

# syzkaller

the next gen kernel fuzzer

Qualcomm Mobile Security Summit 2017

Dmitry Vyukov (dvyukov@), Google

# Agenda

- Kernel sanitizers (KASAN, KMSAN, KTSAN)
- Why new fuzzer? How is it better?
- Operational side
- Tutorial
- Extending syzkaller to fuzz new drivers

# KASAN (KernelAddressSanitizer)

Fast and comprehensive solution for both **UAF** and **OOB**.

- based on compiler instrumentation
- detects both OOB writes and OOB reads
- detects OOB on heap, stack and globals
- strong UAF detection
- double-free detection
- detects bugs at the point of occurrence
- prints informative reports

Upstream: CONFIG\_KASAN + gcc 5.0+

# KASAN Report (CVE-2013-7446)

BUG: KASan: **use-after-free** in remove\_wait\_queue  
Write of size 8 by task syzkaller\_execu/10568

## Call Trace:

```
list_del include/linux/list.h:107
__remove_wait_queue include/linux/wait.h:145
remove_wait_queue+0xfb/0x120 kernel/sched/wait.c:50
...
SYSC_exit_group kernel/exit.c:885
```

## Allocated:

```
kmem_cache_alloc+0x10d/0x140 mm/slub.c:2517
sk_prot_alloc+0x69/0x340 net/core/sock.c:1329
sk_alloc+0x33/0x280 net/core/sock.c:1404
...
SYSC_socketpair net/socket.c:1281
```

## Freed:

```
kmem_cache_free+0x161/0x180 mm/slub.c:2745
sk_prot_free net/core/sock.c:1374
sk_destruct+0x2e9/0x400 net/core/sock.c:1452
...
SYSC_write fs/read_write.c:585
```

# KMSAN (KernelMemorySanitizer)

Detects uses of **uninitialized memory**.

In the context of security: **information leaks** (local or remote).

Most likely it does not what you think it does:

not loads of not-stored-to variables, but uses of uninitialized values (see backup)

No false positives and almost no false negatives.

Not upstreamed yet (on github): `CONFIG_KMSAN` + clang

# KMSAN Report

BUG: KMSAN: **use of uninitialized memory** in nf\_ct\_frag6\_gather

## Call Trace:

```
nf_ct_frag6_gather+0xf5a/0x44a0 net/ipv6/nf_conntrack_reasm.c:577
ipv6_defrag+0x1d9/0x280 net/ipv6/netfilter/nf_defrag_ipv6_hooks.c:68
nf_hook_entry_hookfn ./include/linux/netfilter.h:102
nf_hook_slow+0x13f/0x3c0 net/netfilter/core.c:310
...
Sys_sendto+0xbc/0xe0 net/socket.c:1664
```

## Origin:

```
__alloc_skb+0x2cd/0x740 net/core/skbuff.c:231
alloc_skb ./include/linux/skbuff.h:933
alloc_skb_with_frags+0x209/0xbc0 net/core/skbuff.c:4678
sock_alloc_send_pskb+0x9ff/0xe00 net/core/sock.c:1903
...
Sys_sendto+0xbc/0xe0 net/socket.c:1664
```

# KTSAN (KernelThreadSanitizer)

Detects **data races** (two unsynchronized accesses in different threads, at least one is a write).

Kernel data races represent **security threat**:

- TOCTOU
- uninit/wrong credentials
- racy use-after-frees
- double frees caused by races
- the most frequent type of bug

Not upstreamed yet (prototype on github): `CONFIG_KTSAN`

# KTSAN Report (CVE-2015-7613)

ThreadSanitizer: **data-race** in ipc\_obtain\_object\_check

**Read** at 0x123 of size 8 by thread 234 on CPU 5:

```
ipc_obtain_object_check+0x7d/0xd0 ipc/util.c:621
msg_obtain_object_check ipc/msg.c:90
msgctl_nolock.constprop.9+0x208/0x430 ipc/msg.c:480
SYSC_msgctl ipc/msg.c:538
```

**Previous write** at 0x123 of size 8 by thread 567 on CPU 4:

```
ipc_addid+0x217/0x260 ipc/util.c:257
newque+0xac/0x240 ipc/msg.c:141
ipcget_public ipc/util.c:355
ipcget+0x202/0x280 ipc/util.c:646
SYSC_msgget ipc/msg.c:255
```

**Also:** locked mutexes, thread creation stacks, allocation stack, etc.



fuzzing + sanitizers = perfect fit!

**syzkaller**

# Existing Fuzzers

trinity/iknowthis in essence:

```
for (;;) { syscall(rand(), rand(), rand()); }
```

Do know argument types, so more like:

```
for (;;) { syscall(rand(), rand_fd(), rand_addr()); }
```

Downsides:

- tend to find shallow bugs
- frequently no reproducers
- poorly suitable for regression testing

# Coverage-guided Fuzzing

Code coverage guiding:

- corpus of "interesting" inputs
- mutate and execute inputs from corpus
- if inputs gives new code coverage, add it to corpus

Advantages:

- turns exponential problem into linear (more or less)
- inputs are reproducers
- corpus is perfect for regression testing

But how to apply it to kernel?

# Kernel Code Coverage

- CONFIG\_KCOV=y (upstream)
- based on compiler instrumentation (gcc6+)
- compiler inserts runtime callbacks into every basic block
- coverage per-thread per-syscall

# Syscall Descriptions

Declarative description of all syscalls:

```
open(file filename, flags flags[open_flags],  
      mode flags[open_mode]) fd
```

```
read(fd fd, buf buffer[out], size len[buf])
```

```
close(fd fd)
```

```
open_flags = O_RDONLY, O_WRONLY, O_RDWR, O_APPEND ...
```

# Programs

These descriptions allows to **generate** and **mutate** "programs" in the following form:

```
mmap (& (0x7f0000000000), (0x1000), 0x3, 0x32, -1, 0)  
r0 = open (& (0x7f0000000000) = "./file0", 0x3, 0x9)  
read (r0, & (0x7f0000000000), 42)  
close (r0)
```

# Algorithm

```
Start with empty corpus of programs;
```

```
while (true) {
```

```
    Generate a new program, or
```

```
        choose an existing program from corpus and mutate it;
```

```
    Interpret the program, collect coverage from every syscall;
```

```
    if (program covers code that wasn't covered previously)
```

```
        minimize the program and add it to corpus;
```

```
}
```



Operational Side

Ideally...



# Reality...

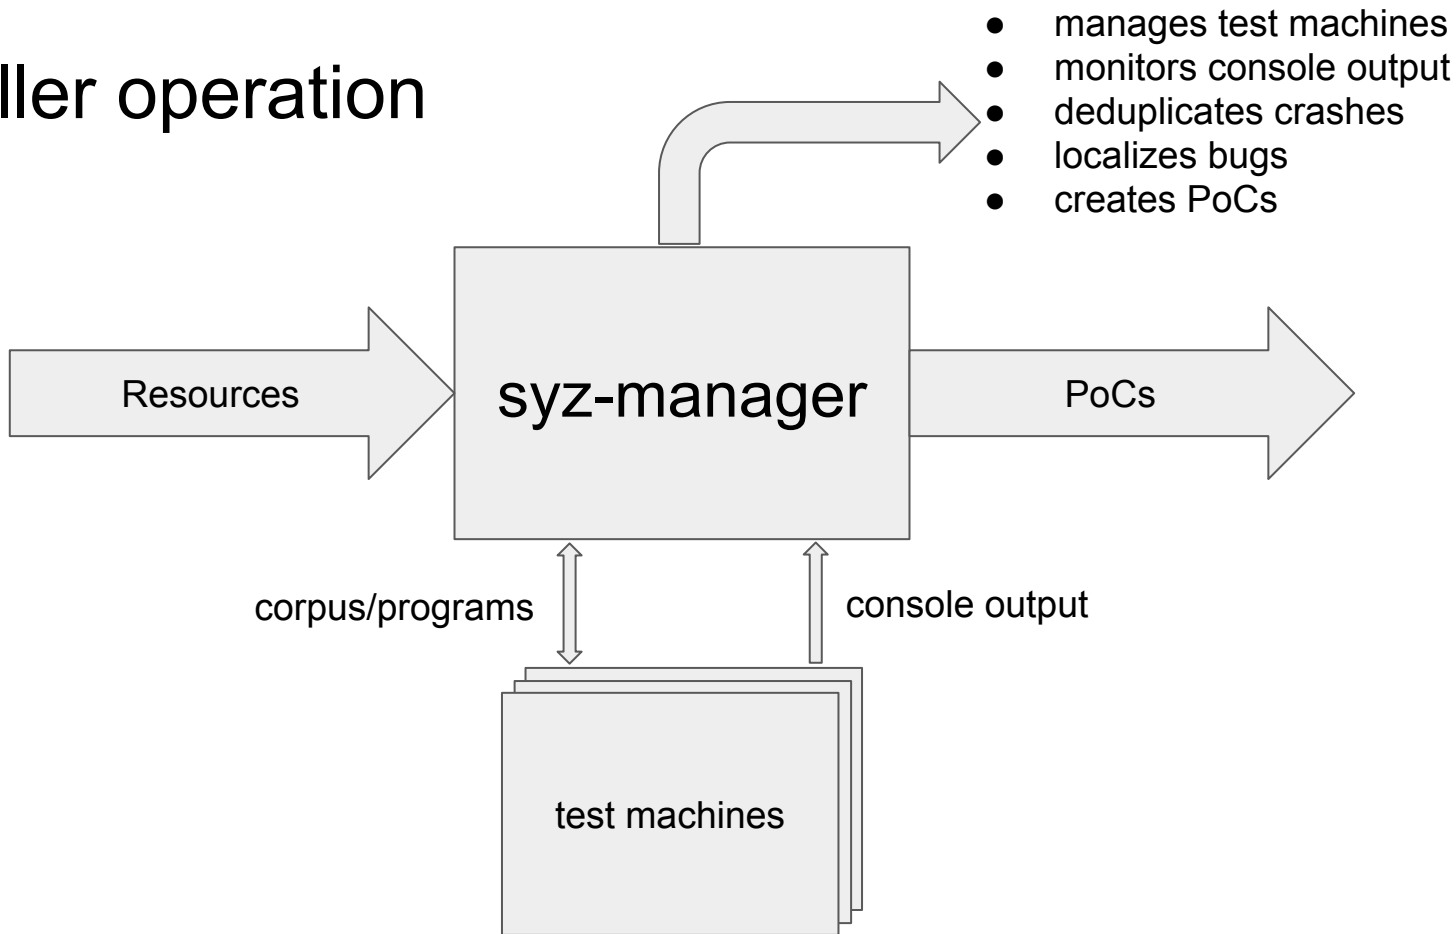
Operation of a typical kernel fuzzer:

- **manually** create a bunch of VMs
- **manually** copy and start the binary
- **manually** monitor console output
- **manually** deduplicate crashes
- **manually** localize and reproduce
- **manually** restart the crashed VM

Works only if you want to find a handful of crashes. Otherwise -- full time job.

Worse for physical devices: you also need to stand there and press power button.

# syzkaller operation



# Test Machines

Currently supports:

- QEMU
- GCE (Google Compute Engine)
- Android devices (with serial cable/Suzy-Q)
- ODROID boards

The list is extensible, to support a new type need to:

- get console output (to grep for crashes)
- reboot/recreate (to repair after a crash)
- copy/run binary (think of scp/ssh)

# Tutorial

# Setting up syzkaller

- build kernel in a slightly special way
- build syzkaller
- write config file
- run syzkaller

# Kernel Build

- Fresh compiler
  - gcc 7.0+ is perfect
  - gcc 6.0+ will work
  - gcc 5.0+ will work with sancov plugin
- CONFIG\_KCOV for code coverage
- CONFIG\_DEBUG\_INFO for report symbolization
- CONFIG\_KASAN for KASAN
- As many debug configs as possible (to get more bugs)
  - CONFIG\_LOCKDEP
  - CONFIG\_DEBUG\_VM
  - CONFIG\_DEBUG\_ATOMIC\_SLEEP



# syzkaller Build

syzkaller is written in Go

- download Go 1.8.1 from [golang.org/dl](http://golang.org/dl)
- unpack it into ~/go1.8
- export PATH=~/go1.8/bin:\$PATH
- go get -d github.com/google/syzkaller/...
- cd go/src/github.com/google/syzkaller
- make

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinuz": "/linux/vmlinuz",  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

**# local address for web UI**  
**# shows crashes, coverage, stats, etc**

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinux": "/linux/vmlinux",  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

**# local dir to save corpus and crashes**  
**# syzkaller checkout dir**

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage", # compressed kernel  
  "vmlinux": "/linux/vmlinux", # kernel object file (for symbolization)  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

# syzkaller config for QEMU

```
{  
"http": "localhost:12345",  
"workdir": "./workdir",  
"syzkaller": ".",  
"kernel": "/linux/arch/x86/boot/bzImage",  
"vmlinuz": "/linux/vmlinuz",  
"image": "/image/wheezy.img",  
"sshkey": "/image/id_rsa",  
"type": "qemu",  
"count": 4,  
"procs": 8,  
"cpu": 2,  
"mem": 2048,  
"sandbox": "none"  
}
```

**# disk image and root ssh key  
# can be created with:  
# ./tools/create-image.sh**

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinuz": "/linux/vmlinuz",  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

**# test machine type**  
**# one of "qemu", "gce", "adb", "odroid"**  
**# other types need slightly different config**

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinuz": "/linux/vmlinuz",  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

**# number of VMs to use**

**# number of parallel tests per VM**

**# CPUs per VM**

**# RAM per VM**

# syzkaller config for QEMU

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
  "kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinuz": "/linux/vmlinuz",  
  "image": "/image/wheezy.img",  
  "sshkey": "/image/id_rsa",  
  "type": "qemu",  
  "count": 4,  
  "procs": 8,  
  "cpu": 2,  
  "mem": 2048,  
  "sandbox": "none"  
}
```

```
# test sandboxing mode  
# "none": run under root as is  
# "setuid": do setuid(nobody)  
# "namespace": run in namespaces
```



# syzkaller config for Android

```
{  
  "http": "localhost:12345",  
  "workdir": "./workdir",  
  "syzkaller": ".",  
"kernel": "/linux/arch/x86/boot/bzImage",  
  "vmlinux": "/linux/vmlinux",  
"image": "/image/wheezy.img",  
"sshkey": "/image/id_rsa",  
  "type": "qemu",  
"count": 4,  
  "procs": 8,  
"cpu": 2,  
"mem": 2048,  
  "sandbox": "none",  
  "type": "adb",  
  "devices": ["ADBDEVID1", "ADBDEVID2", "ADBDEVID3"]  
}
```

Ready to go

```
$ bin/syz-manager -config my.cfg
```

Extending syzkaller

# Steps to extend syzkaller

Let's say we want to fuzz `/dev/ion` device on Android.

1. Create new file for descriptions: `sys/ion.txt`
2. Write descriptions for the new system calls
3. Rebuild
4. Run

## sys/ion.txt: includes

```
# includes are required to extract  
# values of literal constants used in the descriptions.
```

```
# for ioctl-related constants
```

```
include <asm/ioctl.h>
```

```
# for ion-related constants
```

```
include <drivers/staging/android/uapi/ion.h>
```

## sys/ion.txt: resources

```
# resources allow syzkaller to pass output of one syscall  
# as an input to another syscall.
```

```
# these 2 are "derived" from fd
```

```
resource fd_ion[fd]
```

```
resource fd_ion_generic[fd]
```

```
# this is "derived" from int32
```

```
resource ion_handle[int32]
```

## sys/ion.txt: system calls

```
open$ion(file ptr[in, string["/dev/ion"]],  
         flags flags[open_flags], mode const[0]) fd_ion  
  
ioctl$IION_IOC_ALLOC(fd fd_ion, cmd const[IION_IOC_ALLOC],  
                    arg ptr[inout, ion_allocation_data])  
ioctl$IION_IOC_FREE(fd fd_ion, cmd const[IION_IOC_FREE],  
                   arg ptr[in, ion_handle_data])  
ioctl$IION_IOC_MAP(fd fd_ion, cmd const[IION_IOC_MAP],  
                  arg ptr[inout, ion_fd_data])  
... few more of them ...
```

# sys/ion.txt: structs

```
ion_allocation_data {  
    len      intptr  
    align    intptr  
    heapid   int32  
    flags    int32  
    handle   ion_handle  
}
```

```
ion_fd_data {  
    handle ion_handle  
    fd     fd_ion_generic  
}
```



# Rebuilding

```
--- a/extract.sh
+++ b/extract.sh
-ANDROID_FILES="sys/tlk_device.txt"
+ANDROID_FILES="sys/tlk_device.txt sys/ion.txt"
```

```
# updates .const files
```

```
$ make extract ANDROID=/path/to/android/kernel/source
```

```
# rebuilds binaries
```

```
$ make
```

# sys/ion\_arm64.const

```
# AUTOGENERATED FILE
ION_IOC_ALLOC = 3223341312
ION_IOC_CUSTOM = 3222292742
ION_IOC_FREE = 3221506305
ION_IOC_IMPORT = 3221768453
ION_IOC_MAP = 3221768450
ION_IOC_SHARE = 3221768452
ION_IOC_SYNC = 3221768455
__NR_ioctl = 29
```

Ready to go

```
$ bin/syz-manager -config my.cfg
```

tty: fix port buffer locking  
kvm: warning in kvm\_load\_guest\_fpu  
drivers/scsi: GPF in sg\_read  
net/ipv4: use-after-free in ip\_mc\_drop\_socket  
net/ipv6: GPF in rt6\_device\_match  
x86: warning: kernel stack regs has bad 'bp' value  
net/key: slab-out-of-bounds in pfkey\_compile\_policy  
net/ipv6: warning in inet6\_ifa\_finish\_destroy  
net/ipv6: use-after-free in  
\_\_call\_rcu/in6\_dev\_finish\_destroy\_rcu  
net/ipv6: slab-out-of-bounds in ip6\_tnl\_xmit  
net/rose: null-ptr-deref in rose\_route\_frame  
time: hang due to timer\_create/timer\_settime  
net/core: BUG in unregister\_netdevice\_many  
net/xfrm: stack-out-of-bounds in xfrm\_state\_find  
net/bonding: stack-out-of-bounds in bond\_enslave  
net: ipv6: RTF\_PCPU should not be setttable from userspace  
fs/notify/inotify: slab-out-of-bounds write in strcpy  
net/ipv6: slab-out-of-bounds read in seg6\_validate\_srh  
kernel BUG at mm/hugetlb.c:742!  
net/key: slab-out-of-bounds in parse\_ipsecrequests  
net/ipv4: use-after-free in ipv4\_datagram\_support\_cmrg  
net/ipv4: use-after-free in ip\_queue\_xmit  
net: use-after-free in \_\_ns\_get\_path  
net/ipv4: use-after-free in ip\_check\_mc\_rcu  
net/ipv6: use-after-free in ipv6\_sock\_ac\_close  
net/ipv4: use-after-free in ipv4\_mtu  
net/dccp: BUG in tfrc\_rx\_hist\_sample\_rtt  
net/sctp: list double add warning in sctp\_endpoint\_add\_asoc  
kvm: use-after-free in srcu\_reschedule  
ata: WARNING in ata\_bmdma\_qc\_issue  
net/sched: GPF in qdisc\_hash\_add  
sg: random memory corruptions  
fs: GPF in deactivate\_locked\_super  
loop: WARNING in sysfs\_remove\_group  
lib, fs, cgroup: WARNING in percpu\_ref\_kill\_and\_confirm  
ata: WARNING in ata\_qc\_issue  
security, hugetlbfbs: write to user memory in  
hugetlbfbs\_destroy\_inode  
netlink: NULL timer crash  
kvm: use-after-free function call in kvm\_io\_bus\_destroy  
sound: use-after-free in snd\_seq\_cell\_alloc  
usb: use-after-free write in usb\_hcd\_link\_urb\_to\_ep  
net/kcm: double free of kcm inode  
crypto: out-of-bounds write in pre\_crypt  
security: double-free in superblock\_doinit  
kvm: WARNING in kvm\_apic\_accept\_events  
tcp: fix potential double free issue for fastopen\_req  
net/udp: slab-out-of-bounds Read in udp\_rcvmsg  
net: deadlock between ip\_expire/sch\_direct\_xmit  
srcu: BUG in \_\_synchronize\_srcu  
net/sctp: recursive locking in sctp\_do\_peeloff

vmx: WARNING in vmx\_handle\_exit  
futex: use-after-free in futex\_wait\_requeue\_pi  
kvm/arm64: use-after-free in kvm\_vm\_ioctl/vmcache\_update  
kvm/arm64: use-after-free in  
kvm\_unmap\_hva\_handler/unmap\_stage2\_pmds  
local privilege escalation flaw in n\_hdlc CVE-2017-2636  
netlink: GPF in netlink\_unicast  
perf: use-after-free in perf\_release  
net/ipv6: null-ptr-deref in ip6mr\_sk\_done  
bpf: kernel NULL pointer dereference in map\_get\_next\_key  
crypto: deadlock between  
crypto\_alg\_sem/tnl\_mutex/genl\_mutex  
kvm: use-after-free in  
vmx\_check\_nested\_events/vmcs12\_guest\_cr0  
sound: another deadlock in snd\_seq\_pool\_done  
rcu: WARNING in rcu\_seq\_end  
fs: use-after-free in path\_lookupat  
ucount: use-after-free read in inc\_uaccount & dec\_uaccount  
net/ipv4: division by 0 in tcp\_select\_window  
net: heap out-of-bounds in  
fib6\_clean\_node/rt6\_fill\_node/fib6\_age/fib6\_prune\_clone  
mm: use-after-free in zap\_page\_range  
net/kcm: use-after-free in kcm\_wq  
idr: use-after-free write in ida\_get\_new\_above  
sg: stack-out-of-bounds write in sg\_write CVE-2017-7187  
crypto: WARNING in group\_key\_share  
net/rds: use-after-free in rds\_sendmsg\_common  
net: sleeping function called from invalid context in  
net\_enable\_timestamp  
net: use-after-free in neigh\_timer\_handler/sock\_wfreenet/sctp: use-after-free in sctp\_association\_put  
fs: use-after-free in userfaultfd\_exit  
net/ipv4: inconsistent lock state in  
tcp\_conn\_request/inet\_ehash\_insert  
net/ipv4: suspicious RCU usage in ip\_ra\_control  
net/ipv4: deadlock in ip\_ra\_control  
net/dccp: dccp\_create\_openreq\_child freed held lock  
nested\_vmx\_merge\_msr\_bitmap  
ipc: use-after-free in shm\_get\_unmapped\_area  
sounds: deadlocked processed in snd\_seq\_pool\_done  
net/atm: vcc\_sendmsg calls kmem\_cache\_alloc in non-blocking context  
ata: WARNING in ata\_sff\_qc\_issue  
net/rds: use-after-free in inet\_create  
mm: fault in \_\_do\_fault  
kvm: WARNING in nested\_vmx\_vmexit  
net: GPF in rt6\_nexthop\_info  
sound: spinlock lockup in snd\_timer\_user\_tinterrupt  
mm: GPF in bdi\_put  
net/sctp: use-after-free in sctp\_hash\_transport  
net/bridge: warning in br\_fdb\_find  
net/ipv6: null-ptr-deref in ip6\_route\_del/lock\_acquire  
net: possible deadlock in skb\_queue\_tail

# ~14 of the bugs we found

mm: double-free vulnerability (local root) CVE-2017-6074  
net: warning in inet\_sock\_destruct  
net/pppt: use-after-free in dst\_release  
net/udp: slab-out-of-bounds in udp\_rcvmsg/do\_csum  
CVE-2017-6347  
WARNING in skb\_warn\_bad\_offload  
tty: panic in tty\_ldisc\_restore  
net: BUG in \_\_skb\_gso\_segment  
net/dccp: use-after-free in dccp\_feat\_activate\_values  
net/kcm: GPF in kcm\_sendmsg  
net/xfrm: stack out-of-bounds in xfrm\_flowi\_sort  
net/llc: BUG in llc\_sap\_state\_process/skb\_set\_owner\_r  
CVE-2017-6345  
net/llc: bug in llc\_pdu\_init\_as\_xid\_cmd/skb\_over\_panic  
net/packet: use-after-free in packet\_rcv\_fanout  
net: SOFTIRQ-safe -> SOFTIRQ-unsafe lock order detected in  
skb\_array\_produce  
net/ipv4: null-ptr-deref in  
udp\_mmem\_release/sk\_memory\_allocated\_sub  
net/sctp: null-ptr-deref in sctp\_put\_port/sctp\_endpoint\_destroy  
net/ipv4: warning in nf\_nat\_ipv4\_fn  
net/ipv6: double free in ipip6\_dev\_free  
sound: use-after-free in snd\_seq\_queue\_alloc  
tcp: divide error in transfer\_xor  
net/xfrm: use-of-uninit spinlock in xfrm\_policy\_flush  
net/dccp: use-after-free in dccp\_sendmsg  
net: round trip to head in  
net/icmp: null-ptr-deref in ping\_v4\_push\_pending\_frames  
net/kcm: WARNING in kcm\_write\_msgs  
tcp: avoid infinite loop in tcp\_splice\_read() CVE-2017-6214  
tun: read vnet\_hdr\_sz once  
macvtap: read vnet\_hdr\_size once  
udp: properly cope with csum errors  
ipv6: tcp: add a missing tcp\_v6\_restore\_cb()  
ip6\_gre: fix ip6gre\_err() invalid reads CVE-2017-5897  
ipv4: keep skb->dst around in presence of IP options  
CVE-2017-5970  
net: use a work queue to defer net\_disable\_timestamp() work  
netlabel: out of bound access in cipso\_v4\_validate()  
ipv6: pointer math error in ip6\_tnl\_parse\_tlv\_enc\_lim()  
net: heap out-of-bounds in ip6\_fragment  
tcp: fix 0 divide in \_\_tcp\_select\_window()  
keys: GPF in request\_key  
net/tcp: warning in sha512\_mmb\_mgr\_get\_comp\_job\_avx2  
crypto: NULL deref in sha512\_mmb\_mgr\_get\_comp\_job\_avx2  
sound: unable to handle kernel paging request  
snd\_seq\_prioq\_cell\_out  
scsi: BUG in scsi\_init\_io  
mm: sleeping function called from invalid context  
shmem\_undo\_range  
timerfd: use-after-free in timerfd\_remove\_cancel  
scsi: use-after-free in sg\_start\_req

mm: deadlock between get\_online\_cpus/pcpu\_alloc  
BUG at net/sctp/socket.c:7425  
kvm: use-after-free in irq\_bypass\_register\_consumer  
net: suspicious RCU usage in nf\_hook  
kvm: fix page struct leak in handle\_vmon CVE-2017-2596  
ipv6: fix ip6\_tnl\_parse\_tlv\_enc\_lim()  
kvm: WARNING in mmu\_spte\_clear\_track\_bits  
perf: use-after-free in perf\_event\_for\_each  
net: use-after-free in tw\_timer\_handler  
namespace: deadlock in dec\_pid\_namespaces  
sctp: kernel memory overwrite attempt detected in  
sctp\_getsockopt\_assoc\_stats  
kvm: deadlock in kvm\_vgic\_map\_resources  
net/atm: warning in alloc\_tlv/\_might\_sleep  
net/ipv6: use-after-free in sock\_wfreenet/kvm: kvm: BUG in loaded\_vmcs\_init  
kvm: NULL deref in vcpu\_enter\_guest  
kvm: use-after-free in complete\_emulated\_mmio  
CVE-2017-2584  
kvm: BUG in kvm\_unload\_vcpu\_mmu  
x86: warning in unwind\_get\_return\_address  
ipc: BUG: sem\_unlock unlocks non-locked lock  
kvm: WARNING in mmu\_spte\_clear\_track\_bits  
sctp: suspicious rcu\_dereference\_check() usage in  
sctp\_epaddr\_lookup\_transport  
usb/core: warning in  
usb\_create\_ep\_devs/sysfs\_create\_dir\_ns  
usb/gadget: warning in  
ep\_write\_iter/\_alloc\_pages\_nodemask  
kvm: use-after-free in process\_srcu  
kvm: assorted bugs after OOMs  
kvm: deadlock between  
kvm\_io\_bus\_register\_dev/kvm\_hv\_set\_msr\_common  
netlink: GPF in netlink\_dump  
fs, net: deadlock between bind/splice on af\_unix  
usb/gadget: slab-out-of-bounds write in dev\_config  
usb/gadget: warning in dummy\_free\_request  
usb/gadget: use-after-free in gadgetfs\_setup  
usb/gadget: GPF in usb\_gadget\_unregister\_driver  
net: use-after-free in worker\_thread  
net: signed overflows in SO\_[SEND|RCV]BUFFERFORCE sockopts  
CVE-2016-9793 CVE-2012-6704  
usb/gadget: warning in dev\_config/memdup\_user  
net/can: warning in raw\_setsockopt/\_alloc\_pages\_slowpath  
net/ipv6: null-ptr-deref in ip6\_rt\_cache\_alloc  
net/dccp: use-after-free in dccp\_invalid\_packet  
net/sctp: vmalloc allocation failure in  
sctp\_setsockopt/xt\_alloc\_table\_info  
net: BUG in unix\_notiflight  
net: GPF in eth\_header CVE-2016-9755  
net: deadlock on genl\_mutex  
net: GPF in rt6\_get\_cookie

Thank you!

Q&A

[github.com/google/syzkaller](https://github.com/google/syzkaller)  
[syzkaller@googlegroups.com](mailto:syzkaller@googlegroups.com)

Dmitry Vyukov, [dvyukov@](mailto:dvyukov@)

# Links

<https://github.com/google/syzkaller>

<https://www.kernel.org/doc/html/latest/dev-tools/kasan.html>

<https://github.com/google/kmsan>

<https://github.com/google/ktsan>

Backup

# Our Team: Dynamic Testing Tools

- User-space tools: ASAN, MSAN, TSAN
- Kernel tools: KASAN, KMSAN, KTSAN
- Hardening: CFI, SafeStack
- Fuzzing: LibFuzzer, syzkaller, OSS-Fuzz



# KMSAN: What one might think it does

```
int a;  
b = c + a; // report reading of uninit a
```

or:

```
int a = x;  
copy_to_user(p, &a, 4); // don't report since a was initied
```

This is useless: both false positives and false negatives!

# KMSAN: What is does

```
int a;  
if (flag)  
    a = 1;           // initialize  
b = c + a;         // not a "use"  
if (flag)  
    copy_to_user(p, &b, 4); // use: don't report
```

or:

```
int x;              // not initialized  
int a = x;         // still not initialized  
copy_to_user(p, &a, 4); // use: report
```

# Coverage-guiding in action

```
if (input[0] == '{') {  
    if (input[1] == 'i' && input[2] == 'f') {  
        if (input[3] == '(') {  
            input[input[4]] = input[5]; // potential OOB write  
        }  
    }  
}
```

Requires "{if(" input to crash,  $\sim 2^{32}$  guesses to crack when blind.

Coverage-guiding:

Guess "{" in  $\sim 2^8$ , add to corpus.

Guess "{i" in  $\sim 2^8$ , add to corpus.

Guess "{if" in  $\sim 2^8$ , add to corpus.

Guess "{if(" in  $\sim 2^8$ , add to corpus.

Total:  $\sim 2^{10}$  guesses.

# Syscall Descriptions: discriminated syscalls

Can describe families of discriminated syscalls:

```
fcntl$setflags(fd fd, cmd const[F_SETFD], flags flags[fcntl_flags])  
fcntl$setstatus(fd fd, cmd const[F_SETFL], flags flags[fcntl_status])  
fcntl$getown(fd fd, cmd const[F_GETOWN]) pid  
fcntl$setown(fd fd, cmd const[F_SETOWN], pid pid)  
fcntl$setsig(fd fd, cmd const[F_SETSIG], sig signalno)  
fcntl$setlease(fd fd, cmd const[F_SETLEASE], typ flags[flock_type])
```

# Syscall Descriptions: structs/unions

```
fcntl$getownex(fd fd, cmd const[F_GETOWN_EX], arg ptr[out, f_owner_ex])
```

```
# struct:
```

```
f_owner_ex {  
    type    flags[f_owner_type, int32]  
    pid     pid  
}
```

```
# union:
```

```
tun_buffer [  
    pi     tun_pi  
    hdr    virtio_net_hdr  
] [varlen]
```

# Syscall Descriptions: resources

```
resource fd_bpf_map[fd]
```

```
resource fd_bpf_prog[fd]
```

```
bpf$MAP_CREATE(cmd const[BPF_MAP_CREATE], ...) fd_bpf_map
```

```
bpf_map_lookup_arg {  
    map      fd_bpf_map  
    key      buffer[in]  
    val      buffer[out]  
}
```