Wrong Assumption: Atomicity in Syscall


```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
```



Adapted from perf_copy_attr in file kernel/events/core.c

When The Exploit Happens

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 copy_to_user(ubuf, attr, attr->size);
```



Adapted from perf_copy_attr in file kernel/events/core.c

Why Double-Fetch is Prevalent in Kernels?

1. Size checking
2. Dependency look-up
3. Protocol/signature check
4. Information guessing
5.

Double-Fetch: Dependency Lookup

```
1 void mptctl_simplified(unsigned long arg) {
2     mpt_ioctl_header khdr, __user *uhdr = (void __user *) arg;
3     MPT_ADAPTER *iocp = NULL;
4
5     // first fetch
6     if (copy_from_user(&khdr, uhdr, sizeof(khdr)))
7         return -EFAULT;
8
9     // dependency lookup
10    if (mpt_verify_adapter(khdr.iocnum, &iocp) < 0 || iocp == NULL)
11        return -EFAULT;
12
13    // dependency usage
14    mutex_lock(&iocp->ioctl_cmds.mutex);
15    struct mpt_fw_xfer kfwdl, __user *ufwdl = (void __user *) arg;
16
17    // second fetch
18    if (copy_from_user(&kfwdl, ufwdl, sizeof(struct mpt_fw_xfer)))
19        return -EFAULT;
20
21    // BUG: kfwdl.iocnum might not equal to khdr.iocnum
22    mptctl_do_fw_download(kfwdl.iocnum, .....);
23    mutex_unlock(&iocp->ioctl_cmds.mutex);
24 }
```

Adapted from `__mptctl_ioctl` in file `drivers/message/fusion/mptctl.c`

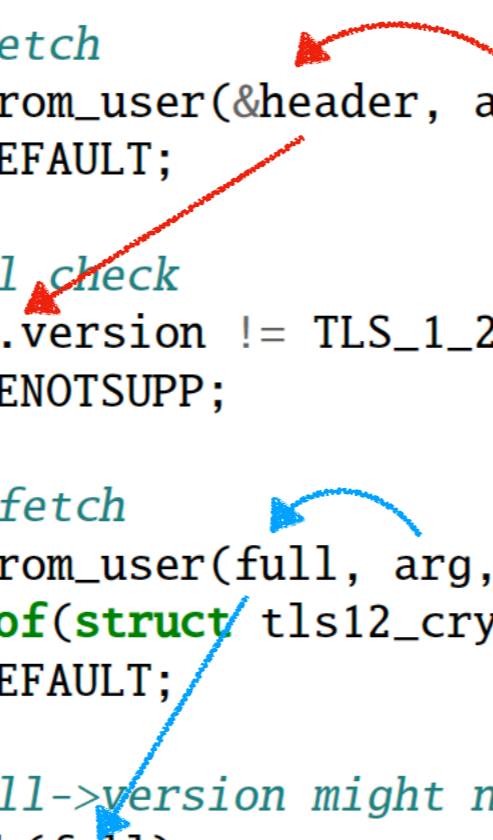
Double-Fetch: Dependency Lookup

```
1 void mptctl_simplified(unsigned long arg) {
2     mpt_ioctl_header khdr, __user *uhdr = (void __user *) arg;
3     MPT_ADAPTER *iocp = NULL;
4
5     // first fetch
6     if (copy_from_user(&khdr, uhdr, sizeof(khdr)))
7         return -EFAULT;
8
9     // dependency lookup
10    if (mpt_verify_adapter(khdr.iocnum, &iocp) < 0 || iocp == NULL)
11        return -EFAULT;
12
13    // dependency usage
14    Acquire mutex lock for ioc 01
15
16
17    // second fetch
18    if (copy_from_user(&kfwdl, ufwdl, sizeof(struct mpt_fw_xfer)))
19        return -EFAULT;
20
21    Do do_fw_download for ioc 02
22
23    Release mutex lock for ioc 01
24 }
```

Adapted from `__mptctl_ioctl` in file `drivers/message/fusion/mptctl.c`

Double-Fetch: Protocol/Signature Check

```
1 void tls_setsockopt_simplified(char __user *arg) {
2     struct tls_crypto_info header, *full = /* allocated before */;
3
4     // first fetch
5     if (copy_from_user(&header, arg, sizeof(struct tls_crypto_info)))
6         return -EFAULT;
7
8     // protocol check
9     if (header.version != TLS_1_2_VERSION)
10        return -ENOTSUPP;
11
12    // second fetch
13    if (copy_from_user(full, arg,
14        sizeof(struct tls12_crypto_info_aes_gcm_128)))
15        return -EFAULT;
16
17    // BUG: full->version might not be TLS_1_2_VERSION
18    do_sth_with(full);
19 }
```



Adapted from do_tls_setsockopt_txZ in file net/tls/tls_main.c

Prior Works

	Bochspwn (BlackHat'13)	DECAF (arXiv'17)	Pengfei et. al., (Security'17)	
Kernel	Windows	Linux	Linux and FreeBSD	
Analysis	Dynamic	Dynamic	Static	
Method	VMI	Kernel fuzzing	Lexical Code Matching	
Patten	Memory access timing	Cache side channel	Size checking	
Code Coverage	Low	Low	High	
Manual Effort	Large	Large	Large	

Prior Works

	Bochspwn (BlackHat'13)	DECAF (arXiv'17)	Pengfei et. al., (Security'17)	Deadline (Our work)
Kernel	Windows	Linux	Linux and FreeBSD	Linux and FreeBSD
Analysis	Dynamic	Dynamic	Static	Static
Method	VMI	Kernel fuzzing	Lexical Code Matching	Symbolic Checking
Patten	Memory access timing	Cache side channel	Size checking	Formal Definitions
Code Coverage	Low	Low	High	High
Manual Effort	Large	Large	Large	Small

Double-Fetch Bugs: Towards A Formal Definition

Fetch: A pair (A, S) , where

A - the starting address of the fetch,

S - the size of memory copied into kernel.

Overlapped-fetch: Two fetches, (A_0, S_0) and (A_1, S_1) , where

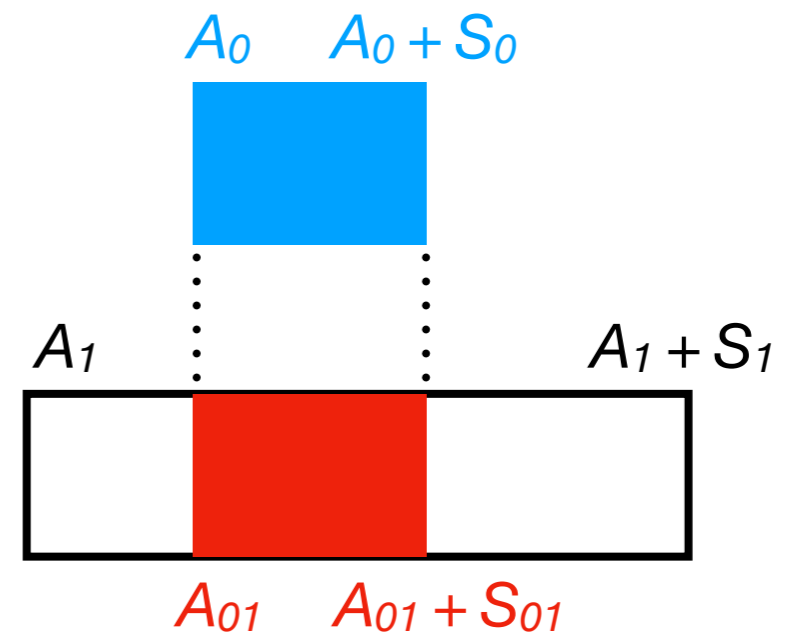
$$A_0 \leq A_1 < A_0 + S_0 \quad || \quad A_1 \leq A_0 < A_1 + S_1$$

- The overlapped memory region is marked as (A_{01}, S_{01}) .
- The copied value during 1st fetch is $(A_{01}, S_{01}, 0)$
- The copied value during 2nd fetch is $(A_{01}, S_{01}, 1)$.

Overlapped-Fetch Case 1

`get_user(attr, &uptr->attr)`

`copy_from_user(kptr, uptr, size)`




$(A_{01}, S_{01}, 0)$ attr

$(A_{01}, S_{01}, 1)$ kptr->attr

Overlapped-Fetch Case 2

```
copy_from_user(  
    khdr, uptr, sizeof(struct hdr)  
)
```

```
copy_from_user(  
    kmsg, uptr, khdr->size  
)
```



$(A_{01}, S_{01}, 0)$ khdr->size, khdr->type, ...

$(A_{01}, S_{01}, 1)$ kmsg->size, kmsg->type, ...

Double-Fetch Bugs: Towards A Formal Definition

Control dependence: A variable $V \in (A_{01}, S_{01})$ and V must satisfy a set of constraints before the second fetch can happen.

Double-Fetch Bugs: Towards A Formal Definition

Control dependence: A variable $V \in (A_{01}, S_{01})$ and V must satisfy a set of constraints before the second fetch can happen.

```
1 void tls_setsockopt_simplified(char __user *arg) {
2     struct tls_crypto_info header, *full = /* allocated before */;
3
4     // first fetch
5     if (copy_from_user(&header, arg, sizeof(struct tls_crypto_info)))
6         return -EFAULT;
7
8     // protocol check
9     if (header.version != TLS_1_2_VERSION)
10        return -ENOTSUPP;
11
12    // second fetch
13    if (copy_from_user(full, arg,
14        sizeof(struct tls12_crypto_info_aes_gcm_128)))
15        return -EFAULT;
16
17    // BUG: full->version might not be TLS_1_2_VERSION
18    do_sth_with(full);
19 }
```

Overlapped variable V :
header.version

The constraint it must satisfy:
header.version == TLS_1_2_VERSION

Expect:
full->version == TLS_1_2_VERSION

Double-Fetch Bugs: Towards A Formal Definition

Data dependence: A variable $V \in (A_{01}, S_{01})$ and V is consumed before or on the second fetch (e.g., involved in calculation, passed to function calls, etc).

Double-Fetch Bugs: Towards A Formal Definition

Data dependence: A variable $V \in (A_{01}, S_{01})$ and V is consumed before or on the second fetch.

```
1 void mptctl_simplified(unsigned long arg) {
2   mpt_ioctl_header khdr, __user *uhdr = (void __user *) arg;
3   MPT_ADAPTER *iocp = NULL;
4
5   // first fetch
6   if (copy_from_user(&khdr, uhdr, sizeof(khdr)))
7     return -EFAULT;
8
9   // dependency lookup
10  if (mpt_verify_adapter(khdr.iocnum, &iocp) <
11    return -EFAULT;
12
13  // dependency usage
14  mutex_lock(&iocp->ioctl_cmds.mutex);
15  struct mpt_fw_xfer kfwdl, __user *ufwdl = (void __user *) arg;
16
17  // second fetch
18  if (copy_from_user(&kfwdl, ufwdl, sizeof(struct mpt_fw_xfer)))
19    return -EFAULT;
20
21  // BUG: kfwdl.iocnum might not equal to khdr.iocnum
22  mptctl_do_fw_download(kfwdl.iocnum, .....);
23  mutex_unlock(&iocp->ioctl_cmds.mutex);
24 }
```

Overlapped variable V :
khdr.iocnum

Data dependence:
mpt_verify_adapter(khdr.iocnum, &iocp)

Expect:
kfwdl.iocnum == khdr.iocnum

Double-Fetch Bugs: Towards A Formal Definition

1. Two fetches from userspace memory that cover an **overlapped** region.
2. A relation must exist on the overlapped region between the two fetches. The relation can be either **control-dependence** or **data-dependence**.
3. We cannot **prove** that the relation established after first fetch still holds after the second fetch.

If all conditions are satisfied: a user thread might race condition to change the content in the overlapped region, and thus, to destroy the relation.

How to Find Double-Fetch Bugs?

How to Find Double-Fetch Bugs?

1. Find as many double-fetch pairs as possible, construct the code paths associated with each pair.
2. Symbolically check each code path and determine whether the two fetches makes a double-fetch bug.

Fetch Pair Collection

Goal: Statically enumerate all pairs of fetches that could possibly occur.

Fetch Pairs Collection

Goal: Statically enumerate all pairs of fetches that could possibly occur.

Ideal solution (top-down):

- ✓ 1. Identify all fetches in the kernel
- ✗ 2. Construct a complete, inter-procedural CFG for the whole kernel
- ✗ 3. Perform pair-wise reachability tests for each pair of fetches

Fetch Pairs Collection

Goal: Statically enumerate all pairs of fetches that could possibly occur.

Ideal solution (top-down):

- ✓ 1. Identify all fetches in the kernel
- ✗ 2. Construct a complete, inter-procedural CFG for the whole kernel
- ✗ 3. Perform pair-wise reachability tests for each pair of fetches

Our solution (bottom-up):

- ✓ 1. Identify all fetches in the kernel
- ✓ 2. For each fetch, within the function it resides in, scan its reaching instructions for fetches or fetch-involved functions

Bottom-up Fetch Pairs Collection

```
static void enclosing_function(  
    struct msg_hdr __user *uptr,  
    struct msg_full *kptr  
) {  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    if (copy_from_user(kptr, uptr, size))  
        return -EFAULT;  
    ...  
}
```

Start from a fetch →

Bottom-up Fetch Pairs Collection

```
static void enclosing_function(  
    struct msg_hdr __user *uptr,  
    struct msg_full *kptr  
) {  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    if (copy_from_user(kptr, uptr, size))  
        return -EFAULT;  
    ...  
}
```

Search through the
reaching instructions



Bottom-up Fetch Pairs Collection

[Case 1]
Found another fetch
==>
found a fetch pair

```
static void enclosing_function(  
    struct msg_hdr __user *uptr,  
    struct msg_full *kptr  
) {  
    ...  
    ...  
    if (get_user(size, &uptr->size))  
        return -EFAULT;  
    ...  
    ...  
    if (copy_from_user(kptr, uptr, size))  
        return -EFAULT;  
    ...  
}
```

Bottom-up Fetch Pairs Collection

[Case 2]
Found a fetch-involved
function
==>
inline the function,
found a fetch pair

```
static void enclosing_function(  
    struct msg_hdr __user *uptr,  
    struct msg_full *kptr  
) {  
    ...  
    ...  
    size = get_size_from_user(uptr);  
    ...  
    ...  
    ...  
    ...  
    if (copy_from_user(kptr, uptr, size))  
        return -EFAULT;  
    ...  
}
```

Bottom-up Fetch Pairs Collection

[Case 3]
No fetch-related
instruction
==>
Not a double-fetch



```
static void enclosing_function(  
    struct msg_hdr __user *uptr,  
    struct msg_full *kptr  
) {  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    ...  
    if (copy_from_user(kptr, uptr, size))  
        return -EFAULT;  
    ...  
}
```

How to Find Double-Fetch Bugs?

- ✓ 1. Find as many double-fetch pairs as possible, construct the code paths associated with each pair.
2. Symbolically check each code path and determine whether the two fetches makes a double-fetch bug.

Symbolic Checking

Goal: Symbolically execute the code path that connects two fetches and determine whether the two fetches satisfy all the criteria set in formal definition of double-fetch bug, i.e.

- Overlapp
- Have a relation (control or data dependence)
- We cannot prove the relation still holds after second fetch

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```


Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
```

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
5 // first fetch
6 fetch(F1): {A = $0 + 4, S = 4}
7 $2 = @0(4, 7, U0), @2 = nil // size
8 ---
```

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
5 // first fetch
6 fetch(F1): {A = $0 + 4, S = 4}
7 $2 = @0(4, 7, U0), @2 = nil // size
8 ---
9 // sanity checks
10 assert $2 <= PAGE_SIZE
11 assert $2 >= PERF_ATTR_SIZE_VER0
12 ---
```

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
5 // first fetch
6 fetch(F1): {A = $0 + 4, S = 4}
7 $2 = @0(4, 7, U0), @2 = nil // size
8 ---
9 // sanity checks
10 assert $2 <= PAGE_SIZE
11 assert $2 >= PERF_ATTR_SIZE_VER0
12 ---
13 // second fetch
14 fetch(F2): {A = $0, S = $2}
15 @1(0, $2 - 1, K) = @0(0, $2 - 1, U1)
16 ---
```

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9     return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14    return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18    return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
5 // first fetch
6 fetch(F1): {A = $0 + 4, S = 4}
7 $2 = @0(4, 7, U0), @2 = nil // size
8 ---
9 // sanity checks
10 assert $2 <= PAGE_SIZE
11 assert $2 >= PERF_ATTR_SIZE_VER0
12 ---
13 // second fetch
14 fetch(F2): {A = $0, S = $2}
15 @1(0, $2 - 1, K) = @0(0, $2 - 1, U1)
16 ---
17 // check fetch overlap
18 assert F2.A <= F1.A < F2.A + F2.S
19        OR F1.A <= F2.A < F1.A + F1.S
20 [solve]
21 --> satisfiable with @0(4, 7, U)
```

Symbolic Checking

```
1 static int perf_copy_attr_simplified
2   (struct perf_event_attr __user *uattr,
3    struct perf_event_attr *attr) {
4
5   u32 size;
6
7   // first fetch
8   if (get_user(size, &uattr->size))
9       return -EFAULT;
10
11  // sanity checks
12  if (size > PAGE_SIZE ||
13      size < PERF_ATTR_SIZE_VER0)
14      return -EINVAL;
15
16  // second fetch
17  if (copy_from_user(attr, uattr, size))
18      return -EFAULT;
19
20  .....
21 }
22
23 // BUG: when attr->size is used later
24 memcpy(buf, attr, attr->size);
```

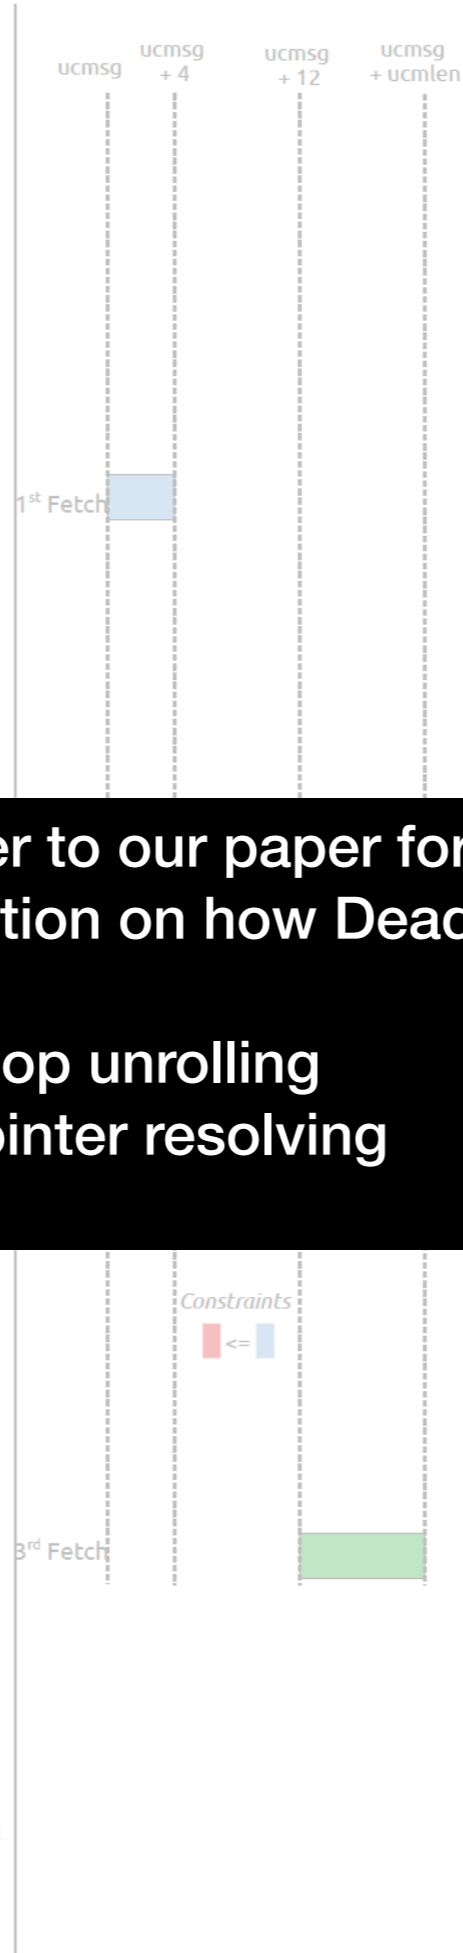
```
1 // init root SR
2 $0 = PARM(0), @0 = UMEM(0) // uattr
3 $1 = PARM(1), @1 = KMEM(1) // attr
4 ---
5 // first fetch
6 fetch(F1): {A = $0 + 4, S = 4}
7 $2 = @0(4, 7, U0), @2 = nil // size
8 ---
9 // sanity checks
10 assert $2 <= PAGE_SIZE
11 assert $2 >= PERF_ATTR_SIZE_VER0
12 ---
13 // second fetch
14 fetch(F2): {A = $0, S = $2}
15 @1(0, $2 - 1, K) = @0(0, $2 - 1, U1)
16 ---
17 // check fetch overlap
18 assert F2.A <= F1.A < F2.A + F2.S
19      OR F1.A <= F2.A < F1.A + F1.S
20 [solve]
21   --> satisfiable with @0(4, 7, U)
22 // check double-fetch bug
23 [prove] @0(4, 7, U0) == @0(4, 7, U1)
24   --> fail: no constraints on @0(4, 7, U1)
```

```

1 int cmsghdr_from_user_compat_to_kern
2 (struct msg_hdr *kmsg, char *kbuf) {
3
4 struct compat_cmsghdr __user *ucmsg;
5 compat_size_t ucmlen;
6 struct cmsghdr *kcmsg;
7 __kernel_size_t kcmlen, tmp;
8
9 // 1st loop: calculate message length
10 kcmlen = 0;
11 ucmsg = kmsg->msg_control;
12 while (ucmsg != NULL) {
13 // first batch of fetches
14 if (get_user(ucmlen, &ucmsg->cmsg_len))
15 return -EFAULT;
16
17 tmp = ucmlen + sizeof(struct cmsghdr)
18 - sizeof(struct compat_cmsghdr);
19
20 kcmlen += tmp;
21 ucmsg = (char *)ucmsg + ucmlen;
22 }
23
24 // 2nd loop: copy the whole message
25 kcmsg = kbuf;
26 ucmsg = kmsg->msg_control;
27 while (ucmsg != NULL) {
28 // second batch of fetches
29 if (get_user(ucmlen, &ucmsg->cmsg_len))
30 return -EFAULT;
31
32 tmp = ucmlen + sizeof(struct cmsghdr)
33 - sizeof(struct compat_cmsghdr);
34
35 // sanity check, but insufficient
36 if (kbuf + kcmlen - (char *)kcmsg < tmp)
37 return -EINVAL;
38
39 // irrelevant fetch
40 if (copy_from_user(
41 (char *)kcmsg + sizeof(*kcmsg),
42 (char *)ucmsg + sizeof(*ucmsg),
43 (ucmlen - sizeof(*ucmsg))))
44 return -EFAULT;
45
46 kcmsg = (char *)kcmsg + tmp;
47 ucmsg = (char *)ucmsg + ucmlen;
48 }
49
50 // BUG: the actual message length != kcmlen
51 kmsg->msg_controllen = kcmlen;
52 return 0;
53 }

```

(a) C source code



(b) Memory access patterns

```

1 // init root SR
2 $0 = $PARM(0), @0 = $KMEM(0) // kmsg
3 $1 = $PARM(1), @1 = $KMEM(1) // kbuf
4 ---
5 // prepare for the 1st batch of fetches
6 $2 = 0, @2 = nil // kcmlen_0
7 $3 = @0(48, 55, K), @3 = $UMEM(0) // ucmsg_0
8 ---
9 // unroll 1st loop
10 assert $2 != NULL
11 fetch(F1) is {A = $3 + 0, S = 4}
12 $4 = @3(0, 3, U0), @4 = nil // ucmlen_0
13 $5 = $4 - 12 + 16, @5 = nil // tmp_0
14 $6 = $2 + $5, @6 = nil // kcmlen_1
15 $7 = $3 + $4, @7 = $UMEM(1) // ucmsg_1
16 assert $7 == NULL (i.e., @7 = nil) // exit loop
17 ---
18 // prepare for the 2nd batch of fetches
19 $8 = $1 @8 = $KMEM(1) // kcmsg_0
20 $9 = @0(48, 55, K) == $3, @9 = @3 // ucmsg_2
21 ---
22 // unroll 2nd loop
23
24 }
25 @0 = nil // ucmlen_1
26 @1 = nil // tmp_1
27
28 > @3(0, 3, U0) >= @3(0, 3, U1)
29
30 $10 - 12}
31 @10 - 13, U0)
32
33 @2 = $KMEM(2) // kcmsg_1
34 @3 = $UMEM(3) // ucmsg_3
35
36 assert $13 == NULL (i.e., @13 = nil) // exit loop
37 ---
38 // check fetch overlap
39 assert F2.A <= F1.A < F2.A + F2.S
40 AND F1.A <= F2.A < F1.A + F1.S
41 // --> satisfiable with @3(0, 3, U)
42
43 assert F3.A <= F1.A < F3.A + F3.S
44 AND F1.A <= F3.A < F1.A + F1.S
45 // --> unsatisfiable
46
47 assert F3.A <= F2.A < F3.A + F3.S
48 AND F2.A <= F3.A < F2.A + F2.S
49 // --> unsatisfiable
50
51 // check double-fetch bug
52 prove @3(0, 3, U0) == @3(0, 3, U1)
53 // --> fail, as @3(0, 3, U0) >= @3(0, 3, U1)

```

(c) Symbolic representation and checking

Please refer to our paper for a comprehensive demonstration on how Deadline handles

1. Loop unrolling
2. Pointer resolving

Findings

- 24 bugs found in total
 - 23 bugs in Linux kernel and 1 in FreeBSD kernel
- 9 bugs have been patched with the fix we provide
- 4 bugs are acknowledged, we are still working on the fix
- 9 bugs are pending for review
- 2 bugs are marked as “won’t fix”

Double-Fetch Bug Mitigations

- The basic idea is to re-assure the control-dependence and data-dependence between the two fetches. In other words, the **automaticity** in user space memory fetches during the execution of the syscall.

Double-Fetch Bug Mitigations

- The basic idea is to re-assure the control-dependence and data-dependence between the two fetches. In other words, the **automaticity** in user space memory fetches during the execution of the syscall.
- Based on our experience and our communications with kernel developers, we found **four patterns** in patching double-fetch bugs.

Double-Fetch Bug Mitigations

1. Override after second fetch.

```
1 kernel/events/core.c | 2 ++
2 1 file changed, 2 insertions(+)
3
4 diff --git a/kernel/events/core.c b/kernel/events/core.c
5 index ee20d4c..c0d7946 100644
6 --- a/kernel/events/core.c
7 +++ b/kernel/events/core.c
8 @@ -9611,6 +9611,8 @@ static int perf_copy_attr(struct perf_event_attr __user *uattr,
9     if (ret)
10         return -EFAULT;
11
12 + attr->size = size;
13 +
14     if (attr->__reserved_1)
15         return -EINVAL;
```

Override the overlapped memory (`attr->size`) with the value from the first fetch (`size`).

Double-Fetch Bug Mitigations

2. Abort on change detected.

```
1 net/compat.c | 7 ++++++
2 1 file changed, 7 insertions(+)
3
4 diff --git a/net/compat.c b/net/compat.c
5 index 6ded6c8..2238171 100644
6 --- a/net/compat.c
7 +++ b/net/compat.c
8 @@ -185,6 +185,13 @@ int cmsghdr_from_user_compat_to_kern(struct msghdr *kmsg, struct sock *sk,
9         ucmsg = cmsg_compat_nxthdr(kmsg, ucmsg, ucmlen);
10     }
11
12 + /*
13 +  * check the length of messages copied in is the same as the
14 +  * what we get from the first loop
15 +  */
16 + if ((char *)kmsg - (char *)kmsg_base != kcmLen)
17 +     goto Eival;
18 +
19 /* Ok, looks like we made it. Hook it up and return success. */
20 kmsg->msg_control = kmsg_base;
21 kmsg->msg_controllen = kcmLen;
```

Compare the new message length (kmsg - kmsg_base) with the value from the first fetch (kcmLen).

Double-Fetch Bug Mitigations

3. Refactor overlapped copies into incremental copies.

```
1 block/scsi_ioctl.c | 8 ++++++--
2 1 file changed, 7 insertions(+), 1 deletion(-)
3
4 diff --git a/block/scsi_ioctl.c b/block/scsi_ioctl.c
5 index 7440de4..8fe1e05 100644
6 --- a/block/scsi_ioctl.c
7 +++ b/block/scsi_ioctl.c
8 @@ -463,7 +463,13 @@ int sg_scsi_ioctl(struct request_queue *q, struct gendisk *disk, fmode_t mode,
9      */
10     err = -EFAULT;
11     req->cmd_len = cmdlen;
12 -    if (copy_from_user(req->cmd, sic->data, cmdlen))
13 +
14 +    /*
15 +     * avoid copying the opcode twice
16 +     */
17 +    memcpy(req->cmd, &opcode, sizeof(opcode));
18 +    if (copy_from_user(req->cmd + sizeof(opcode),
19 +                      sic->data + sizeof(opcode), cmdlen - sizeof(opcode)))
20         goto error;
21
22     if (in_len && copy_from_user(buffer, sic->data + cmdlen, in_len))
```

When copying the whole message, skip the information copied in the first fetch (+ sizeof(opcode)).

Double-Fetch Bug Mitigations

4. Refactor overlapped copies into a single-fetch.

```

1 drivers/isdn/i4l/isdn_ppp.c | 37 ++++++-----
2 1 file changed, 25 insertions(+), 12 deletions(-)
3
4 diff --git a/drivers/isdn/i4l/isdn_ppp.c b/drivers/isdn/i4l/isdn_ppp.c
5 index 6c44609..cd2b3c6 100644
6 --- a/drivers/isdn/i4l/isdn_ppp.c
7 +++ b/drivers/isdn/i4l/isdn_ppp.c
8 @@ -825,7 +825,6 @@ isdn_ppp_write(int min, struct file *file, const char __user *buf, int count)
9     isdn_net_local *lp;
10     struct ippp_struct *is;
11     int proto;
12 -    unsigned char protobuf[4];
13
14     is = file->private_data;
15
16 @@ -839,24 +838,28 @@ isdn_ppp_write(int min, struct file *file, const char __user *buf, int count)
17     if (!lp)
18         printk(KERN_DEBUG "isdn_ppp_write: lp == NULL\n");
19     else {
20 -        /*
21 -         * Don't reset huptimer for
22 -         * LCP packets. (Echo requests).
23 -         */
24 -        if (copy_from_user(protobuf, buf, 4))
25 -            return -EFAULT;
26 -        proto = PPP_PROTOCOL(protobuf);
27 -        if (proto != PPP_LCP)
28 -            lp->huptimer = 0;
29 +        if (lp->isdn_device < 0 || lp->isdn_channel < 0) {
30 +            unsigned char protobuf[4];
31 +            /*
32 +             * Don't reset huptimer for
33 +             * LCP packets. (Echo requests).
34 +             */
35 +            if (copy_from_user(protobuf, buf, 4))
36 +                return -EFAULT;
37 +
38 +            proto = PPP_PROTOCOL(protobuf);
39 +            if (proto != PPP_LCP)
40 +                lp->huptimer = 0;
41
42 -        if (lp->isdn_device < 0 || lp->isdn_channel < 0)
43             return 0;
44 +    }
45
46     if ((dev->drv[lp->isdn_device]->flags & DRV_FLAG_RUNNING) &&
47         lp->dialstate == 0 &&
48         (lp->flags & ISDN_NET_CONNECTED)) {
49         unsigned short hl;
50         struct sk_buff *skb;
51 +        unsigned char *cpy_buf;
52         /*
53          * we need to reserve enough space in front of
54          * sk_buff. old call to dev_alloc_skb only reserved
55 @@ -869,11 +872,21 @@ isdn_ppp_write(int min, struct file *file, const char __user *buf, int count)
56             return count;
57         }
58         skb_reserve(skb, hl);
59 -        if (copy_from_user(skb_put(skb, count), buf, count))
60 +        cpy_buf = skb_put(skb, count);
61 +        if (copy_from_user(cpy_buf, buf, count))
62         {
63             kfree_skb(skb);
64             return -EFAULT;
65         }
66 +
67 +        /*
68 +         * Don't reset huptimer for
69 +         * LCP packets. (Echo requests).
70 +         */
71 +        proto = PPP_PROTOCOL(cpy_buf);
72 +        if (proto != PPP_LCP)
73 +            lp->huptimer = 0;
74 +
75         if (is->debug & 0x40) {
76             printk(KERN_DEBUG "ppp xmit: len %d\n", (int) skb->len);
77             isdn_ppp_frame_log("xmit", skb->data, skb->len, 32, is->unit, lp->ppp_slot);

```

Such a strategy is usually very complex and requires careful refactoring.

Double-Fetch Bug Mitigations

Unfortunately, not all double-fetch bugs can be patched with these patterns. Some requires heavy refactoring of existing codebase or re-designing of structs, which requires substantial manual effort.

Double-Fetch Bug Mitigations

Unfortunately, not all double-fetch bugs can be patched with these patterns. Some requires heavy refactoring of existing codebase or re-designing of structs, which requires substantial manual effort.

Recently, DECAF has provided a promising solution in using TSX-based techniques to ensure user space memory access **automaticity** in syscall execution.

Limitations of Deadline

- **Source code coverage**
 - Files not compilable under LLVM.
 - Special combination of kernel configs (e.g., CONFIG_*).
- **Execution path construction**
 - Limit on total number of paths explored per fetch pair (4096).
 - Loop unrolling (limited to unroll once only).
- **Symbolic checking**
 - Ignores inline assemblies.
 - Imprecise pointer to memory object mapping.
 - Assumption on enclosing function.

Conclusion

- Detecting double-fetch bugs without a precise and formal definition has led to many false alerts and tremendous manual effort.
- Deadline is based on a precise modeling of double-fetch bugs and achieves both high accuracy and high scalability.
- Application beyond kernels: hypervisors, browsers, TEE, etc.
- Logic bugs are on the rise! We hope that more logic bugs can be modeled and checked systematically.

<https://github.com/sslabs-gatech/deadline>