



black hat[®]
EUROPE 2018
DECEMBER 3-6, 2018
EXCEL LONDON / UNITED KINGDOM



Drill the Apple Core: Up & Down

Fuzz Apple Core Component in Kernel and User Mode for Fun and Profit

 #BHEU / @BLACKHATEVENTS



Juwei Lin

- @panicall
- Joined TrendMicro Since 2013
- Windows Kernel/Rootkit/Bootkit
- Ransomware Decryption
- iOS/Android/Mac Vulnerability Hunting



Lilang Wu

- @Lilang_Wu
- Joined Trend Micro Since 2016
- Mac/iOS Vulnerability/Malware
- iOS/Android Exploit Detection



Moony Li

- @Flyic
- 8 years security
- Sandcastle
- Deep Discovery
- Exploit Detection
- Mac/Windows Kernel
- iOS/Android Vulnerability

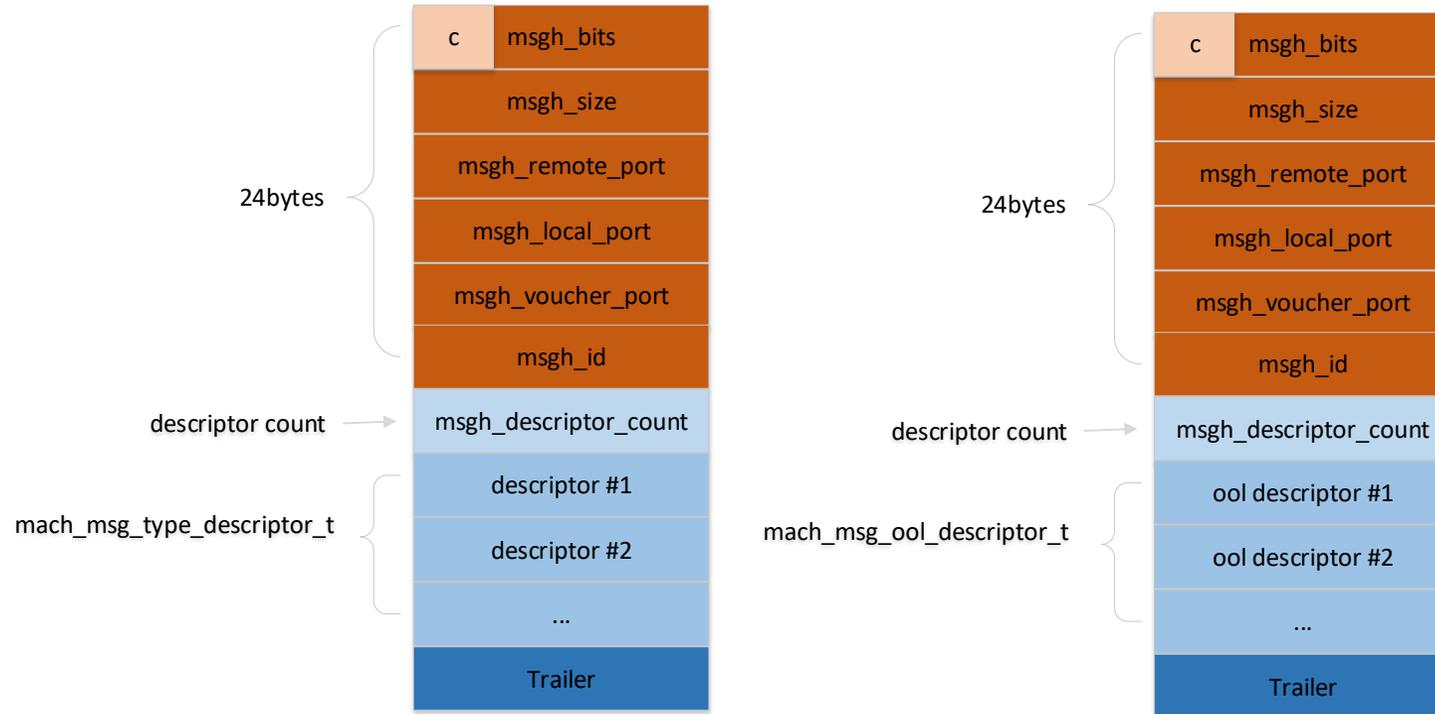
- Smart Fuzz XPC
 - XPC Internals
 - Fuzz Strategy
 - Reproduce Strategy
 - Case study

- Smart Fuzz XNU
 - Introduction
 - Architecture and Sanitizer Support
 - Syntax Engine and Corpus
 - Sanitizers
 - Root Case Study

Smart Fuzzing XPC

• What is XPC?

- low-level (libSystem) interprocess communication mechanism
- simple messages and complex messages





- Message Binary Format

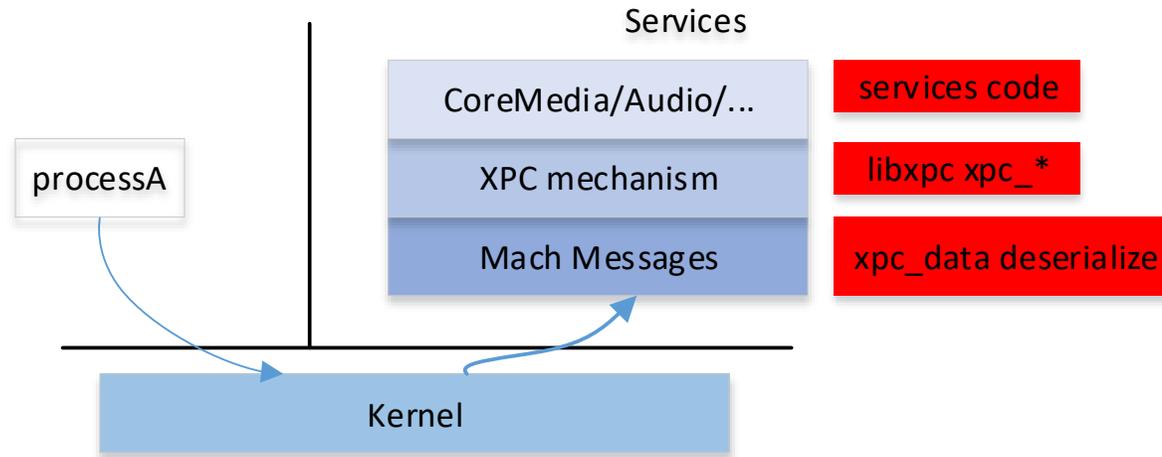
```
(lldb) c
Process 84781 resuming
Process 84781 stopped
* thread #1, queue = 'com.apple.main-thread', stop reason = breakpoint 1.1
  frame #0: 0x00007fff5c41f6e8 libsystem_kernel.dylib`mach_msg
libsystem_kernel.dylib`mach_msg:
-> 0x7fff5c41f6e8 <+0>: pushq  %rbp
   0x7fff5c41f6e9 <+1>: movq   %rsp, %rbp
   0x7fff5c41f6ec <+4>: pushq %r15
   0x7fff5c41f6ee <+6>: pushq %r14
Target 0: (nsxpc_client) stopped.
(lldb) x/10g $rdi mach_msg_header_t mach_msg_type_descriptor_t
0x100204728: 0x0000007480110013 0x00000000000001003
0x100204738: 0x10000000000001807 0x0000130700000001
0x100204748: 0x0011000000000000 0x00000000540585043
0x100204758: 0x00000003c0000f000 0x00000000000000003
0x100204768: 0x0000000100004000 0x746f6f7200000000
```

这里可以看到，这个serial...
0x34000。而事实上，在d...
mach_msg_port_descripto...
mach_msg_body_tort...
magic value /version
dictionary data



• Attack Surface

- serialize/deserialize
- libxpc
- services code



- How to trigger these bugs?



Crafted Mach Message



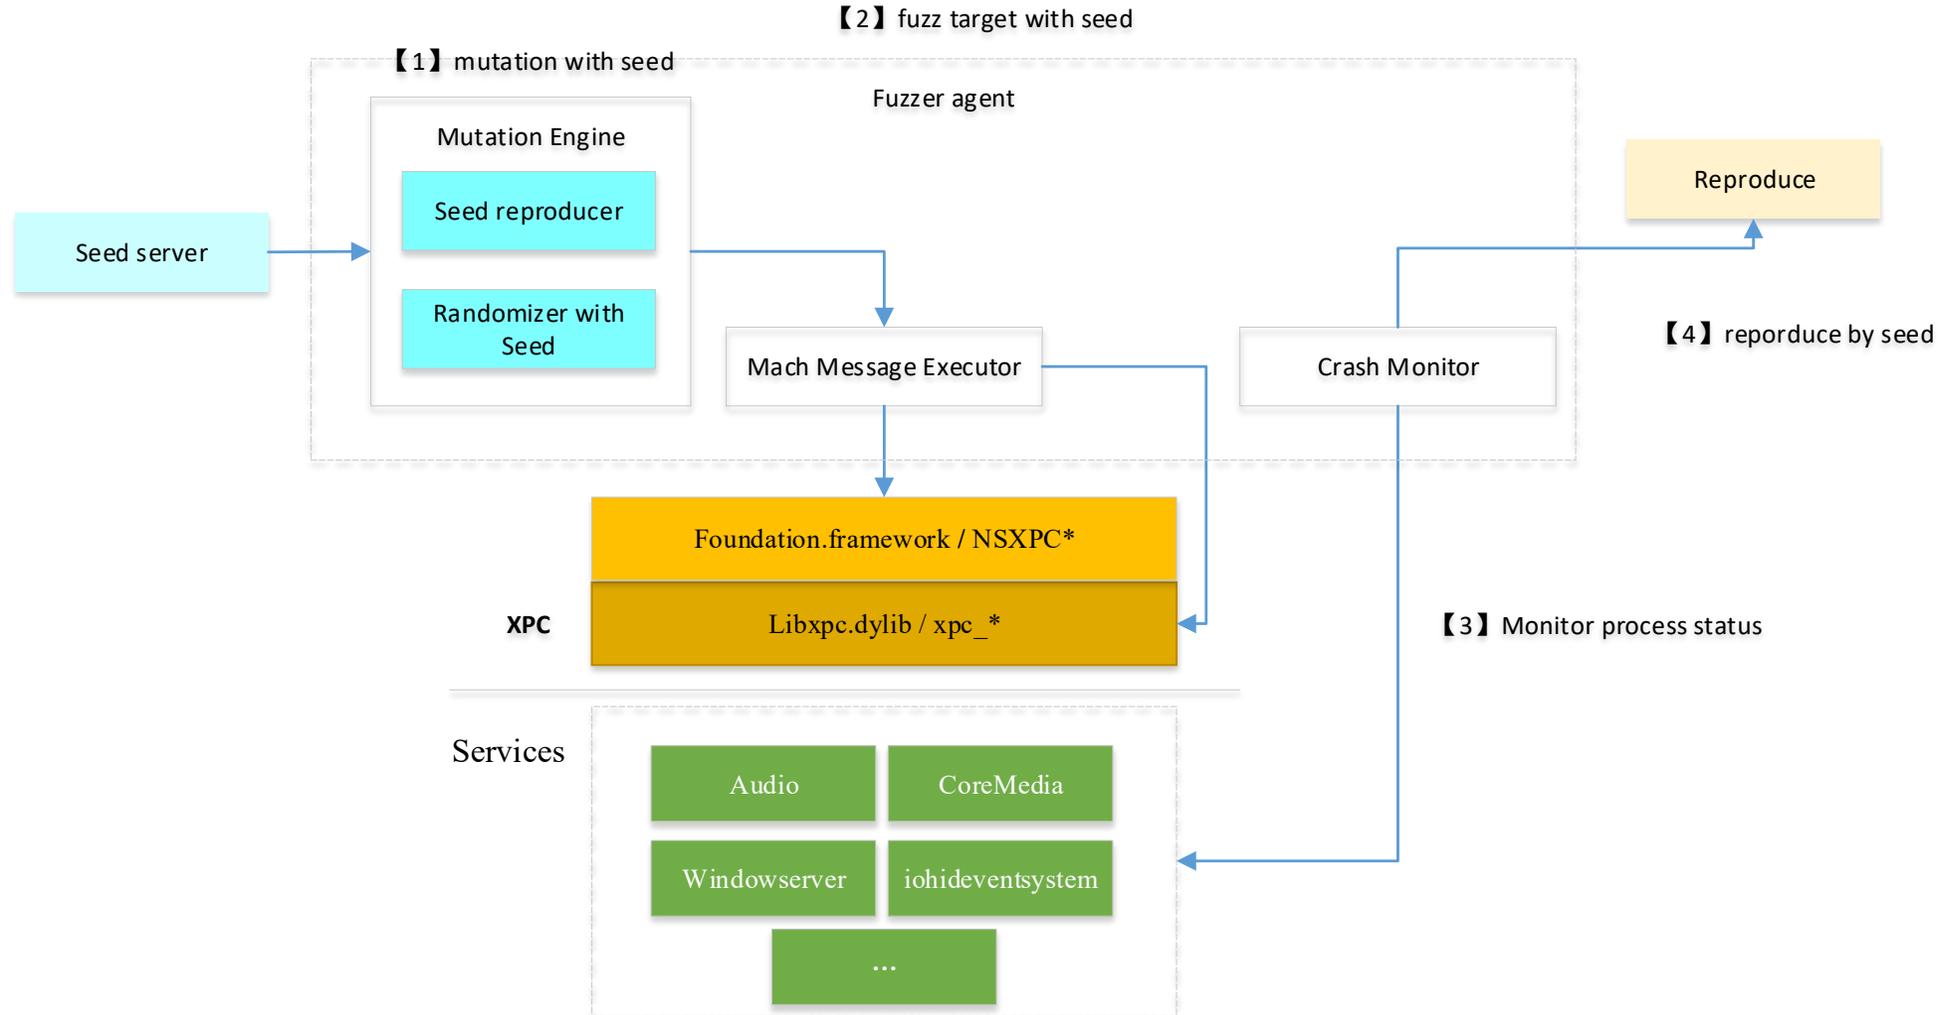
- Proactive fuzz

```
Process 84781 resuming
Process 84781 stopped
* thread #1, queue = 'com.apple.main-thread', stop reason = breakpoint 1.1
  frame #0: 0x00007fff5c41f6e8 libsystem_kernel.dylib`mach_msg
libsystem_kernel.dylib`mach_msg:
-> 0x7fff5c41f6e8 <+0>: pushq  %rbp
   0x7fff5c41f6e9 <+1>: movq   %rsp, %rbp
   0x7fff5c41f6ec <+4>: pushq  %r15
   0x7fff5c41f6ee <+6>: pushq  %r14
Target 0: (nsxpc_client) stopped.
(lldb) x/10g $rdi
0x100204728: 0x0000007480110013 0x00000000000001003
0x100204738: 0x10000000000001807 0x00000130700000001
0x100204748: 0x0011000000000000 0x00000000540585043
0x100204758: 0x0000003c00001000 0x00000006600000003
0x100204768: 0x0000000100000400 0x700000005f720000000
(lldb)
0x100204778: 0x0000000000000000 0x00000000000000000
0x100204788: 0x0000000000000000 0x00000000000000000
0x100204798: 0x0000000000000000 0x00000000000000000
0x1002047a8: 0x0000000000000000 0x00000000000000000
0x1002047b8: 0x0000000000000000 0x00000000000000000
```

- 1) body count
- 2) message descriptor
- 3) dictionary data

- Fuzz Strategy
 - Easy to control
 - Easy to mutate
 - Easy to monitor
 - Easy to reproduce

XPC Fuzz Architecture



- Fuzzing Target

- XPC Services

- `launchctl dumpstate`

```
xpc_connection_t client_con = xpc_connection_create_mach_service("com.apple.xpc.example", NULL, 0);
xpc_connection_set_event_handler(client_con, ^(xpc_object_t event) {
    //event handler
});
xpc_connection_resume(client_con);
```

```
services = {
  xpc 0 0 com.apple.wifiFirmwareLoader
  __init__ 64 - com.apple.uninstall
  OsxFuzz 0 py - com.apple.tzlinkd
  run.sh 0 - com.apple.storedownload.daemon
  xpcsConf 0 - com.apple.rpmuxd
  __init__ 65 - com.apple.nis.ybind
  fianlService 0.txt - com.apple.kextd
  fontdConf 66 y - com.apple.Kerberos.digest-service
  generalXpc 0 conf.py - com.apple.kcproxy
  sysmondConf 0 py - com.apple.fsevents
  xpcConf 0 (pe) com.apple.diagnosticextensions.osx.timemachine.helper
  XPCService 0ool.py 0 com.apple.diagnosticextensions.osx.spotlight.helper
  __init__ 0 0 com.apple.CoreRAID
  launchctl_du 0state.txt com.apple.CoreAuthentication.daemon
  machMsg.py 0 0 com.apple.DesktopServicesHelper.151FBB7D-869B-49E0-8EB2-2F509E9F92A6
  OsxFuzz 41779 (pe) com.apple.DesktopServicesHelper.726D2776-BA99-4F51-B49E-06474EF7B673
  com.apple.systempreferences.cacheAssistant
  com.apple.TrustEvaluationAgent.system
  com.apple.newsyslog
  com.apple.mediaremoted
  eservicesd
  amountd
  d
  clientDisplayAgent
  chbarserver
  rmald
  kgated Provider
  reagent.daemon
  oteDesktop.PrivilegeProxy
  id
  com.apple.mds.trampoline
  com.apple.GSSCred
  com.apple.FileCoordination
  com.apple.FileCoordination
  com.apple.colorsync.displayservices
  com.apple.avbdevice
  com.apple.audio.systemsoundserv
  com.apple.DesktopServicesHelper.79AD4457-5922-46E8-8A5B-90069A99D910
  com.apple.DesktopServicesHelper.3EA89BE3-09D0-4298-AA2A-5D1D384733CF
  com.apple.signpost.signpost-notificati
  org.macosforge.xquartz.privileged_startx
  com.apple.xpc.smd
  com.apple.wifilocityd
  com.apple.systemstats.analysis
  com.apple.installandsetup.systemmigrationd
  com.apple.gkreport
  com.apple.FontWorker
  com.apple.eapolcfg_auth
  com.apple.diagnosticd
  com.apple.rpmuxd = {
    active count = 0
    path = /System/Library/L
    state = waiting
    program = /usr/libexec/r
    arguments = {
      /usr/libexec/rpmuxd
    }
    launch environment = {
      PATH => /usr/bin:/bi
    }
    xpc_service_name =>
  }
  domain = com.apple.xpc.t
  minimum runtime = 10
  exit timeout = 5
  runs = 0
  successive crashes = 0
  excessive crashing = 0
  last exit code = (never
  event triggers = {
  }
  endpoints = {
    "com.apple.rpmuxd" =
      port = 0x1be03
      active = 0
      managed = 1
      reset = 0
      hide = 0
  }
  dynamic endpoints = {
  }
  provider
  pid-local endpoints = {
  }
  instance-specific endpoi
}
```

- Fuzz Controller

- ✓ Wrap the xpc interfaces by python

```
BOOST_PYTHON_MODULE(xpcconnection) {  
    PyEval_InitThreads();  
  
    class <XpcConnection boost::noncopyable>("XpcConnection", init<std:  
        .def("XpcCreateConnection", &XpcConnection::XpcCreateConnection)  
        .def("mach_connect", &XpcConnection::mach_connect_)  
        .def("XpcHandler", pure_virtual(&XpcConnection::handler))  
  
        .def("mach_msg", &XpcConnection::mach_msg_)  
        .def("XpcSendMessage", &XpcConnection::XPCSendMessage)  
    ;  
};
```

- ✓ Fuzz by python script

```
XpcConnectionBase::XpcConnectionBase(std::string target) :  
    serviceName(target) {  
    dispatchQueue = dispatch_queue_create(serviceName.c_str(), 0);  
    xpcConnection = xpc_connection_create_mach_service(serviceName.c_str(), dispatchQueue, XPC_CONNE  
  
    xpc_connection_set_event_handler(xpcConnection, ^(xpc_object_t event) {  
        // this seems to fix a segfault (it's never released, so this should probably be a problem..  
        auto gstate = PyGILState_Ensure();  
  
        xpc_retain(event);  
  
        // handle the event  
        handleEvent(event);  
        PyGILState_Release(gstate);  
    });  
  
    xpc_connection_resume(xpcConnection);  
};
```

• Crash Monitor

- Monitor the processes IDs cluster status
- Monitor exits signal value

```
zuffdemac-pro:~ zuff$ launchctl list
PID      Status  Label
-        0       com.apple.SafariHistoryServiceAgent
307      0       com.apple.Finder
336      0       com.apple.homed
578      0       com.apple.SafeEjectGPUAgent
-        0       com.apple.quicklook
-        0       com.apple.parentalcontrols.check
-        0       com.apple.PackageKit.InstallStatus
345      0       com.apple.mediaremoteagent
-        0       com.apple.FontWorker
321      0       com.apple.bird
-        0       com.apple.familycontrols.useragent
-        0       com.apple.AssetCache.agent
666      0       com.apple.universalaccessAuthWarn
312      0       com.apple.nsurlsessiond
-        0       com.apple.mobileactivationd
-        0       com.apple.syncservices.uihandler
352      0       com.apple.iconservices.iconservicesagent
```

No	Name	Default Action	Description
1	SIGHUP	terminate process	terminal line hangup
2	SIGINT	terminate process	interrupt program
3	SIGQUIT	create core image	quit program
4	SIGILL	create core image	illegal instruction
5	SIGTRAP	create core image	trace trap
6	SIGABRT	create core image	abort program (formerly SIGIOT)
7	SIGEMT	create core image	emulate instruction executed
8	SIGFPE	create core image	floating-point exception
9	SIGKILL	terminate process	kill program
10	SIGBUS	create core image	bus error
11	SIGSEGV	create core image	segmentation violation
12	SIGSYS	create core image	non-existent system call invoked
13	SIGPIPE	terminate process	write on a pipe with no reader
14	SIGALRM	terminate process	real-time timer expired
15	SIGTERM	terminate process	software termination signal
16	SIGURG	discard signal	urgent condition present on socket
17	SIGSTOP	stop process	stop (cannot be caught or ignored)
18	SIGTSTP	stop process	stop signal generated from keyboard
19	SIGCONT	discard signal	continue after stop
20	SIGCHLD	discard signal	child status has changed
21	SIGTTIN	stop process	background read attempted from control terminal
22	SIGTTOU	stop process	background write attempted to control terminal
23	SIGIO	discard signal	I/O is possible on a descriptor (see <code>fcntl(2)</code>)
24	SIGXCPU	terminate process	cpu time limit exceeded (see <code>setrlimit(2)</code>)



- Comparison between different Reproduce Methods

	Typical Example	Storage Cost	Speed Cost	Support Complex Scenario	Reproduce Rate	Dev Effort
Log	Trinity	High (Execution Log)	High	Low	Low	Low
Case(File)	AFL	Middle (Files Causing Crash)	Low	Middle	Middle	High
Crash Dump	-	High (Every Crash Context)	High	-	Very Low	No
Seed	JS Fun Fuzz	Low (Integer)	Low	High	High	Low



- Case Study - CVE-2018-4411

- libATSServer can read out of boundary for the latest MacOS

Target 0: (fontd) stopped.

(lldb) bt

* thread #1, queue = 'com.apple.main-thread', stop reason = EXC_BAD_ACCESS (code=1, address=0x7ffee1934000)

* frame #0: 0x00007fff55a06f49 libsystem_platform.dylib`_platform_memmove\$VARIANT\$Haswell + 41

frame #1: 0x00007fff2b8b597a libATSServer.dylib`FODBWriteToAnnex + 246

frame #2: 0x00007fff2b8d0157 libATSServer.dylib`HandleFontManagementMessage + 5403

frame #3: 0x00007fff2b8cd2d1 libATSServer.dylib`serverMainHandler(__CFMachPort*, FontMgrMessage*, long, void*) + 263

frame #4: 0x00007fff2d3e4596 CoreFoundation`__CFMachPortPerform + 310

frame #5: 0x00007fff2d3e4449 CoreFoundation`__CFRUNLOOP_IS_CALLING_OUT_TO_A_SOURCE1_PERFORM_FUNCTION__ + 41

frame #6: 0x00007fff2d3e4395 CoreFoundation`__CFRunLoopDoSource1 + 533

frame #7: 0x00007fff2d3dbf50 CoreFoundation`__CFRunLoopRun + 2848

frame #8: 0x00007fff2d3db1a3 CoreFoundation`CFRunLoopRunSpecific + 483

frame #9: 0x00007fff2d419c33 CoreFoundation`CFRunLoopRun + 99

frame #10: 0x00007fff2b8cc91c libATSServer.dylib`main_handler + 4510

frame #11: 0x00007fff556f5015 libdyld.dylib`start + 1

frame #12: 0x00007fff556f5015 libdyld.dylib`start + 1



```

// (__CFHachPort*, FontMgrMessage*, long, void*)
void __fastcall serverMainHandler(double a1, __int64 a2, __int64 a3)
{
    // ...
}
else
{
    v4 = HandleFontManagementMessage((FILE *)a3, &v10, a1); // a3=msg
    FOrmoveExceptionFrame(&v8, &v10);
    v5 = 1;
}

// ...
goto LABEL_044;
case 0x28:
    v82 = &v238->bf;
    if ( gUseNewFODB == 2 )
    {
        FODDBeginTransactions(9);
        if ( LODWORD(v82->_base) )
        {
            v83 = *(&v238->_lbfsz + 1);
            v84 = *(const char *)(&v238->_bf._base + 4);
            v84 = v238->_lbfsz;
        }
        else
        {
            a2 = (const char *)(&v238->_bf._size + 1);
            v83 = HIWORD(v238->_bf._base);
            v84 = v238->_bf._size;
        }
        FODDBAddAnnex(v83, a2, v84, 0, a3); // a2=buffer, v84=size -----c
        FODDBEndTransactions(9LL);
    }
}

// ...
|| (v12 = __ROL2__(*(__WORD *) (v11 + 22), 8), *(__WORD *) (gFontContainerLis
|| !gAnnexDB && (v8 = FODDBOpenAnnexFile(v10)) != 0 )
{
    result = (unsigned int)v8;
}
else
{
    result = FODDBWriteToAnnex(v7, a2, v6, v5, a5); // a2=buffer, a3=v6=size -----b
}
return result;
}

Microseconds((__int64)v54);
*(__QWORD *) (v12 + 3) = (v16 << 32) | v54[0];
v17 = __ROL2__(v13, 8);
LOWORD(v58) = v17;
v18 = __ROL2__(v14, 8);
HIWORD(v58) = v18;
v12[5] = v58;
v19 = v12;
memcpy(v12 + 6, a2, v53); // v53=a3=size -----a
if ( *(__BYTE *) (gAnnexHUXFile + 12LL) )
{
}

```

Out of boundary

Smart Fuzzing XNU

What I will introduce today

- Our Tool: PanicXNU
 1. Port Syzkaller to Support macOS XNU Fuzzing.
 2. Modify XNU to add support some features.
- Root Case Study



Fuzzer



- 530 BSD API Patterns
- VM Fusion Support
- macOS Executor

XNU

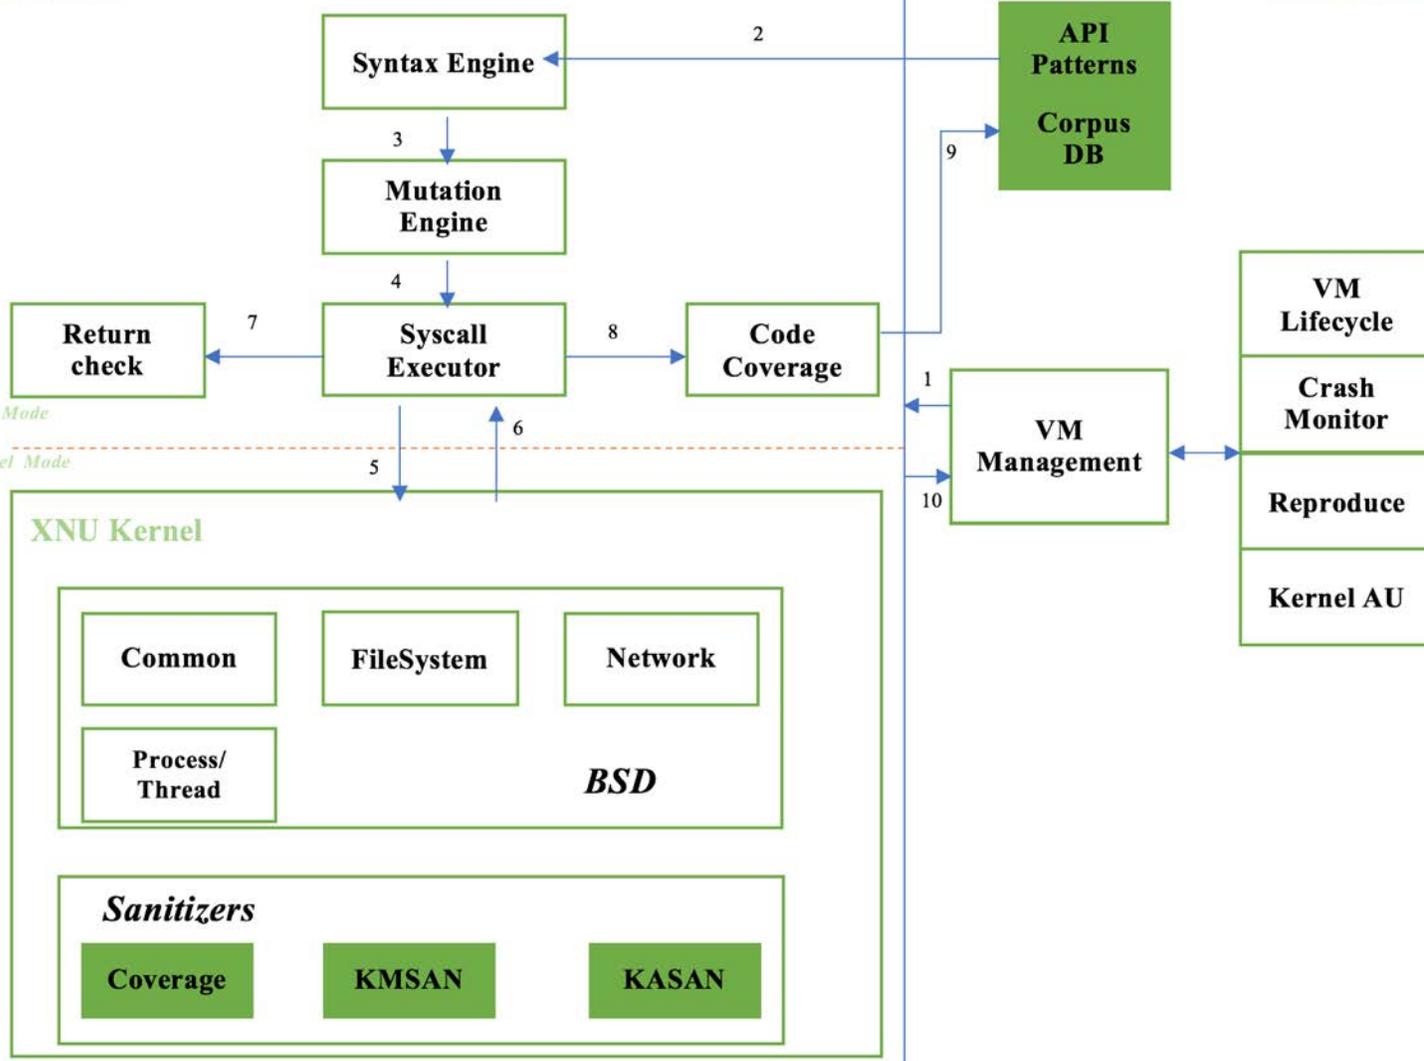


- Add Code Coverage
- Add Kernel Memory Sanitizer
- Enable Kernel Address Sanitizer



Client Side

Server Side



Architecture

1. Key modules are in GREEN
2. Also add some other modules, e.g. vmfusion

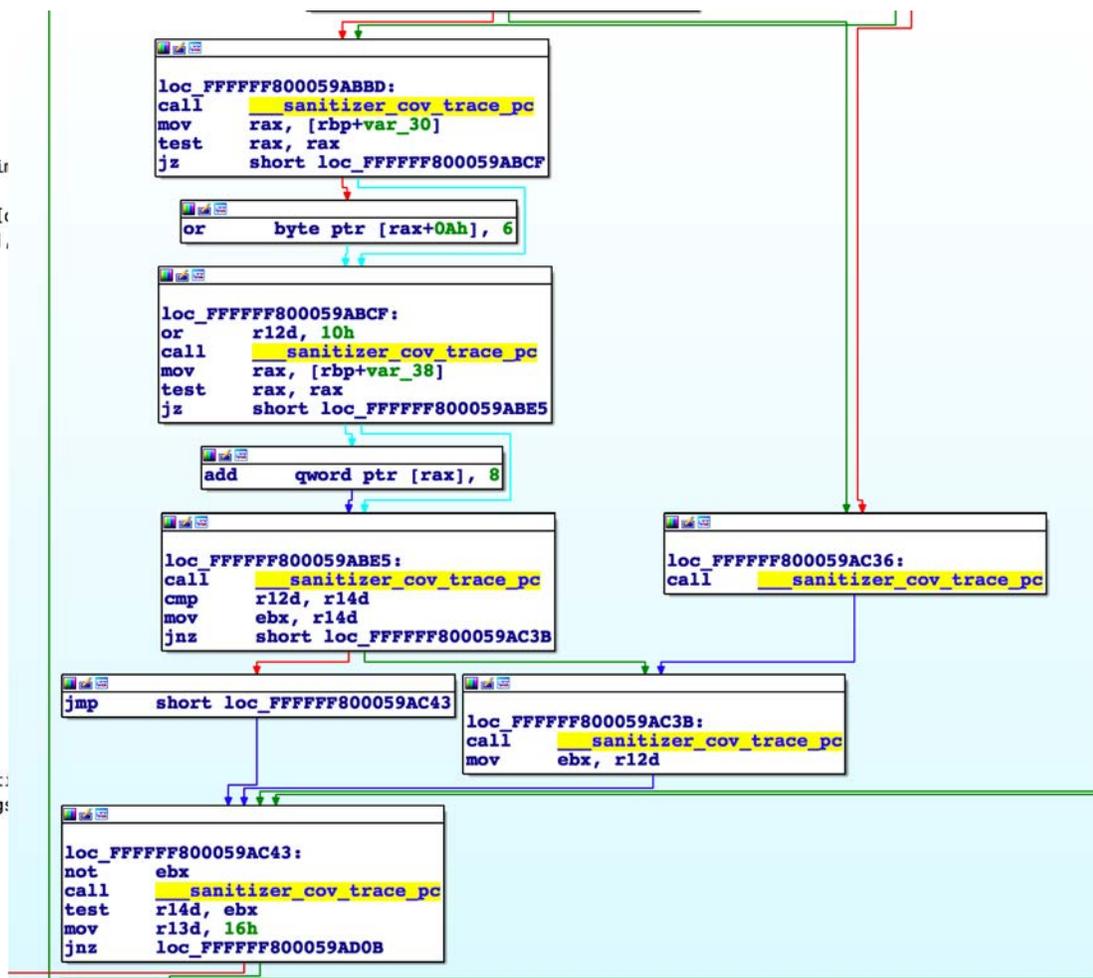


```

4 syscall
  gen
  dev_bpf_amd64.const
  dev_bpf.txt
  dev_dtrace_helper_amd64.const
  dev_dtrace_helper.txt
  dev_ptmx_amd64.const
  dev_ptmx.txt
  init.go
  ipc_amd64.const
  ipc.txt
  posix_fs_amd64.const
  posix_fs.txt
  posix_mm_amd64.const
  posix_mm.txt
  proc_thread_amd64.const
  proc_thread.txt
  ptrace_debug_amd64.const
  ptrace_debug.txt
  socket_amd64.const
  socket_inet_amd64.const
  socket_inet_icmp_amd64.const
  socket_inet_icmp.txt
  socket_inet_tcp_amd64.const
  socket_inet_tcp.txt
  socket_inet_udp_amd64.const
  socket_inet_udp.txt
  socket_inet.txt
  socket_inet6_amd64.const
92 setrlimit(res flags[rlimit_type], rlim ptr[in], rlimit)
93
94 sigaltstack(ss vma, oss ptr[out], intptr, opt)
95 getitimer(which flags[getitimer_which], cur ptr[out], itimerval)
96 setitimer(which flags[getitimer_which], new ptr[in], itimerval), old ptr[out], itir
97 exit(code intptr)
98 wait4(pid pid, status ptr[out], int32, opt), options flags[wait_options], ru ptr[
99 wait4_nocancel(pid pid, status ptr[out], int32, opt), options flags[wait_options],
100
101 kill(pid pid, signal intptr, posix intptr)
102 getlogin()
103 setlogin(name ptr[in], string)
104 acct(file ptr[in], filename)
105 umask(cmask flags[open_mode])
106 reboot(howto flags[reboot_flags])
107 revoke(path ptr[in], filename)
108 swapon(dummy int32)
109 gettid(uidp ptr[out], uid, gidp ptr[out], gid)
110 settid(uid uid, gid gid)
111 setegid(egid gid)
112 seteuid(euid uid)
113 getpriority(which flags[priority_flags], who intptr)
114 setpriority(which flags[priority_flags], who intptr, prio int32)
115 gettimeofday(tp ptr[out], timeval, tzp ptr[out], timezone)
116 gettimeofday(tp ptr[in], timeval, tzp ptr[in], timezone)
117 setsid() pid
118 futimes(fildes fd, times ptr[in], array[timeval, 2])
119 getsid(pid pid)
120 getfh(path ptr[in], filename, fhptr ptr[in], intptr)
121 sigaction(sig flags[sigaction_sig], act ptr[in], sigaction), oact ptr[out], sigact:
122 sigprocmask(how flags[sigprocmask_flags], set ptr[in], sigset), oset ptr[out], sig:
123 sigpending(set ptr[in], sigset)
124 getdtablesize()
125 sigsuspend(set ptr[in], sigset)
126 sigsuspend_nocancel(set ptr[in], sigset)
127 gethostuuid(id int16, wait ptr[in], timespec)

```

API Pattern



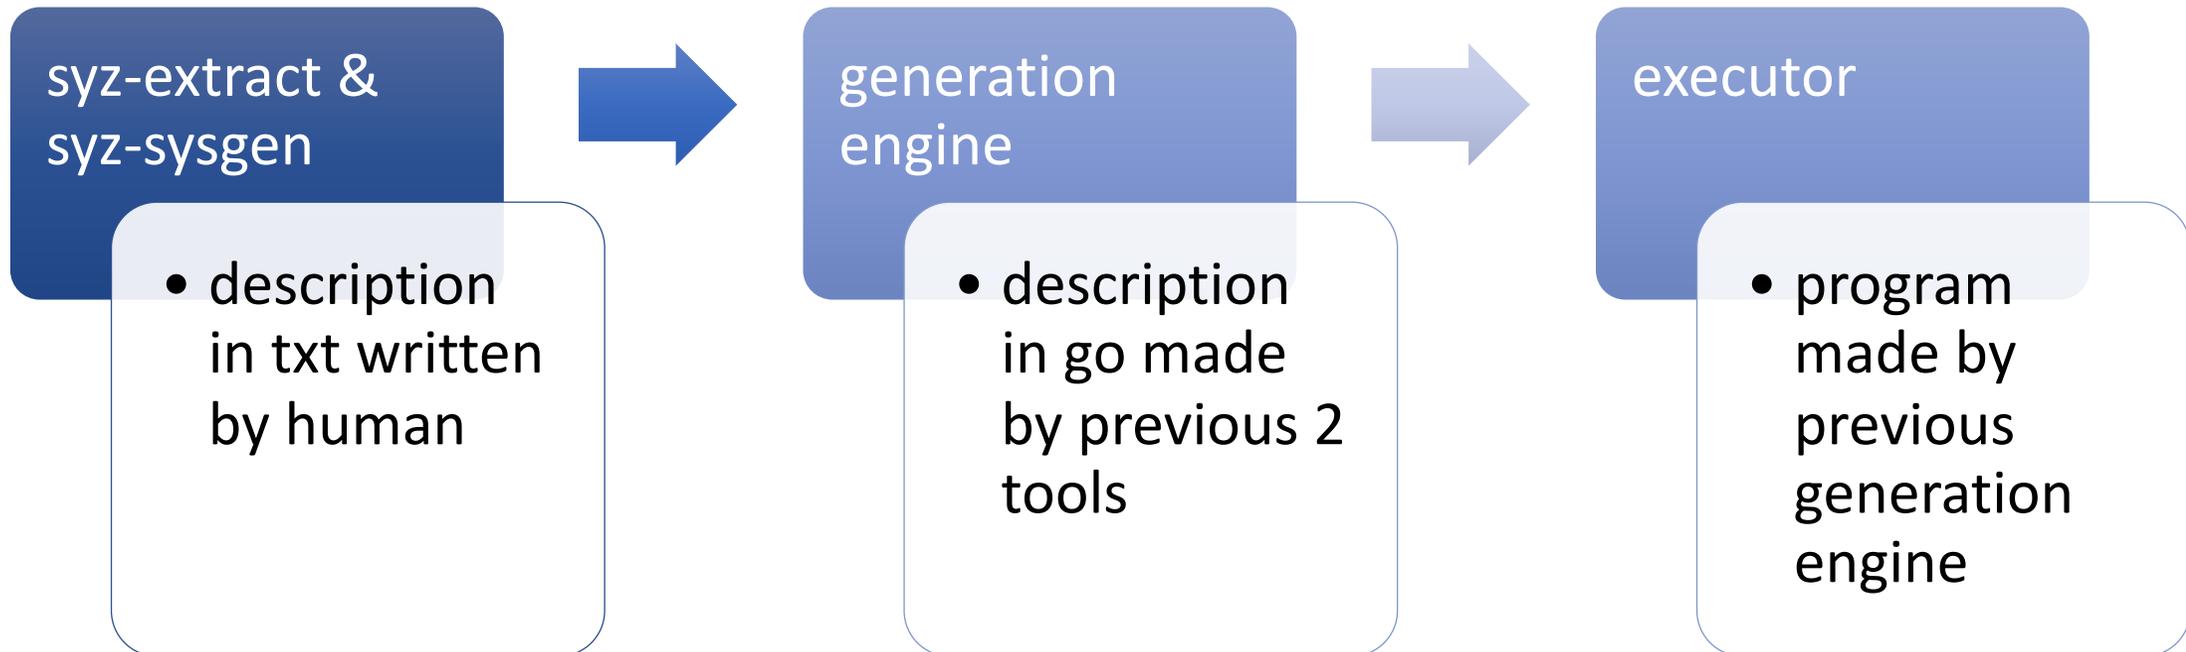
Code Coverage

My Efforts

- Syntax Engine is directly from Syzkaller; But I developed the XNU BSD API patterns.
- KASAN is from XNU, but it does not work well after compilation.
- I developed coverage sanitizer.
- I developed KMSAN.

Syntax Engine & Corpus

Quick glance at syzkaller's syntax engine



Corpus

- More than 500 syscalls in XNU kernel
- Refer to syzkaller's syscall descriptions syntax:
https://github.com/google/syzkaller/blob/master/docs/syscall_descriptions_syntax.md
- Refer to sample txt files in syzkaller project



Sanitizers



Kernel Mode Sanitizers

Name	Features	Comments
Kernel Sanitizer Coverage	<ul style="list-style-type: none">• get function/block/edge coverage	<ul style="list-style-type: none">• Has instrumentations support• NO existing callbacks implementation
KASAN (kernel address sanitizer)	<ul style="list-style-type: none">• Out-of-bounds accesses• Use-after-free• Use-after-return• Use-after-scope• Double-free, invalid free	<ul style="list-style-type: none">• Has instrumentations support• Has callbacks/module support
KMSAN (kernel memory sanitizer)	<ul style="list-style-type: none">• uninitialized reads	<ul style="list-style-type: none">• Not implemented

Sanitizer Coverage

- We need to develop a new module in XNU to:
 - Support sanitizer callback function
 - Read the coverage data back to user fuzzing program

Callback Implementation

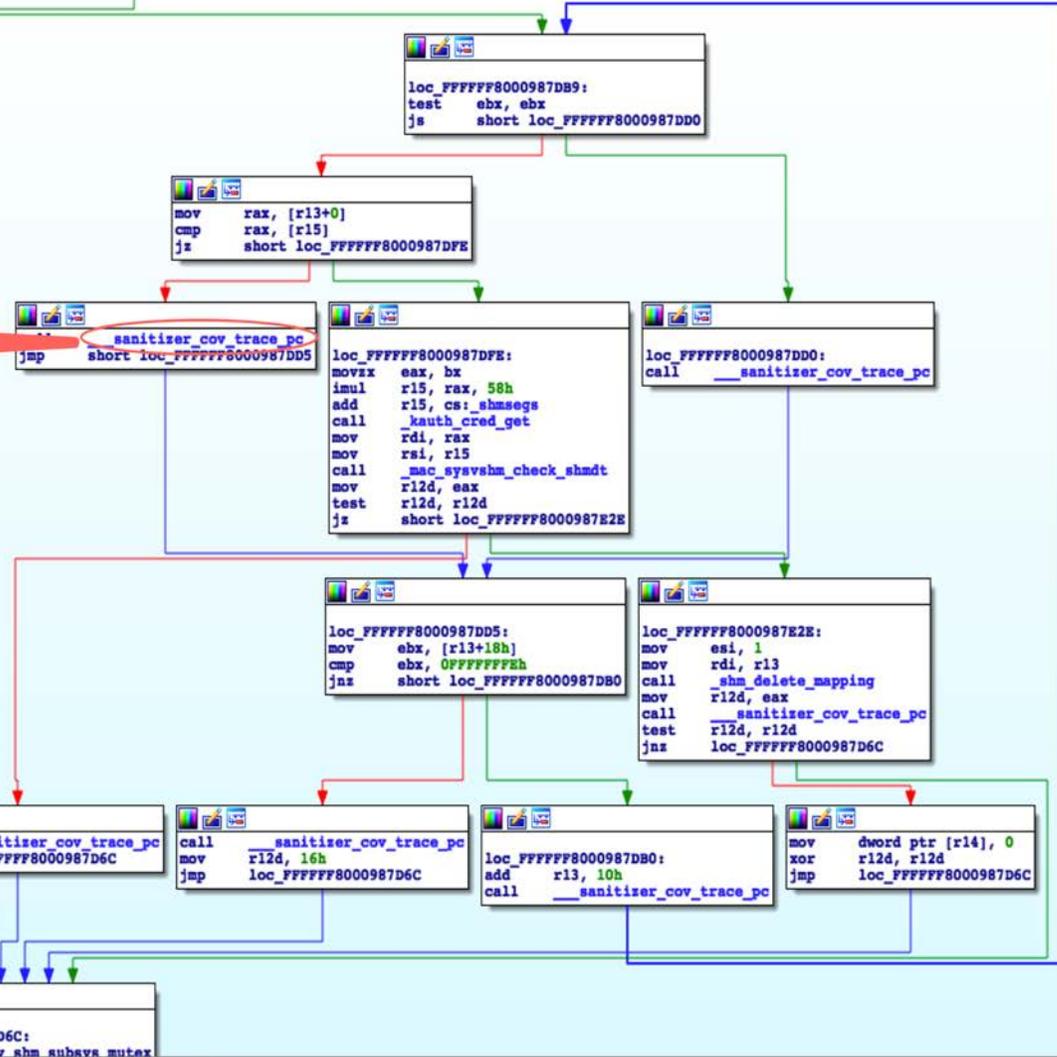
```
struct task {  
    ...  
  
    enum kcov_mode kcov_mode;  
    unsigned      kcov_size;  
    void          *kcov_area;  
    struct kcov *kcov;  
    uint32_t      refcount;  
}  
  
void __attribute__((noinline)) __sanitizer_cov_trace_pc()  
{  
    ...  
}
```

1. callback name:
 __sanitizer_cov_trace_pc
2. just support single-thread mode
3. store coverage structure into task_t



After Compilation

this function is my code coverage trace handler for XNU





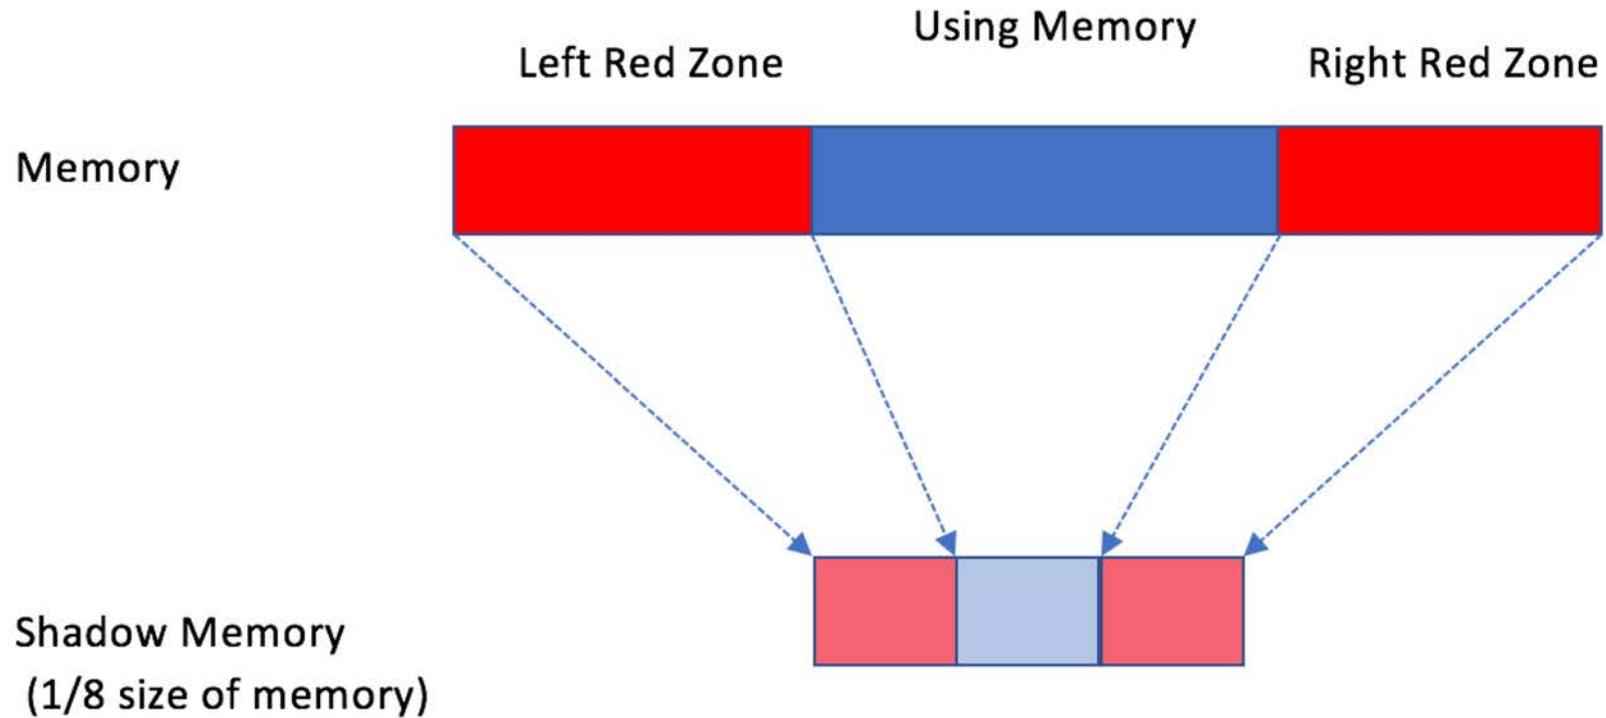
KASAN

- latest XNU has KASAN support
 - KDK now provides kernel.kasan which works well.
 - It does not work if you compile it, VM cannot boot.
- It consists of *guard pages, shadow memory and operations.*
- It can protect Globals, Stack and Heap memory.

How KASAN protects memory

- 1) memory operations are called, e.g. `__asan_strncpy`
- 2) `__asan_strncpy` checks shadow memory
- 3) KASAN panics the kernel if shadow memory is illegal (shadow value < 0)

Guard Pages & Shadow Memory





Operations

Heap Memory Operations	Stack Memory Operations	Other Memory Operations
<code>__asan_bcopy</code>	<code>__asan_stack_malloc_0</code>	<code>__asan_load1</code>
<code>__asan_memmove</code>	<code>__asan_stack_malloc_1</code>	<code>__asan_load2</code>
<code>__asan_memcpy</code>	<code>__asan_stack_malloc_2</code>	<code>__asan_load4</code>
<code>__asan_memset</code>	<code>__asan_stack_malloc_3</code>	<code>__asan_load8</code>
<code>__asan_bzero</code>	<code>__asan_stack_malloc_4</code>	<code>__asan_load16</code>
<code>__asan_bcmp</code>	<code>__asan_stack_malloc_5</code>	<code>__asan_loadN</code>
<code>__asan_memcmp</code>	<code>__asan_stack_malloc_6</code>	
<code>__asan_strncpy</code>	<code>__asan_stack_malloc_7</code>	
<code>__asan_strlcat</code>	<code>__asan_stack_malloc_8</code>	
<code>__asan_strncpy</code>	<code>__asan_stack_malloc_9</code>	
<code>__asan_strncat</code>	<code>__asan_stack_malloc_10</code>	
<code>__asan_strnlen</code>		
<code>__asan_strlen</code>		

`#define strncpy __asan_strncpy`

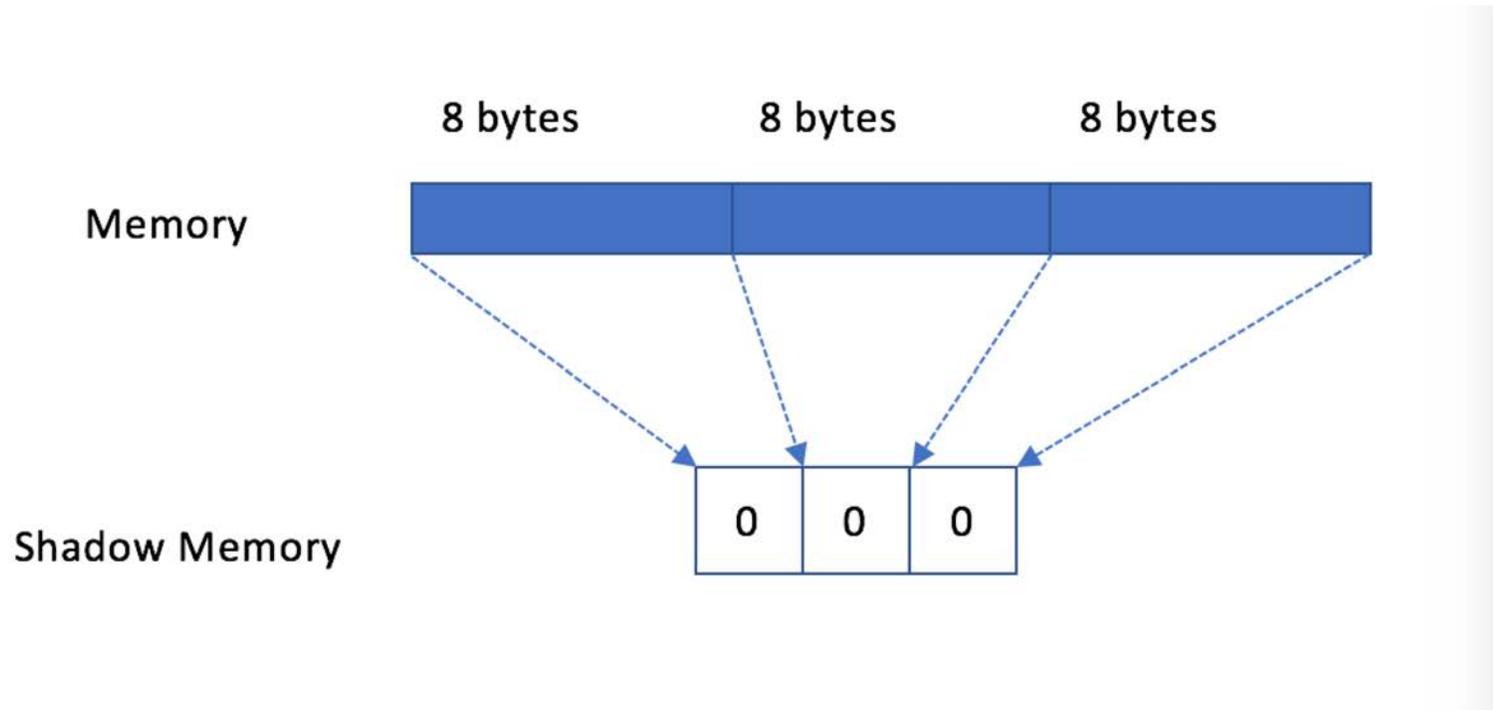
`-fsanitizer=address`

buildin calls in xnu source code



Example: Detect UAF

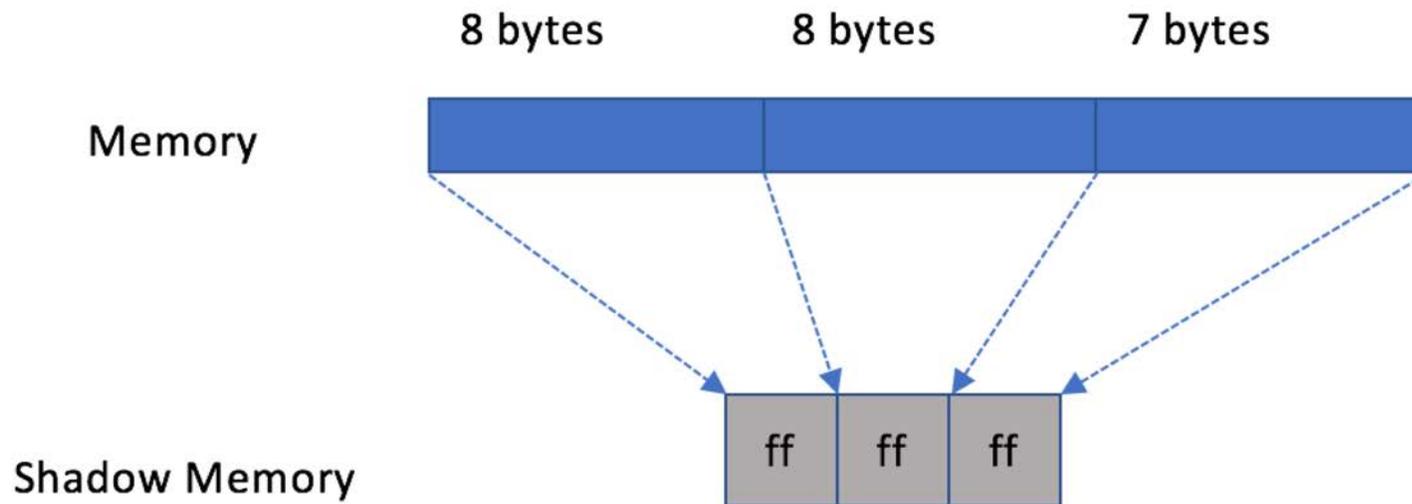
- When new memory is allocated and aligned with 8





Example cont. 1

- When the memory is freed



Example cont. 2

- When the memory is used after free, any related operation will check its shadow memory and then panic the system.
 - 0xff is illegal

KMSAN

- Kernel memory sanitizer is used to detect uninitialized memory.
- We worked on how to initialize all uninitialized memory allocated in kernel, e.g. `kalloc_canblock`

kalloc_canblock

```
    assert(size <= z->elem_size);  
  
#if VM_MAX_TAG_ZONES  
    if (z->tags && site)  
    {  
        tag = vm_tag_alloc(site);  
        if (!canblock && !vm_allocation_zone_totals[tag]) tag = VM_KERN_MEMORY_KALLOC;  
    }  
#endif  
  
    addr = zalloc_canblock_tag(z, canblock, size, tag);  
  
#if KASAN_KALLOC  
    /* fixup the return address to skip the redzone */  
    addr = (void *)kasan_alloc((vm_offset_t)addr, z->elem_size, req_size, KASAN_GUARD_SIZE);  
  
    /* For KASan, the redzone lives in any additional space, so don't  
     * expand the allocation. */  
#else  
    *psize = z->elem_size;  
#endif  
  
    // add by @panicall  
    if (addr)  
        memset(addr, 0xde, *psize);  
    return addr;  
}
```


Conclusion

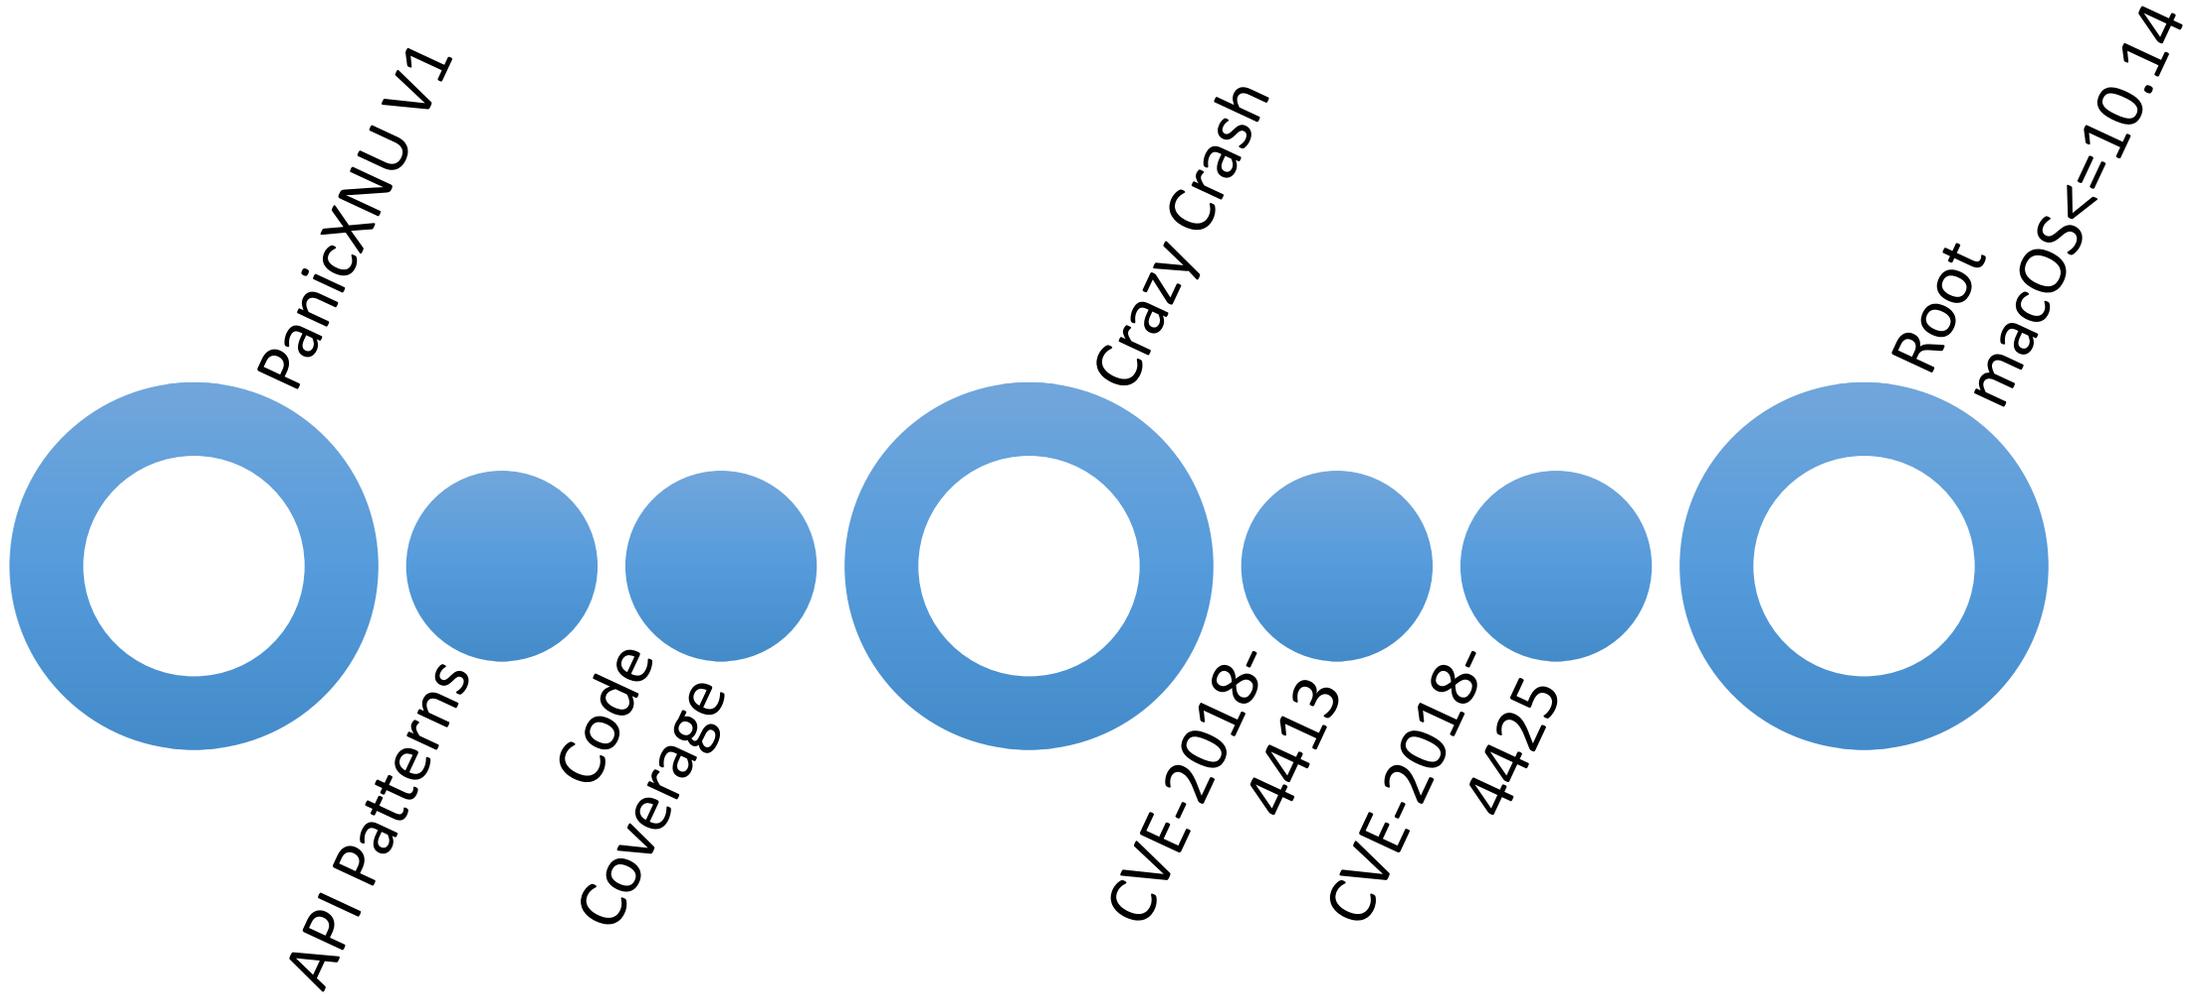


- About 530 API patterns
- Corpus

- Coverage Sanitizer
- KASAN
- KMSAN



macOS Root Case Study





CVE-2018-4413

- Uninitialized heap memory leak
- Fixed in macOS 10.14.1 and iOS 12.1
- Can be used to leak ipc_port object address

CVE-2018-4425

- NECP type confusion
- Fixed in macOS 10.14.1
- Can be used to write arbitrary kernel address
- Can be used to free arbitrary kernel address

```
STATIC int
sysctl_procargsx(int *name, u_int namelen, user_addr_t where,
                 size_t *sizep, proc_t cur_proc, int argc_yes)
{
    ...

    if ((u_int)arg_size > p->p_arghlen)
        arg_size = round_page(p->p_arghlen);           --- (a)

        arg_addr = p->user_stack - arg_size;

    ...

    ret = kmem_alloc(kernel_map, &copy_start, round_page(arg_size), VM_KERN_MEMORY_BSD);
    if (ret != KERN_SUCCESS) {
        vm_map_deallocate(proc_map);
        return(ENOMEM);
    }

    copy_end = round_page(copy_start + arg_size);

    if( vm_map_copyin(proc_map, (vm_map_address_t)arg_addr,
                     (vm_map_size_t)arg_size, FALSE, &tmp) != KERN_SUCCESS) {
        vm_map_deallocate(proc_map);
        kmem_free(kernel_map, copy_start,
                  round_page(arg_size));
        return (EIO);
    }

    /*
     * Now that we've done the copyin from the process'
     * map, we can release the reference to it.
     */
    vm_map_deallocate(proc_map);
}
```

CVE-2018-4413

sysctl_procargsx is used to retrieve process args information by calling sysctl.

at location (a) :

- p->p_arghlen is usually around 0x300;
- I set my arg_size to 0x200 so that arg_size will not be round_paged



CVE-2018-4413

At location (b):

- Stack information is copied to new allocated page at offset 0 with `arg_size` (0x200).
- The new allocated page is not zeroed. So this operation leaves the rest of this page filled with uninitialized heap data.

At location (c):

- `copy_end` is `round_paged`, parameter data points to the last 0x200 bytes of the page.

At location (d):

- `copyout` the 0x200 bytes leaked heap information to user buffer

```
if( vm_map_copy_overwrite(kernel_map,          --- (b)
                          (vm_map_address_t)copy_start,
                          tmp, FALSE) != KERN_SUCCESS) {
    kmem_free(kernel_map, copy_start,
              round_page(arg_size));
    vm_map_copy_discard(tmp);
    return (EIO);
}

if (arg_size > argslen) {
    data = (caddr_t) (copy_end - argslen);
    size = argslen;
} else {
    data = (caddr_t) (copy_end - arg_size);      --- (c)
    size = arg_size;
}

...

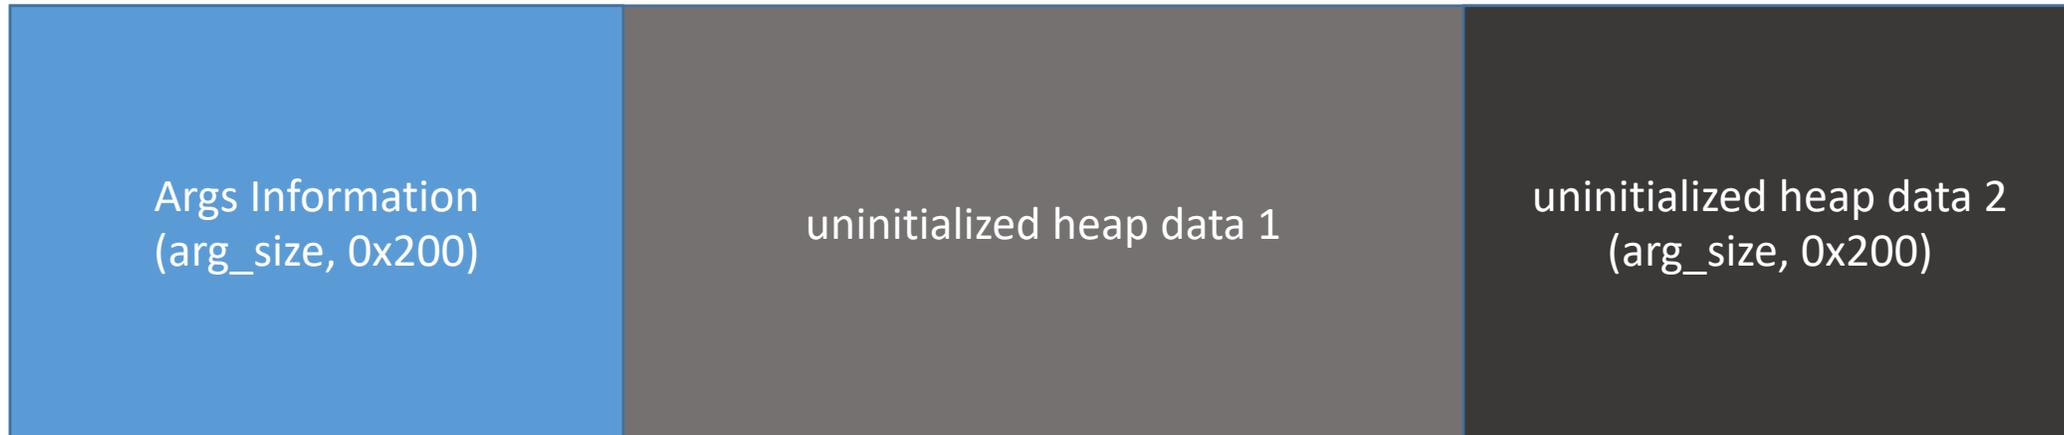
if (argc_yes) {
    /* Put processes argc as the first word in the copyout buffer */
    suword(where, argc);
    error = copyout(data, (where + sizeof(int)), size);
    size += sizeof(int);
} else {
    error = copyout(data, where, size);          --- (d)
}
}
```



page start

data

copy_end

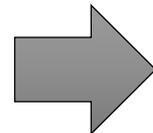


leaked!!!



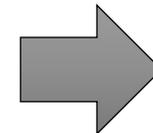
Exploit CVE-2018-4413 to leak ipc_port object address:

```
MACH_MSG_OOL_PORTS_DESCRIPTOR  
  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8
```



Destroy the ports memory:

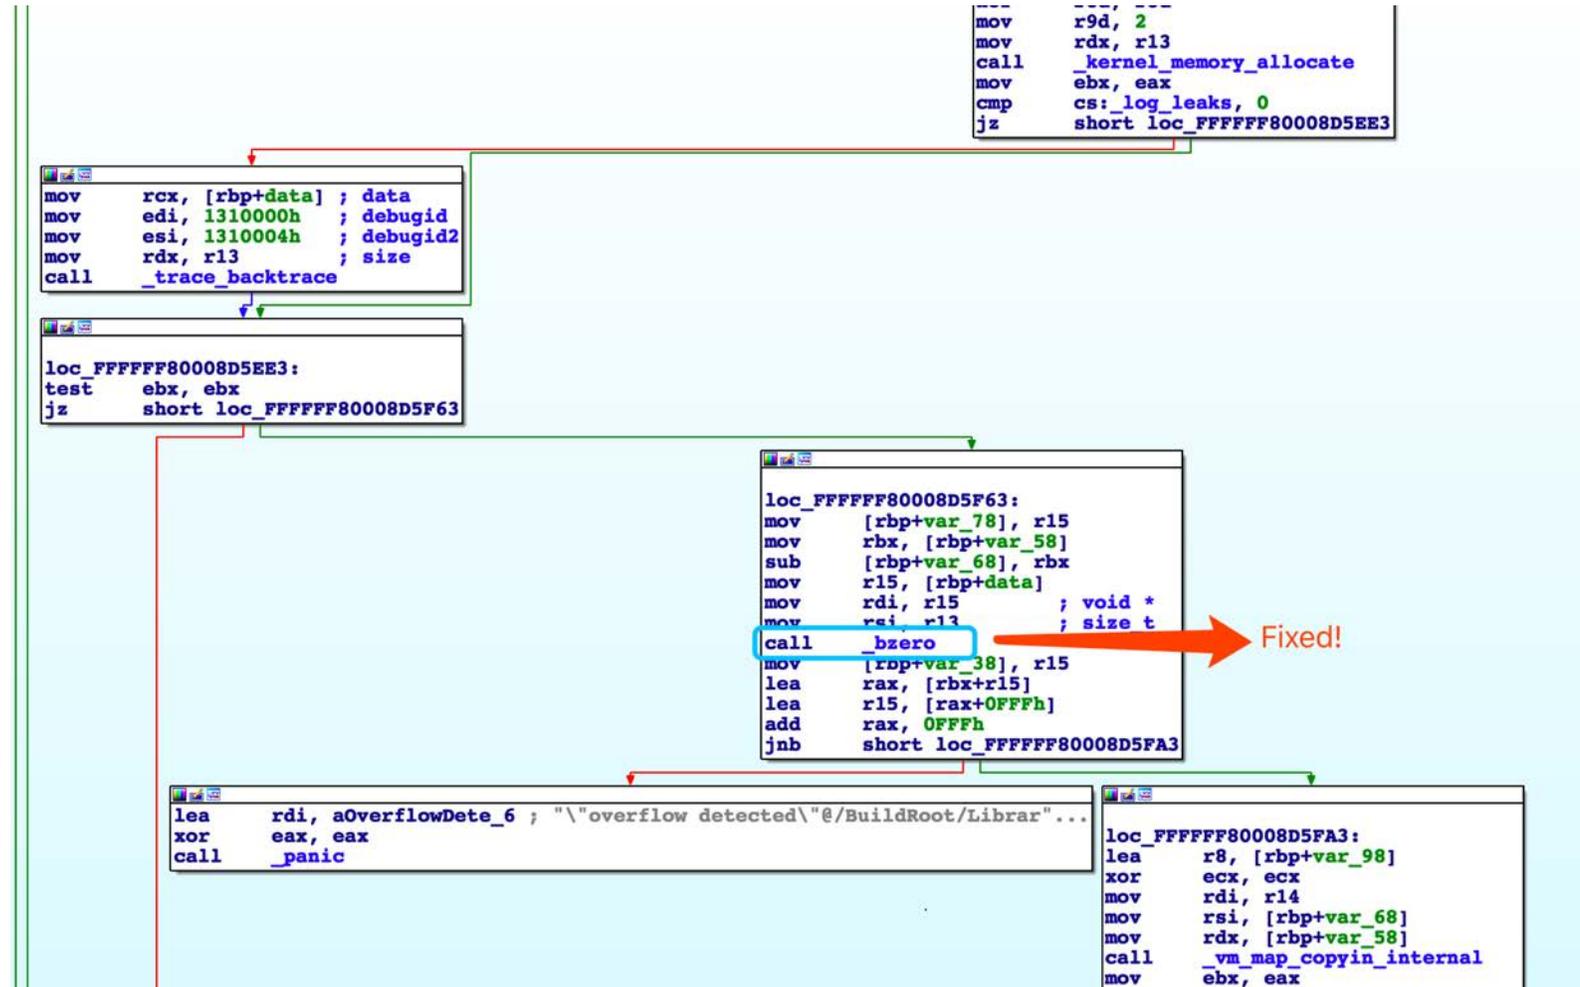
```
mach_port_destroy(mach_task_self(), q);
```



```
Trigger the vulnerability to leak  
the ports memory:  
  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8
```

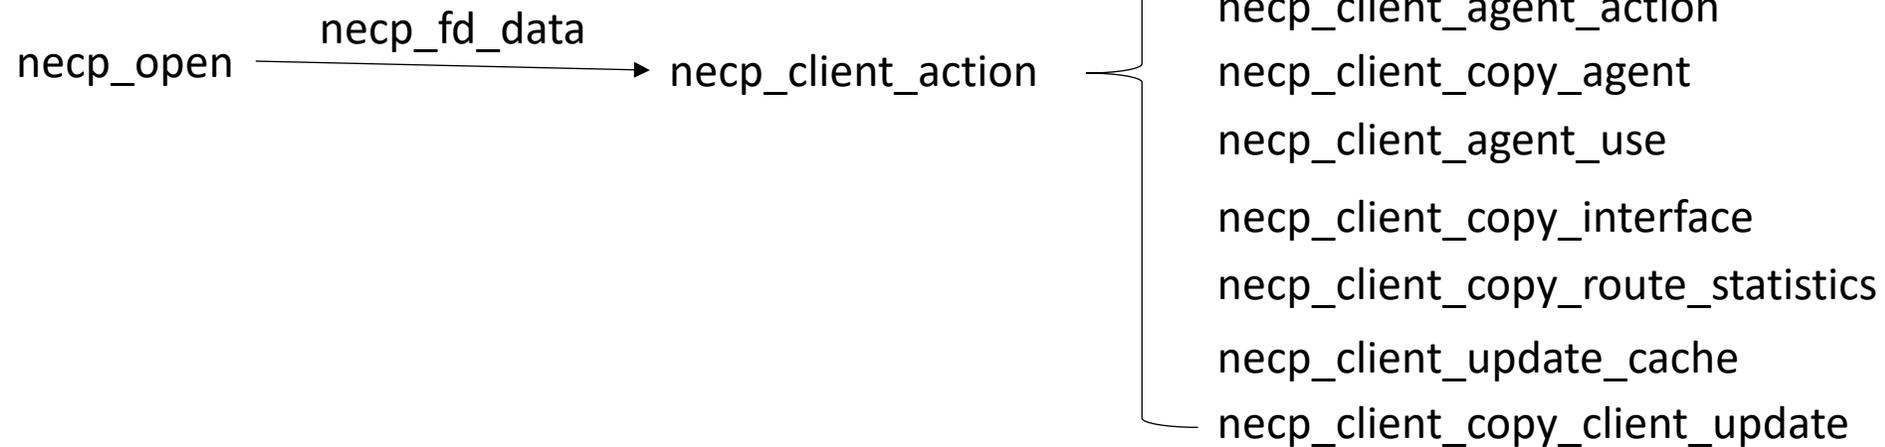
CVE-2018-4413

Apple fixed it by calling bzero.



CVE-2018-4425

NECP Attack Surface 1



```
int
necp_open(struct proc *p, struct necp_open_args *uap, int *retval)
{
#pragma unused(retval)
    int error = 0;
    struct necp_fd_data *fd_data = NULL;
    struct fileproc *fp = NULL;
    int fd = -1;

    ...

    fp->f_fglob->fg_data = fd_data;

    ...
}

struct necp_fd_data {
    +0x00 u_int8_t necp_fd_type;
    +0x08 LIST_ENTRY(necp_fd_data) chain;
    +0x18 struct _necp_client_tree clients;
    +0x20 TAILQ_HEAD(_necp_client_update_list, necp_client_update) update_list;
    +0x30 int update_count;
    +0x34 int flags;
    +0x38 int proc_pid;
    +0x40 decl_lck_mtx_data(, fd_lock);
    +0x50 struct selinfo si;
};
```

CVE-2018-4425

NECP Attack Surface 1

necp_open assigns necp_fd_data to fg_data:

- user-mode syscall gets returned fd handle
- fd is an index to kernel fp object
- fp object contains necp_fd_data object as fg_data



```
int
necp_client_action(struct proc *p, struct necp_client_action_args *uap, int *retval)
{
#pragma unused(p)
    int error = 0;
    int return_value = 0;
    struct necp_fd_data *fd_data = NULL;
    error = necp_find_fd_data(uap->necp_fd, &fd_data); ---(a)
    if (error != 0) {
        NECPLOG(LOG_ERR, "necp_client_action find fd error (%d)", error);
        return (error);
    }

    u_int32_t action = uap->action;
    switch (action) {
        ...
    }
}
```

CVE-2018-4425

NECP Attack Surface 1

necp_client_action operates on fg_data:

- at (a), call necp_find_fd_data to find necp_fd_data with given handle
- dispatch methods operates on necp_fd_data

```
static int
necp_find_fd_data(int fd, struct necp_fd_data **fd_data)
{
    proc_t p = current_proc();
    struct fileproc *fp = NULL;
    int error = 0;

    proc_fdlock_spin(p);
    if ((error = fp_lookup(p, fd, &fp, 1)) != 0) {
        goto done;
    }
    if (fp->f_fglob->fg_ops->fo_type != DTYPE_NETPOLICY) { ---(b)
        fp_drop(p, fd, fp, 1);
        error = ENODEV;
        goto done;
    }
    *fd_data = (struct necp_fd_data *)fp->f_fglob->fg_data;

done:
    proc_fdunlock(p);
    return (error);
}
```

CVE-2018-4425

NECP Attack Surface 1

necp_find_fd_data finds fd_data:

- call fp_lookup to get fp of given fd
- at (b), verify if the fp is of type necp_fd_data by checking fo_type

CVE-2018-4425

NECP Attack Surface 1

Normal Process:

- `necp_open` creates `necp_fd_data` object in kernel and returns handle to user mode
- `necp_client_action` finds the `necp_fd_data` by given handle, it internally checks if corresponding `fo_type` equals `DTYPE_NETPOLICY`
- dispatch methods of `necp_client_action` operates on found `necp_fd_data`

CVE-2018-4425

NECP Attack Surface 2

necp_session_open $\xrightarrow{\text{necp_session}}$ necp_session_action

- necp_session_add_policy
- necp_session_get_policy
- necp_session_delete_policy
- necp_session_apply_all
- necp_session_list_all
- necp_session_delete_all
- necp_session_set_session_priority
- necp_session_lock_to_process
- necp_session_register_service
- necp_session_unregister_service
- necp_session_dump_all

```
int
necp_session_open(struct proc *p, struct necp_session_open_args *uap, int *retval)
{
#pragma unused(uap)
    int error = 0;
    struct necp_session *session = NULL;
    struct fileproc *fp = NULL;
    int fd = -1;

    ...

    fp->f_fglob->fg_data = session;

    ...
}

struct necp_session {
+0x00    u_int8_t    necp_fd_type;
+0x04    u_int32_t   control_unit;
+0x08    u_int32_t   session_priority; // Descriptive priority rating
+0x0c    u_int32_t   session_order;

+0x10    decl_lck_mtx_data(, lock);

+0x20    bool    proc_locked; // Messages must come from proc_uuid
+0x21    uuid_t  proc_uuid;
+0x34    int    proc_pid;

+0x38    bool    dirty;
+0x40    LIST_HEAD(_policies, necp_session_policy) policies;

+0x50    LIST_HEAD(_services, necp_service_registration) services;

+0x60    TAILQ_ENTRY(necp_session) chain;
};
```

CVE-2018-4425

NECP Attack Surface 2

necp_session open assigns necp_session to fg_data:

- user-mode syscall gets returned fd handle
- fd is an index to kernel fp object
- fp object contains necp_session object as fg_data

CVE-2018-4425

NECP Attack Surface 2

```
int
necp_session_action(struct proc *p, struct necp_session_action_args *uap, int *retval)
{
#pragma unused(p)
    int error = 0;
    int return_value = 0;
    struct necp_session *session = NULL;
    error = necp_session_find_from_fd(uap->necp_fd, &session); ---(aa)
    if (error != 0) {
        NECPLOG(LOG_ERR, "necp_session_action find fd error (%d)", error);
        return (error);
    }

    NECP_SESSION_LOCK(session);
    ...
}
```

necp_session_action operates on fg_data:

- at (aa), call necp_session_find_from_fd to find necp_session with given handle
- dispatch methods operates on necp_session object

```
static int
necp_session_find_from_fd(int fd, struct necp_session **session)
{
    proc_t p = current_proc();
    struct fileproc *fp = NULL;
    int error = 0;

    proc_fdlock_spin(p);
    if ((error = fp_lookup(p, fd, &fp, 1)) != 0) {
        goto done;
    }
    if (fp->f_fglob->fg_ops->fo_type != DTYPE_NETPOLICY) { ---(bb)
        fp_drop(p, fd, fp, 1);
        error = ENODEV;
        goto done;
    }
    *session = (struct necp_session *)fp->f_fglob->fg_data;

done:
    proc_fdunlock(p);
    return (error);
}
```

CVE-2018-4425

NECP Attack Surface 2

necp_session_find_from_fd finds fd_data:

- call fp_lookup to get fp of given fd
- at (bb), verify if the fp is of type necp_session by checking fo_type



CVE-2018-4425

NECP Attack Surface 2

Normal Process:

- `necp_session_open` creates `necp_session` object in kernel and returns handle to user mode
- `necp_session_action` finds the `necp_session` by given handle, it internally checks if corresponding `fo_type` equals `DTYPE_NETPOLICY`
- dispatch methods of `necp_session_action` operates on found `necp_session`

CVE-2018-4425

Type Confusion

What we learn so far:

Attack surface 1: if fp->...->fo_type == DTYPE_NETPOLICY , fp is of type **necp_fd_data**

Attack surface 2: if fp->...->fo_type == DTYPE_NETPOLICY , fp is of type **necp_session**

necp_fd_data is totally different from **necp_session!!!**

我和小伙伴们都惊呆了!





CVE-2018-4425

Exploit : arbitrary address free

Method:

1. create `necp_fd_data` object and call `necp_session_action` to operate on it
2. create `necp_session` object and call `necp_client_action` to operate on it



CVE-2018-4425

Exploit : arbitrary address free

Step 1 call `necp_open` to create `necp_fd_data` object:

- `fd_data->update_list` is initialized by `TAILQ_INIT`
 - +20: 0
 - +28: `update_list` address

```
struct necp_fd_data {
    +0x00 u_int8_t necp_fd_type;
    +0x08 LIST_ENTRY(necp_fd_data) chain;
    +0x18 struct _necp_client_tree clients;
    +0x20 TAILQ_HEAD(_necp_client_update_list, necp_client_update) update_list;
    +0x30 int update_count;
    +0x34 int flags;
    +0x38 int proc_pid;
    +0x40 decl_lck_mtx_data(, fd_lock);
    +0x50 struct selinfo si;
};
```

```
int
necp_open(struct proc *p, struct necp_open_args *uap, int *retval)
{
    #pragma unused(retval)
    int error = 0;
    struct necp_fd_data *fd_data = NULL;
    struct fileproc *fp = NULL;
    int fd = -1;

    if (uap->flags & NECP_OPEN_FLAG_OBSERVER) {
        if (necp_skywalk_priv_check_cred(p, kauth_cred_get()) != 0 &&
            priv_check_cred(kauth_cred_get(), PRIV_NET_PRIVILEGED_NETWORK)
            NECPLOG0(LOG_ERR, "Client does not hold necessary entitlement")
            error = EACCES;
        goto done;
    }

    error = falloc(p, &fp, &fd, vfs_context_current());
    if (error != 0) {
        goto done;
    }

    if ((fd_data = zalloc(necp_client_fd_zone)) == NULL) {
        error = ENOMEM;
        goto done;
    }

    memset(fd_data, 0, sizeof(*fd_data));

    fd_data->necp_fd_type = necp_fd_type_client;
    fd_data->flags = uap->flags;
    RB_INIT(&fd_data->clients);
    TAILQ_INIT(&fd_data->update_list);
    lck_mtx_init(&fd_data->fd_lock, necp_fd_mtx_grp, necp_fd_mtx_attr);
    klist_init(&fd_data->si.si_note);
    fd_data->proc_pid = proc_pid(p);

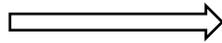
    fp->f_fglob->fg_flag = FREAD;
    fp->f_fglob->fg_ops = &necp_fd_ops;
    fp->f_fglob->fg_data = fd_data;
```



CVE-2018-4425

Exploit : arbitrary address free

necp_open



+0x20: 0

+0x28: update_list address

CVE-2018-4425

Exploit : arbitrary address free

Step 2 call `necp_session_action` on the object

at location (b), if `session->proc_locked` is false(0), `session->proc_uid` and `session->proc_pid` will be updated.

```
int
necp_session_action(struct proc *p, struct necp_session_action_args *uap, int *retval)
{
#pragma unused(p)
    int error = 0;
    int return_value = 0;
    struct necp_session *session = NULL;
    error = necp_session_find_from_fd(uap->necp_fd, &session);
    if (error != 0) {
        NECPLOG(LOG_ERR, "necp_session_action find fd error (%d)", error);
        return (error);
    }

    NECP_SESSION_LOCK(session);

    if (session->proc_locked) {
        // Verify that the calling process is allowed to do actions
        uid_t proc_uid;
        proc_getexecutableuid(current_proc(), proc_uid, sizeof(proc_uid));
        if (uid_compare(proc_uid, session->proc_uid) != 0) {
            error = EPERM;
            goto done;
        }
    } else {
        // If not locked, update the proc_uid and proc_pid of the session
        proc_getexecutableuid(current_proc(), session->proc_uid, sizeof(session->proc_uu
        session->proc_pid = proc_pid(current_proc());    ---(b)
    }

    ...
}
```

CVE-2018-4425

Exploit : arbitrary address free

- session->proc_locked at offset 0x20 overlaps update_list which is 0 in necp_fd_data.
- session->proc_uuid at offset 0x21 is updated with macho UUID
- session->proc_pid is updated with current pid

```
struct necp_session {
    +0x00    u_int8_t    necp_fd_type;
    +0x04    u_int32_t   control_unit;
    +0x08    u_int32_t   session_priority; // Descriptive priority rating
    +0x0c    u_int32_t   session_order;

    +0x10    decl_lck_mtx_data(, lock);

    +0x20    bool    proc_locked; // Messages must come from proc_uuid
    +0x21    uuid_t  proc_uuid;
    +0x34    int    proc_pid;

    +0x38    bool    dirty;
    +0x40    LIST_HEAD(_policies, necp_session_policy) policies;

    +0x50    LIST_HEAD(_services, necp_service_registration) services;

    +0x60    TAILQ_ENTRY(necp_session) chain;
};
```

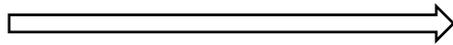


CVE-2018-4425

Exploit : arbitrary address free

+0x20: 0
+0x28: update_list address

necp_session_action



+0x20: 0
+0x21: UUID, low 7Bytes
+0x28: UUID, high 9Bytes
+0x34: pid

CVE-2018-4425

Exploit : arbitrary address free

Step 3 call `necp_client_action` on the object

- we use action 15(`necp_client_copy_client_update`)
- at location (f), `client_update` is freed
- `client_update` is the first element of `update_list` which is UUID now

```
NECP_FD_LOCK(fd_data);
struct necp_client_update *client_update = TAILQ_FIRST(&fd_data->update_list);
if (client_update != NULL) {
    TAILQ_REMOVE(&fd_data->update_list, client_update, chain); ---(c)
    VERIFY(fd_data->update_count > 0);
    fd_data->update_count--;
}
NECP_FD_UNLOCK(fd_data);

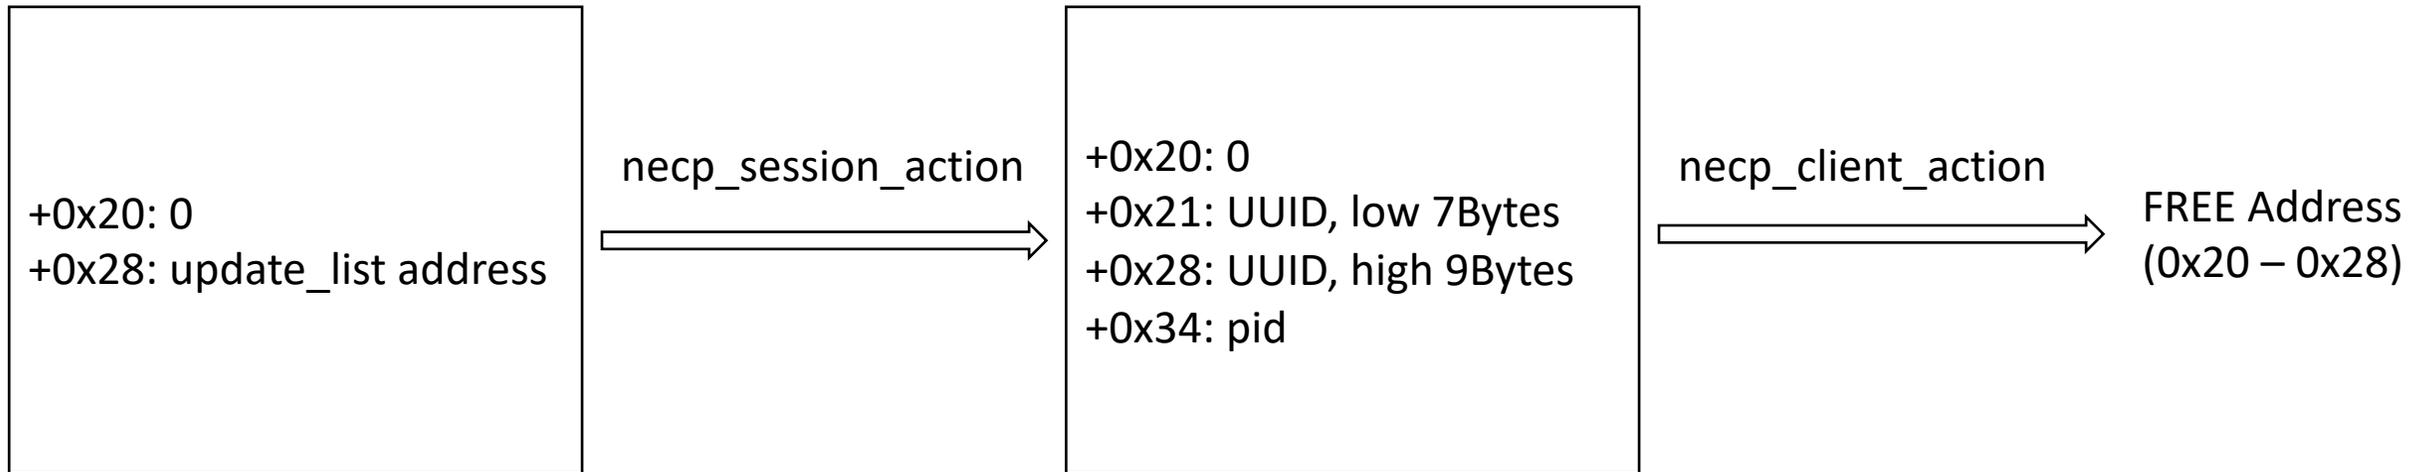
if (client_update != NULL) {
    error = copyout(client_update->client_id, uap->client_id, sizeof(uuid_t)); ---(d)
    if (error) {
        NECPLOG(LOG_ERR, "Copy client update copyout client id error (%d)", error);
    } else {
        if (uap->buffer_size < client_update->update_length) {
            NECPLOG(LOG_ERR, "Buffer size cannot hold update (%zu < %zu)", uap->buffer_
            error = EINVAL;
        } else {
            error = copyout(&client_update->update, uap->buffer, client_update->update_
            if (error) {
                NECPLOG(LOG_ERR, "Copy client update copyout error (%d)", error);
            } else {
                *retval = client_update->update_length;
            }
        }
    }
}

FREE(client_update, M_NECP); ---(f)
client_update = NULL;
} else {
    error = ENOENT;
}

return (error);
}
```

CVE-2018-4425

Exploit : arbitrary address free



For Example, we set MachO UUID(16 bytes) as 41414141414141414141414141414141, here we get 0x414141414141414100 freed. We can control high 7 bytes of the address to be freed.



CVE-2018-4425

Apple Fix

Add sub type check:

necp_session has sub type 1

necp_fd_data has sub type 2

```

0      public _necp_session_action
0      _necp_session_action proc near
0      push    rbp
1      mov     rbp, rsp
4      push    r15
5      push    r14
6      push    r13
7      push    r12
8      push    rbx
9      sub     rsp, 48h
A      mov     [rbp-50h], rdx
B      mov     r13, rsi
C      lea    rax, ___stack_chk_guard
D      mov     rax, [rax]
E      mov     [rbp-30h], rax
F      movsxd  rbx, dword ptr [r13+0]
1      call   _current_proc
2      mov     r15, rax
3      lea    r12, [r15+0C0h]
4      mov     rdi, r12
5      call   _lck_mtx_lock_spin_always
6      mov     r14d, 9
7      test   rbx, rbx
8      js     loc_FFFFFFFF80006BD34B
9      mov     rax, [r15+0E8h]
A      test   rax, rax
B      jz     loc_FFFFFFFF80006BD34B
C      cmp   [rax+48h], ebx
D      jle   loc_FFFFFFFF80006BD34B
E      mov     rcx, [rax]
F      mov     rdx, [rcx+rbx*8]
1      test   rdx, rdx
2      jz     loc_FFFFFFFF80006BD34B
3      mov     rax, [rax+30h]
4      test   byte ptr [rax+rbx], 4
5      jnz   loc_FFFFFFFF80006BD34B
6      inc   dword ptr [rdx+4]
7      mov     rax, [rdx+8]
8      mov     rcx, [rax+28h]
9      cmp   dword ptr [rcx], 9 ; DTYPE_NETPOLICY
A      jnz   loc_FFFFFFFF80006BD336
B      mov     rbx, [rax+38h] ; fg_data
C      mov     r14d, 16h
D      cmp   byte ptr [rbx], 1 ; sub_type check Fixed!
E      jnz   loc_FFFFFFFF80006BD34B
F      mov     rdi, r12
1      call   _lck_mtx_unlock
2      lea   r15, [rbx+18h]

```

BD235: _necp_session_action+A5 (Synchronized with Hex View-1)

```

)      public _necp_client_action
)      _necp_client_action proc near
)      push    rbp
1      mov     rbp, rsp
2      push    r15
3      push    r14
4      push    r13
5      push    r12
6      push    rbx
7      sub     rsp, 448h
8      mov     [rbp-428h], rdx
9      mov     r13, rdi
A      lea    rax, ___stack_chk_guard
B      mov     rax, [rax]
C      mov     [rbp-30h], rax
D      mov     [rbp-418h], rsi
E      movsxd  r14, dword ptr [rsi]
F      call   _current_proc
1      mov     r12, rax
2      lea    r15, [r12+0C0h]
3      mov     rdi, r15
4      call   _lck_mtx_lock_spin_always
5      mov     ebx, 9
6      test   r14, r14
7      js     loc_FFFFFFFF80006DE5E3
8      mov     rax, [r12+0E8h]
9      test   rax, rax
A      jz     loc_FFFFFFFF80006DE5E3
B      cmp   [rax+48h], r14d
C      jle   loc_FFFFFFFF80006DE5E3
D      mov     rcx, [rax]
E      mov     rdx, [rcx+r14*8]
F      test   rdx, rdx
1      jz     loc_FFFFFFFF80006DE5E3
2      mov     rax, [rax+30h]
3      test   byte ptr [rax+r14], 4
4      jnz   loc_FFFFFFFF80006DE5E3
5      inc   dword ptr [rdx+4]
6      mov     rax, [rdx+8]
7      mov     rcx, [rax+28h]
8      cmp   dword ptr [rcx], 9 ; DTYPE_NETPOLICY
9      jnz   loc_FFFFFFFF80006DE5CE
A      mov     r14, [rax+38h] ; fg_data
B      mov     ebx, 16h
C      cmp   byte ptr [r14], 2 ; sub_type check Fixed!
D      jnz   loc_FFFFFFFF80006DE5E3
E      mov     rdi, r15
F      call   _lck_mtx_unlock
1      mov     r12, [rbp-418h]

```

DE4E3: _necp_client_action+B3 (Synchronized with Hex View-1)

macOS <= 10.14 Root

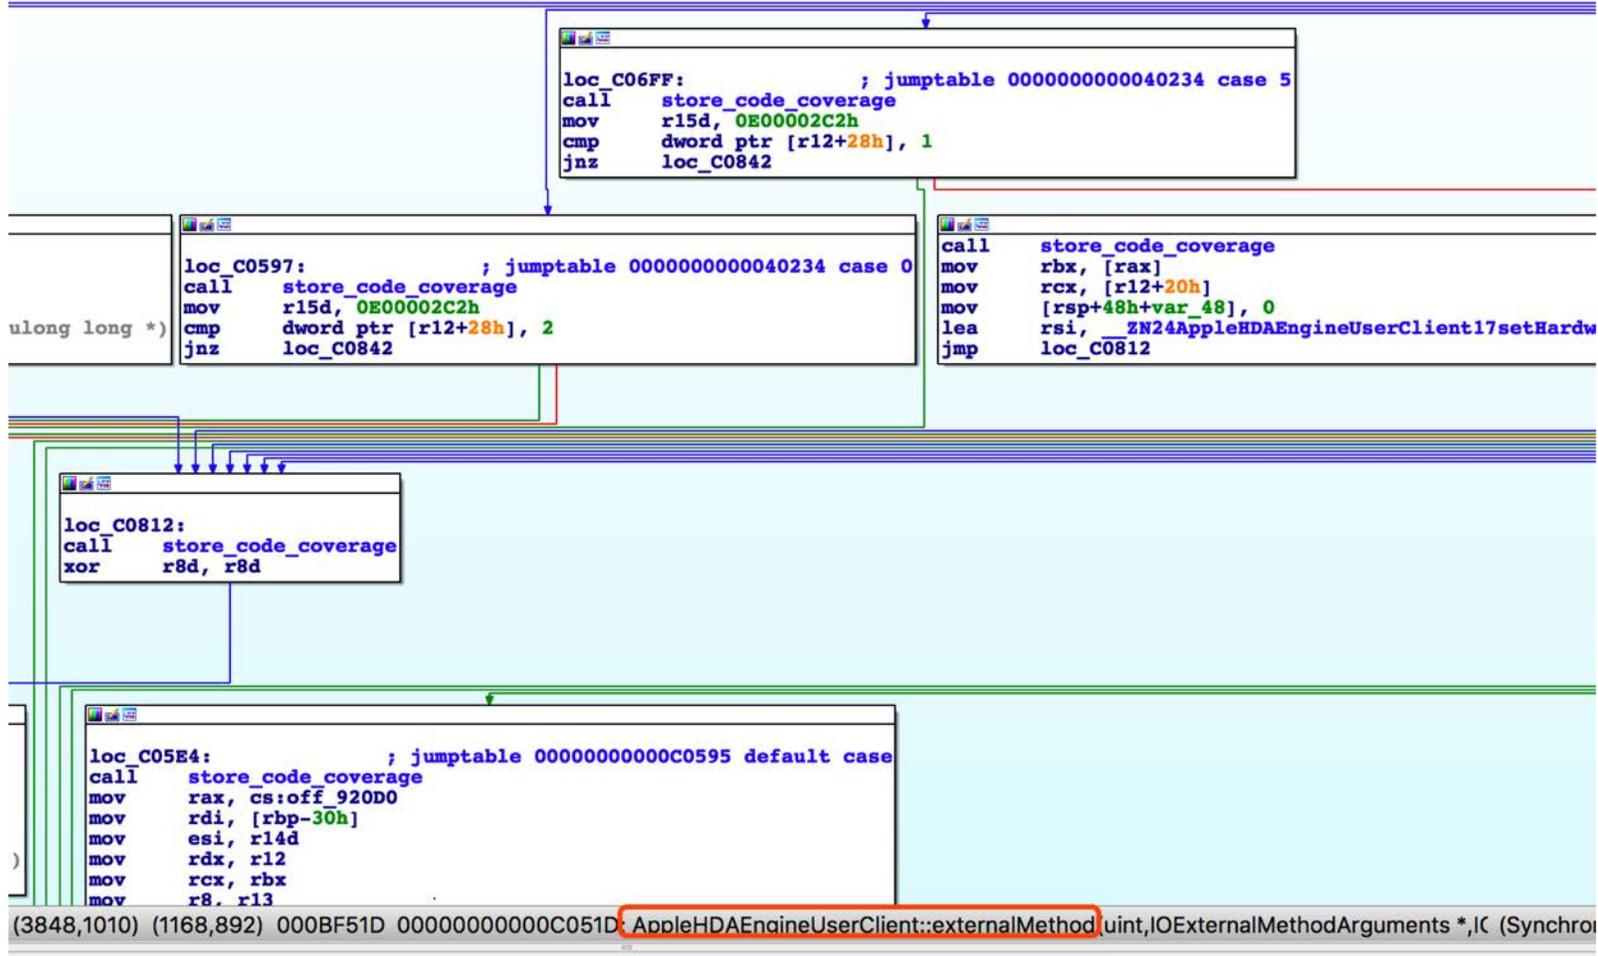
- Root = CVE-2018-4413 + CVE-2018-4425 + mach-portal
- mach_portal: all details <https://bugs.chromium.org/p/project-zero/issues/detail?id=1417>
- Demo(10.13.6)

Future Plan of Our Fuzzing Tool

- Support kernel extension
- Support IOKit(+code coverage)
- Support Passive Fuzzing
- More and More Corpus



IOKit Code Coverage Example



More Information

- follow me on twitter: @panicall



Acknowledge

- Google Project Syzkaller¹
- Ian Beer for his exploit technique

ANY QUESTIONS?