

# Fuzzing the kernel

Andrew Turner

[andrew@FreeBSD.org](mailto:andrew@FreeBSD.org)

# Who am I?

- FreeBSD Committer (andrew@)
- Research Associate in the University of Cambridge
- Sometimes a Freelance Software Engineer

Sanitizers

# Sanitizers

- A tool from the compiler to instrument code
- Add function calls on in interesting points in the code, e.g.:
  - Entry to a basic block
  - On comparison operations
  - Every memory access
- The compiler provides a runtime for userspace
- We need our own runtime in the kernel

Undefined Behaviour

# KUBSAN – Undefined Behaviour Sanitizer

- Instruments code that may be undefined depending on input, e.g.
  - Misaligned or NULL pointer
  - Shift out of bounds
- From the NetBSD  $\mu$ UBsan
  - imported 3 August 2018
- Imported into FreeBSD 6 November 2018
  - Off by default (large increase in kernel file size)
- Imported into OpenBSD 18 March 2019

# KUBSAN Reports

- Misaligned memory access:
  - UBSan: Undefined Behavior in `.../sys/vm/uma_core.c:1746:8`, member access within misaligned address `0xffff8087ffde7c0` for type `'struct uma_zone'` which requires 128 byte alignment
- NULL pointer dereference:
  - UBSan: Undefined Behavior in `.../sys/contrib/ck/src/ck_epoch.c:143:1`, member access within null pointer of type `'struct ck_epoch_record'`
- Shift out of bounds:
  - UBSan: Undefined Behavior in `.../sys/cddl/contrib/opensolaris/uts/common/fs/zfs/vdev_label.c:410:14`, shift exponent 64 is too large for 64-bit type `'unsigned long long'`

Coverage

# KCOV – Coverage Sanitizer

- Coverage sanitizer
- Inserts function calls to trace:
  - The start of basic blocks
  - On comparison operations
- Comparison tracing includes values being compared
  - Useful for finding what input data to try changing
- Committed to:
  - OpenBSD on 19 August 2018
  - FreeBSD on 12 January 2019
  - NetBSD on 23 February 2019

# KCOV – PC Tracing

Count	PC	PC	PC	...
348	0xffffffff81595cf0	0xffffffff8155c2f0	0xffffffff8155ccb0	...

- Starts with the number of entries
- Each entry contains an address in the basic block
  - Probably the return address of the inserted function
- Each field is:
  - `uintptr_t` on OpenBSD
  - `uint64_t` on FreeBSD and NetBSD (as `kcov_int_t`)

# KCOV – Comparison Tracing

Count	Type	Arg	Arg	PC	Type	Arg	Arg	PC	...
348	0x2	0x10	0x20	0xffffffff...	0x7	0x8080	0x8080	0xffffffff...	...

- Starts with the number of entries
- Each entry contains:
  - A comparison type – encodes width and if comparing with a const
  - Two arguments
  - An address near the comparison
    - Probably the return address of the inserted function
- Each field is:
  - `uintptr_t` on OpenBSD
  - `uint64_t` on FreeBSD and NetBSD (as `kcov_int_t`)

# KCOV – User interface

1. User opens `/dev/kcov`
2. Sets the buffer size with an `ioctl`
3. `mmaps` the buffer
4. Enables tracing within the thread being traces
  - May not be the same thread (or process) as opened the device
5. Zeros the first entry in the buffer
6. Runs the traced operations
7. Disables tracing
8. May repeat from 4
9. `Unmaps` the buffer
10. Closes the device

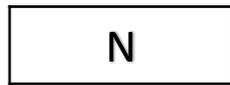
# Address Space

# KASAN – Address Space Sanitizer

- Checks memory accesses are in bounds
- Uses a shadow map to mark valid memory
  - 1 shadow byte for each 8 bytes
- Can mark the first 1-8 bytes as valid
  - Must be contiguous
- Can mark all bytes as invalid
  - Includes data on why, e.g. stack padding

# KASAN – Address Space Sanitizer

1 byte in the shadow map (signed)



$N > 0$ , e.g.  $N = 3$



$N = 0$



$N < 0$



Valid memory:



Invalid memory:



# KASAN – Address Space Sanitizer

- All allocations are now aligned to at least an 8 byte boundary
- Allocations are rounded up to an 8 byte boundary
  - Memory past the end of the allocation is marked as unusable
- One or more 8 byte blocks of unusable memory after the allocation

# KASAN – Example

```
void get_data(int *output, size_t count);
```

```
int example(void) {  
    int ret, *data = malloc(sizeof(int), M_TEMP, M_WAITOK);  
  
    get_data(data, 1);  
    ret = *data;  
  
    free(data, M_TEMP);  
    return (ret);  
}
```

# KASAN – Example

Allocation ↓

8	9	10	11	12	13	14	15
0	1	2	3	4	5	6	7

- Allocated 4 bytes
- Aligned data to an 8 byte boundary
- Padded to 8 bytes
- Allocated 8 bytes of invalid data

# KASAN – Example

Allocation ↓

Shadow map →

8	9	10	11	12	13	14	15
0	1	2	3	4	5	6	7

...
0xFF
0x04
...

- Allocated up to 8 bytes
- Padded up to 8 bytes
- 8 bytes of padding after the allocation

# KASAN – Example

Allocation ↓

Shadow map →

8	9	10	11	12	13	14	15
0	1	2	3	4	5	6	7

...
0xFF
0x04
...

- A load or store that includes bytes 4-15 will be detected
  - Can warn with printf or panic
- A load or store past byte 15 may or may not be detected
  - It depends on if it has been allocated

# KASAN – Example

```
void get_data(int *output, size_t count);
```

```
int example(void) {  
    int ret, *data = malloc(sizeof(int), M_TEMP, M_WAITOK);  
  
    get_data(data, 1);  
    ret = data[1];          /* Out of bounds: access past the end of data */  
  
    free(data, M_TEMP);  
    return (ret);  
}
```

# KASAN

- Committed to NetBSD 20 August 2018
- FreeBSD has a Google Summer of Code student working on it

# KHWASAN – Hardware Assisted ASAN

- An arm64 specific extension
- Enable Top Byte Ignore in the kernel
  - The top 8 bits of a pointer are ignored
- Store an 8 bit tag in the top byte
  - One tag is reserved for free memory
- Store the same tag in the shadow map
- Allocate random tags when allocating memory

# KHWASAN – Example

Initial state  
No allocations

0
1
2
3
4
5
6
7
8
9
A
B

First allocation

0
1
2
3
4
5
6
7
8
9
A
B

More allocations

0
1
2
3
4
5
6
7
8
9
A
B

Load from location 3  
with a blue pointer



# KHWASAN – Hardware Assisted ASAN

- The shadow map is 1 byte per 16 allocated bytes
- Allocations are 16 byte aligned and sized
- Not able to detect slightly out of bounds access
  - Within a 16 byte block
- Can probabilistically detect larger out of bounds access of allocated memory
  - Just under 1/256 probability of an incorrect tag match

# CHERI

- Add bounds and permissions to pointers to create capabilities
- Capabilities are non-forgable
- Can only reduce bounds and permissions
- Capabilities can only be derived from other capabilities
  - At least one will be provided to the kernel on boot
- With CheriABI all pointers are capabilities
  - See Brooks' talk tomorrow for more information

# CHERI

- CHERI with CheriABI can detect all out-of-bounds access
  - Not just slightly out of bounds like KASAN
  - No tag collision like KHWASAN
- Research on narrowing bounds more
  - Stopping buffer overflows for variables within a struct

Memory

# KMSAN – Memory Sanitizer

- Checks use of uninitialized values in the kernel
- Use is defined as:
  - Conditionals
  - Pointer dereference
  - Copied to userspace

# KMSAN – Memory Sanitizer

```
int a;
```

```
int b = a; /* Not a use */
```

```
copyout(&b, u, sizeof(b)); /* An uninit use */
```

# KMSAN – Memory Sanitizer

```
int a;
```

```
if (flag)
```

```
    a = 1;
```

```
c = a + b;    /* Not a use of a */
```

```
if (flag)
```

```
    copyout(&c, u, sizeof(c)); /* A use, don't report */
```

# KMSAN – Memory Sanitizer

```
struct config *example(void)
{
    struct config *conf;

    conf = malloc(sizeof(*conf), M_TEMP, M_WAITOK);
    init_config(conf);
}

void init_config(struct config *conf)
{
    if (conf->flag != 0)      /* An uninit use */
        do_something();
}
```

# KMSAN – Memory Sanitizer

- As with KASAN it uses a shadow map
- 1 bit per byte
  - Set when poisoned (uninitialized)
- Memory is poisoned by default
- `malloc(9)` will unpoison memory with `M_ZERO`
- Writing a constant unpoisons memory
- Shadow state is propagated

# KLEAK – NetBSD memory leak detector

- Similar in concept to KMSAN
  - Find copying uninitialized memory to userspace
- Uses in-band signalling
  - Use a magic value, then check for it when copying to userspace
- Uses the coverage sanitizer instrumentalization to poison the stack
- Prone to false positives
  - In-band value may be legitimately in the data
  - Cleaver choice of value to reduce this chance

# Threading

# KTSAN – Thread Sanitizer

- Find data races
- Still a work in progress by Google in Linux
  - May be on hold

Why add Kernel Sanitizers?

# Why add Kernel Sanitizers?

- Find and fix more bugs
- KCOV allows kernel fuzzing
- Improves fuzzing by making bugs easier to find

# Syzkaller

- A system call fuzzer from Google
- Supports many kernels including FreeBSD, NetBSD, and OpenBSD
- Finds new ways to panic the kernel from userspace
- Google hosts an instance on their infrastructure

[fixed bugs \(16\)](#)**Instances:**

Name	Active	Uptime	Corpus	Coverage	Crashes	Execs	Kernel build			syzkaller build		
							Commit	Freshness	Status	Commit	Freshness	Status
ci-freebsd-main	now	5h46m	8936	<a href="#">81982</a>	137	1563723	<a href="#">85eaade9</a>	13h46m		<a href="#">2755003a</a>	9d02h	

**upstream (47):**

<u>Title</u>	<u>Repro</u>	<u>Bisected</u>	<u>Count</u>	<u>Last</u>	<u>Reported</u>
<a href="#">panic: inp_leave_group: imf_sources not empty (2)</a>	<code>syz</code>		2	2h31m	<a href="#">3h01m</a>
<a href="#">Fatal trap 12: page fault in uma_dbg_free</a>			1	1d05h	<a href="#">1d05h</a>
<a href="#">panic: Duplicate free of ADDR from zone ADDR(16) slab ADDR(241)</a>			1	1d16h	<a href="#">1d16h</a>
<a href="#">Fatal trap 12: page fault in _mtx_assert (2)</a>			1	1d18h	<a href="#">1d18h</a>
<a href="#">Fatal trap 9: general protection fault in sys_nlm_syscall</a>			2	2d03h	<a href="#">4d22h</a>
<a href="#">Fatal trap 12: page fault in fifo_close</a>			17	17h43m	<a href="#">5d15h</a>
<a href="#">panic: Assertion lock == sq-&gt;sq_lock failed at /syzkaller/managers/m...</a>			1	6d03h	<a href="#">6d03h</a>
<a href="#">panic: Memory modified after free ADDR(256) val=0 @ ADDR</a>	<code>syz</code>		4	1d05h	<a href="#">6d11h</a>
<a href="#">Fatal trap 12: page fault in uma_dbg_alloc</a>			1	6d12h	<a href="#">6d12h</a>
<a href="#">panic: mtx_lock() of destroyed mutex at sys/kern/sys_socket.c:LINE</a>	<code>syz</code>		1	6d19h	<a href="#">6d19h</a>
<a href="#">panic: mtx_lock() of spin mutex (null) @ /syzkaller/managers/main/k...</a>			1	6d19h	<a href="#">6d19h</a>
<a href="#">Fatal trap 12: page fault in link_elf_search_symbol</a>			1	7d00h	<a href="#">7d00h</a>
<a href="#">Fatal trap 12: page fault in ip_output</a>			2	6d14h	<a href="#">7d13h</a>
<a href="#">panic: Most recently used by temp</a>			4	3d12h	<a href="#">7d23h</a>
<a href="#">panic: Bad tailq NEXT(ADDR-&gt;tqh_last) != NULL</a>			1	8d10h	<a href="#">8d10h</a>
<a href="#">panic: Most recently used by tty</a>	<code>syz</code>		11	4d00h	<a href="#">9d00h</a>
<a href="#">panic: Most recently used by ip6opt</a>			1	11d	<a href="#">11d</a>

[fixed bugs \(1240\)](#)**Instances:**

Name	Active	Uptime	Corpus	Coverage	Crashes	Execs	Kernel build			syzkaller build		
							Commit	Freshness	Status	Commit	Freshness	Status
ci-upstream-bpf-kasan-gce	now	1h53m	10023	<a href="#">284396</a>	163	2141078	<a href="#">d72386fe</a>	1d08h		<a href="#">95dfd515</a>	2h21m	
ci-upstream-bpf-next-kasan-gce	now	1h53m	10406	<a href="#">287727</a>	249	2742186	<a href="#">35c99ffa</a>	2d15h		<a href="#">95dfd515</a>	2h21m	
ci-upstream-gce-leak	now	24m	11460	<a href="#">330887</a>	11	87189	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kasan-gce	now	39m	72945	<a href="#">4602079</a>	466	912856	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kasan-gce-386	now	57m	22792	<a href="#">430747</a>	14	1914041	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kasan-gce-root	now	14m	74600	<a href="#">4827346</a>	440	892084	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kasan-gce-selinux-root	now	30m	73966	<a href="#">4668186</a>	282	1229467	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kasan-gce-smack-root	now	48m	72906	<a href="#">5032570</a>	236	1588729	<a href="#">a6a4b66b</a>	10h06m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-kmsan-gce	now	1h07m	56237	<a href="#">3125847</a>	614	465976	<a href="#">2b51a114</a>	4h01m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-linux-next-kasan-gce-root	now	1h53m	72275	<a href="#">5189609</a>	156	1956910	<a href="#">b1d6682e</a>	8h14m		<a href="#">95dfd515</a>	2h21m	
ci-upstream-net-kasan-gce	now	1h53m	19059	<a href="#">429673</a>	49	4444259	<a href="#">35c99ffa</a>	2d15h		<a href="#">95dfd515</a>	2h21m	
ci-upstream-net-this-kasan-gce	now	1h53m	17762	<a href="#">409496</a>	36	3111905	<a href="#">510e2ced</a>	15h46m		<a href="#">95dfd515</a>	2h21m	
ci2-upstream-usb	now	5h28m	770	<a href="#">32489</a>	466	747197	<a href="#">43151d6c</a>	580d		<a href="#">2755003a</a>	9d02h	

**open (505):**

Title	Repro	Bisected	Count	Last	Reported
<a href="#">general protection fault in ext4 mb initialize context</a>			3	18d	<a href="#">1d01h</a>
<a href="#">KASAN: use-after-free Write in xfrm hash rebuild</a>			1	7d05h	<a href="#">1d01h</a>
<a href="#">KASAN: use-after-free Write in xfrm policy unlink (2)</a>			2	4d09h	<a href="#">1d01h</a>
<a href="#">WARNING: locking bug in inet autobind</a>			1	1d11h	<a href="#">1d06h</a>
<a href="#">WARNING: locking bug in udpv6 pre connect</a>			1	4d21h	<a href="#">1d06h</a>

# Syzkaller

- Will combine system calls to try finding new paths through the kernel
- Understands arguments
  - E.g. read takes a file description, a pointer, and a length
- Will try to mix syscalls in interesting ways
  - Pass a socket into something that doesn't take a socket
- Very good at panicking the kernel
- Will try to find a reproducer
- Adding a sanitizer makes it easier to find memory issues



**panic: ffs\_blkfree\_cg: freeing free block**

Status: [fixed on 2019/04/29 23:55](#)

Reported-by: [syzbot+36fd786cb3ab88f18c9b@syzkaller.appspotmail.com](mailto:syzbot+36fd786cb3ab88f18c9b@syzkaller.appspotmail.com)

Fix commit: [a7a455c2 Optimize lseek\(SEEK\\_DATA\) on UFS.](#)

First crash: 60d, last: 60d

**similar bugs (1):**

<a href="#">Kernel</a>	<a href="#">Title</a>	<a href="#">Repro</a>	<a href="#">Bisected</a>	<a href="#">Count</a>	<a href="#">Last</a>	<a href="#">Reported</a>	<a href="#">Patched</a>	<a href="#">Status</a>
freebsd	<a href="#">panic: ffs_blkfree_cg: freeing free block (2)</a>			1	17d	<a href="#">17d</a>	0/1	<a href="#">upstream: reported on 2019/04/30 12:16</a>

**Sample crash report:**

```
panic: ffs_blkfree_cg: freeing free block
cpuid = 0
time = 1552872502
KDB: stack backtrace:
db_trace_self_wrapper() at db_trace_self_wrapper+0x47/frame 0xfffffe0020dfc150
vpanic() at vpanic+0x1e0/frame 0xfffffe0020dfc1b0
panic() at panic+0x43/frame 0xfffffe0020dfc210
ffs_blkfree_cg() at ffs_blkfree_cg+0x6e9/frame 0xfffffe0020dfc2d0
ffs_blkfree() at ffs_blkfree+0x15e/frame 0xfffffe0020dfc350
ffs_indirtrunc() at ffs_indirtrunc+0x724/frame 0xfffffe0020dfc450
ffs_indirtrunc() at ffs_indirtrunc+0x856/frame 0xfffffe0020dfc530
ffs_truncate() at ffs_truncate+0x17c3/frame 0xfffffe0020dfc720
ufs_setattr() at ufs_setattr+0x918/frame 0xfffffe0020dfc7c0
VOP_SETATTR_APV() at VOP_SETATTR_APV+0xc2/frame 0xfffffe0020dfc7f0
vn_truncate() at vn_truncate+0x23f/frame 0xfffffe0020dfc930
kern_ftruncate() at kern_ftruncate+0x13b/frame 0xfffffe0020dfc980
amd64_syscall() at amd64_syscall+0x436/frame 0xfffffe0020dfcab0
fast_syscall_common() at fast_syscall_common+0x101/frame 0xfffffe0020dfcab0
--- syscall (0, FreeBSD ELF64, nosys), rip = 0x42132a, rsp = 0x7fffffff88, rbp = 0x2 ---
KDB: enter: panic
[ thread pid 762 tid 100093 ]
Stopped at      kdb_enter+0x6a: movq    $0,kdb_why
```

**All crashes (5):**

<a href="#">Manager</a>	<a href="#">Time</a>	<a href="#">Kernel</a>	<a href="#">Commit</a>	<a href="#">Syzkaller</a>	<a href="#">Config</a>	<a href="#">Log</a>	<a href="#">Report</a>	<a href="#">Syz repro</a>	<a href="#">C repro</a>
ci-freebsd-main	2019/03/18 01:31	freebsd	<a href="#">8b17fbc2</a>	<a href="#">f8757044</a>		<a href="#">log</a>	<a href="#">report</a>	<a href="#">syz</a>	<a href="#">C</a>
ci-freebsd-main	2019/03/18 01:14	freebsd	<a href="#">8b17fbc2</a>	<a href="#">f8757044</a>		<a href="#">log</a>	<a href="#">report</a>		
ci-freebsd-main	2019/03/18 05:00	freebsd	<a href="#">8b17fbc2</a>	<a href="#">f8757044</a>		<a href="#">log</a>	<a href="#">report</a>		

# Syzkaller

- Emails a per-project mailing list with new issues
- Fixes should be tagged in the commit
- Will check the issue is fixed
  
- Join the appropriate list if you care about kernel quality

# AFL – American Fuzzy Lop

- A file format fuzzer
  - Can change a file and see if any new paths are found
- Test patches for KCOV to support AFL
- Before starting clear the buffer
- On each basic block:
  - Calculate  $(\text{hash}(\text{old\_ptr}) \wedge \text{hash}(\text{new\_ptr})) \% \text{buffer\_length}$
  - Increment this entry
- Patched AFL to talk to kcov

# AFL – Fuzzing UFS

- Tried fuzzing a 128K UFS image
- Just mount and unmount the image
- Very slow
  - ~60 mounts/second
  - Around 12 days to try all single bitflips

Conclusion

# Conclusion – Sanitizers

- FreeBSD, NetBSD, and OpenBSD have KCOV and KUBSAN
- NetBSD has KASAN, with it planned for FreeBSD
- Other sanitizers need work
- Will make bugs easier to find

# Conclusion – Fuzzing

- Google runs a syscall fuzzer on FreeBSD, NetBSD, and OpenBSD
- Look through the reports & fix the code
- AFL may be useful in the future, but currently is too slow

Questions?