



**black hat**<sup>®</sup>  
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

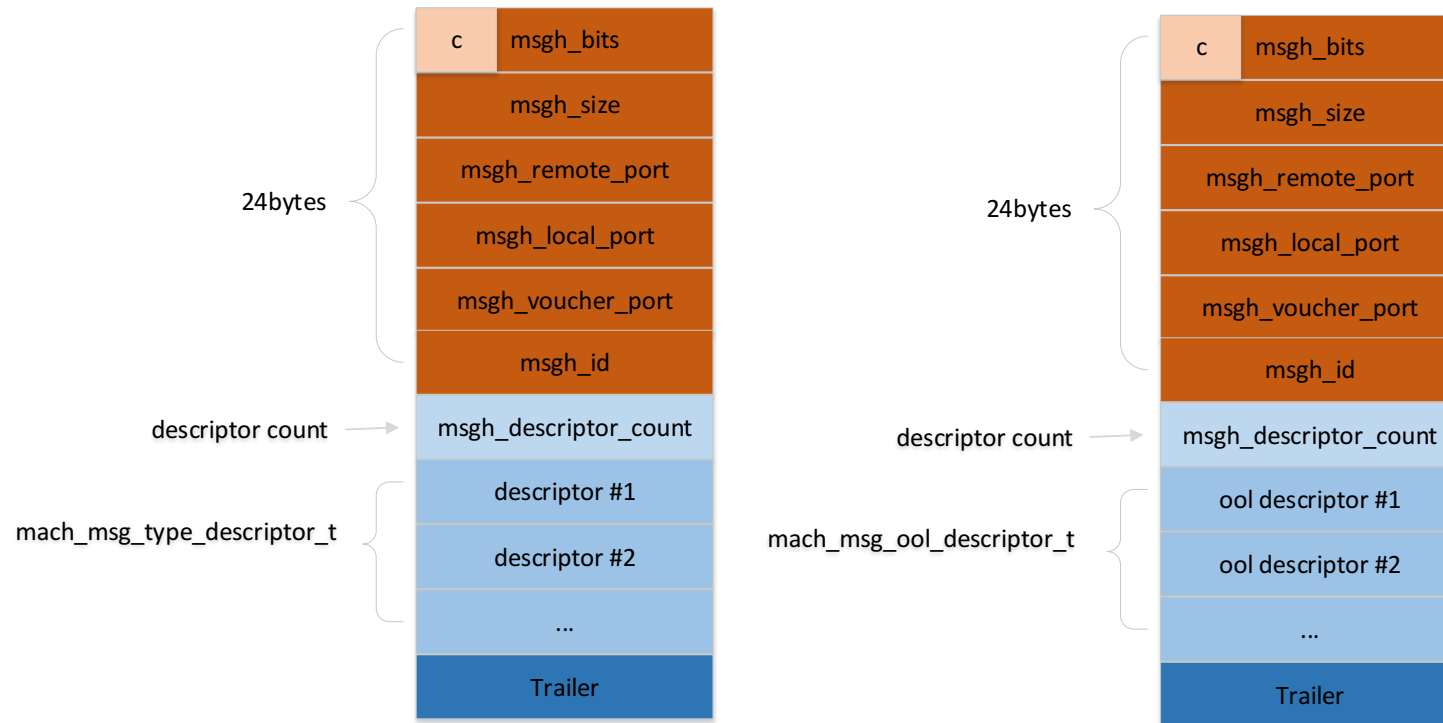
- Smart Fuzz XPC
  - XPC Internals
  - Fuzz Strategy
  - Reproduce Strategy
  - Output

- 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



- XPC Services
  - `launchctl dumpstate`

```

services = {
  xpc = {
    __init__.py = 64 -
    OsxFuzz-0.py = 0 -
    run.sh = 0 -
    xpcsConf = 0 -
    __init__.py = 65 -
    fianlService0.txt = 0 -
    fontdCor66.y = -
    generalXpc0nf.py = 0 -
    sysmondConf.py = 0 -
    xpcConf.p0 (pe) = 0 -
    XPCService0ool.py = 0 -
    __init__.py = 0 -
    launchctl_dum0state.txt = 0 -
    machMsg.py = 0 -
    OsxFuzz.p = 41779 (pe) = 0 -
    Server_osx = 141 -
    seeds = 42075 (pe) = 0 -
    com.ap1695irPlay0gent.xpc = 0 -
    com.apple0irPlayAgent.xpc = 0 -
    com.apple0irPlayXPCHelper = 269 -
    com.ap2190utoUnLock.System = 0 -
    com.apple0allHiseryPlugin = 0 -
    com.apple0oreAuthenticatio = 0 -
    com.apple0oreAuthenticatio = 0 -
    com.a50661ore (pe) tion.agen = 0 -
    com.ap1780oreServices.core = 0 -
    com.apple0ras0porterSup = 0 -
    346341966JZZ (pe) MODE = 0 -
    com.apple.wifiFirmwareLoader
    com.apple.uninstalld
    com.apple.tzlinkd
    com.apple.storedownloadd.daemon
    com.apple.rpmuxd
    com.apple.nis.yppbind
    com.apple.kextd
    com.apple.Kerberos.digest-service
    com.apple.kcproxy
    com.apple.fseventsd
    com.apple.diagnosticextensions.osx.timemachine.helper
    com.apple.diagnosticextensions.osx.spotlight.helper
    com.apple.CoreRAID
    com.apple.CoreAuthentication.daemon
    com.apple.DesktopServicesHelper.151FBB7D-869B-49E0-8EB2-2F509E9F92A6
    com.apple.DesktopServicesHelper.726D2776-BA99-4F51-B49E-06474EF7B673
    com.apple.systempreferences.cacheAssistant
    com.apple.TrustEvaluationAgent.system
    com.apple.newsyslog
    com.apple.mediaremoted
    com.apple.coreservicesd
    com.apple.automountd
    com.apple.adid
    com.apple.AmbientDisplayAgent
    com.vix.cron
    com.apple.touchbarserver
    com.apple.thermald
    com.apple.taskgated.Provider
    com.apple.storeagent.daemon
    com.apple.RemoteDesktop.PrivilegeProxy
    com.apple.MRTd
    com.apple.mbusertrampoline
    com.apple.GSSCred
    com.apple.FileCoordination
    com.apple.colorsync.displayservices
    com.apple.avbdeviced
    com.apple.audio.systemsoundsserverd
    com.apple.rpmuxd = {
      active count = 0
      path = /System/Library/L
      state = waiting

      program = /usr/libexec/r
      arguments = {
        /usr/libexec/rpmuxd
      }

      default environment = {
        PATH => /usr/bin:/bi
      }

      XPC_SERVICE_NAME =>

      domain = com.apple.xpc.l
      minimum runtime = 10
      exit timeout = 5
      runs = 0
      successive crashes = 0
      excessive crashing = 0
      last exit code = (never

      event triggers = {
      }

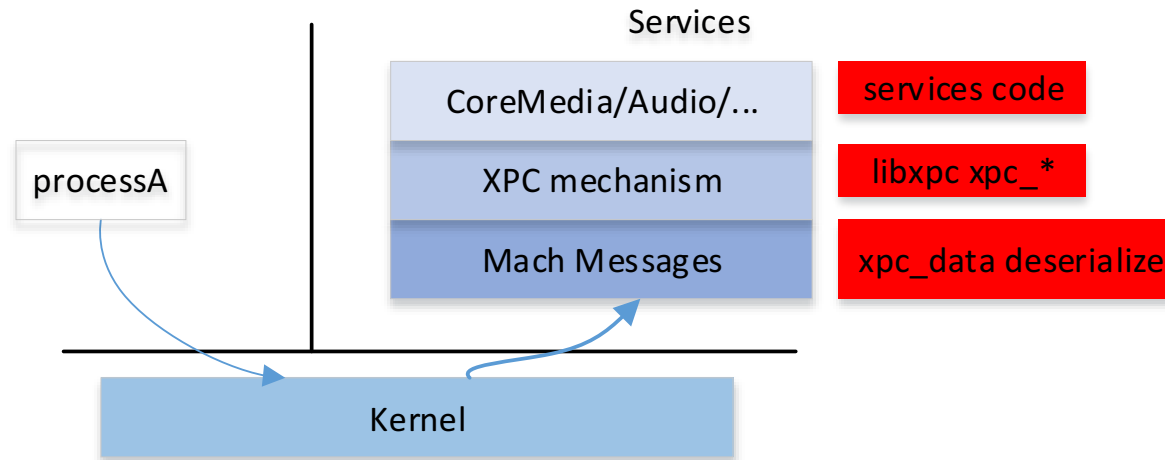
      endpoints = {
        "com.apple.rpmuxd" =
          port = 0x1be03
          active = 0
          managed = 1
          reset = 0
          hide = 0
      }
    }
  }
}

```



## • Attack Surface

- serialize/deserialize
- libxpc
- services code





- How to trigger these bugs?



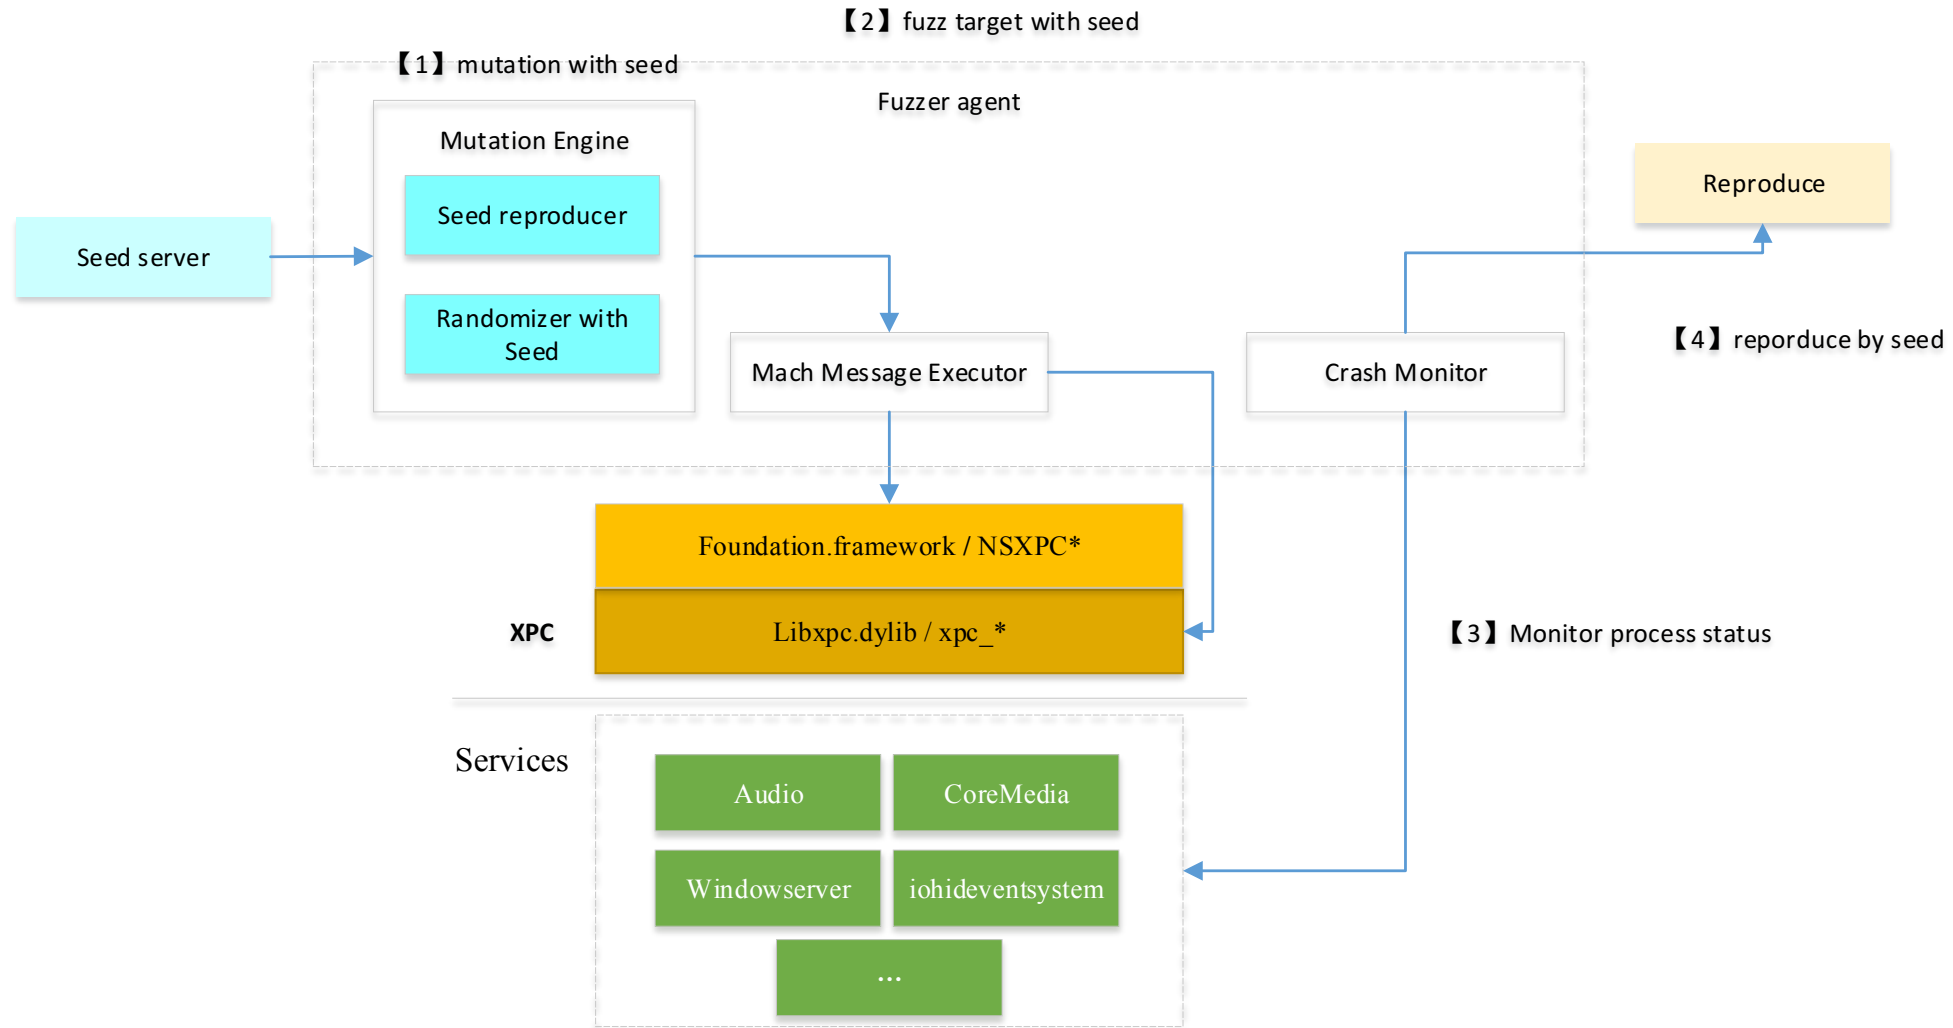
**Crafted Mach Message**





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

# XPC Fuzz Architecture





- Fuzz Controller

- ✓ Wrap the xpc interfaces by python

```
BOOST_PYTHON_MODULE(xpcconnection) {  
    PyEval_InitThreads();  
  
    class_<XpcConnection, boost::noncopyable>("XpcConnection", init<std::string>())  
    .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)  
    ;  
}
```

- ✓ Python fuzz Engine





## • 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
```



- Comparison between different Reproduce Methods

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



- Case Study - CVE-2018-4411

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



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

```
gULU LABEL_044,
case 0x28:
    v82 = &v238->bf;
    if ( gUseNewFODB == 2 )
    {
        FODDBeginTransactions(9);
        if ( LODWORD(v82->_base) )
        {
            v83 = *(&v238->_lbfssize + 1);
            a2 = *(const char *)((char *)&v238->_bf._base + 4);
            v84 = v238->_lbfssize;
        }
        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
        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
}
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
IF ( *(__BYTE *) (gHnnexHUXF11e + 12LL) )
{
```

Out of boundary

# Smart Fuzzing XNU

# Smart Fuzzing XNU

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

## What I will introduce today

1. Port Syzkaller to Support macOS XNU Fuzzing.
2. Modify XNU to add support some features.



# Fuzzer



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

# XNU

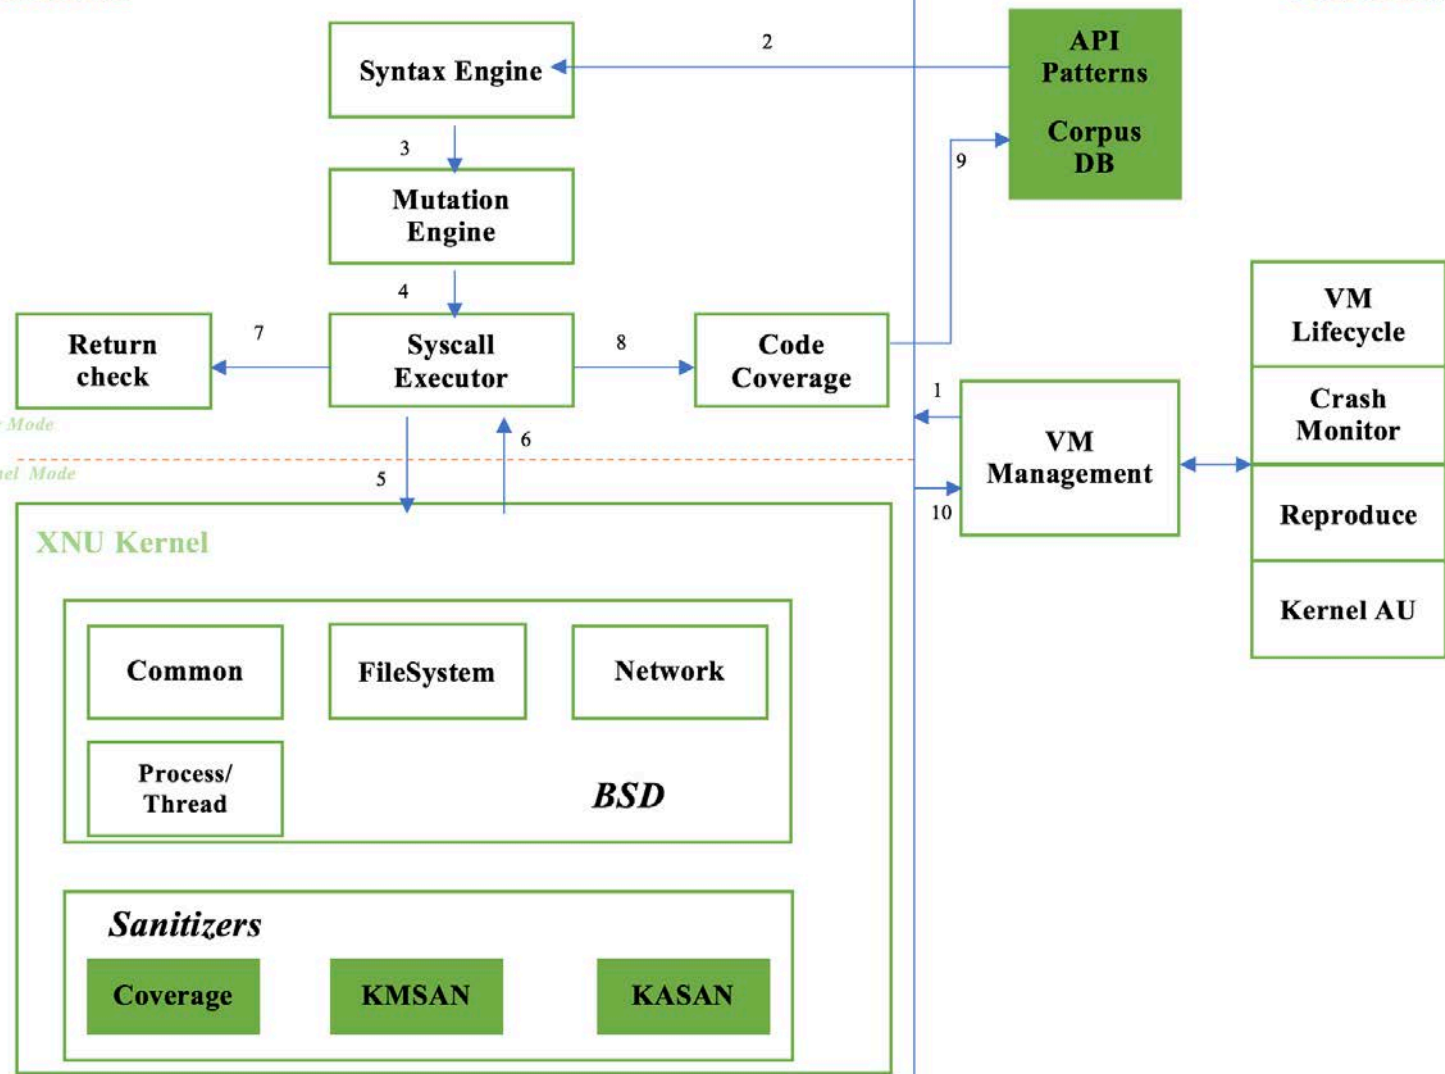


- Add Code Coverage
- Add Kernel Memory Sanitizer
- Enable Kernel Address Sanitizer



Client Side

Server Side



# Architecture

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



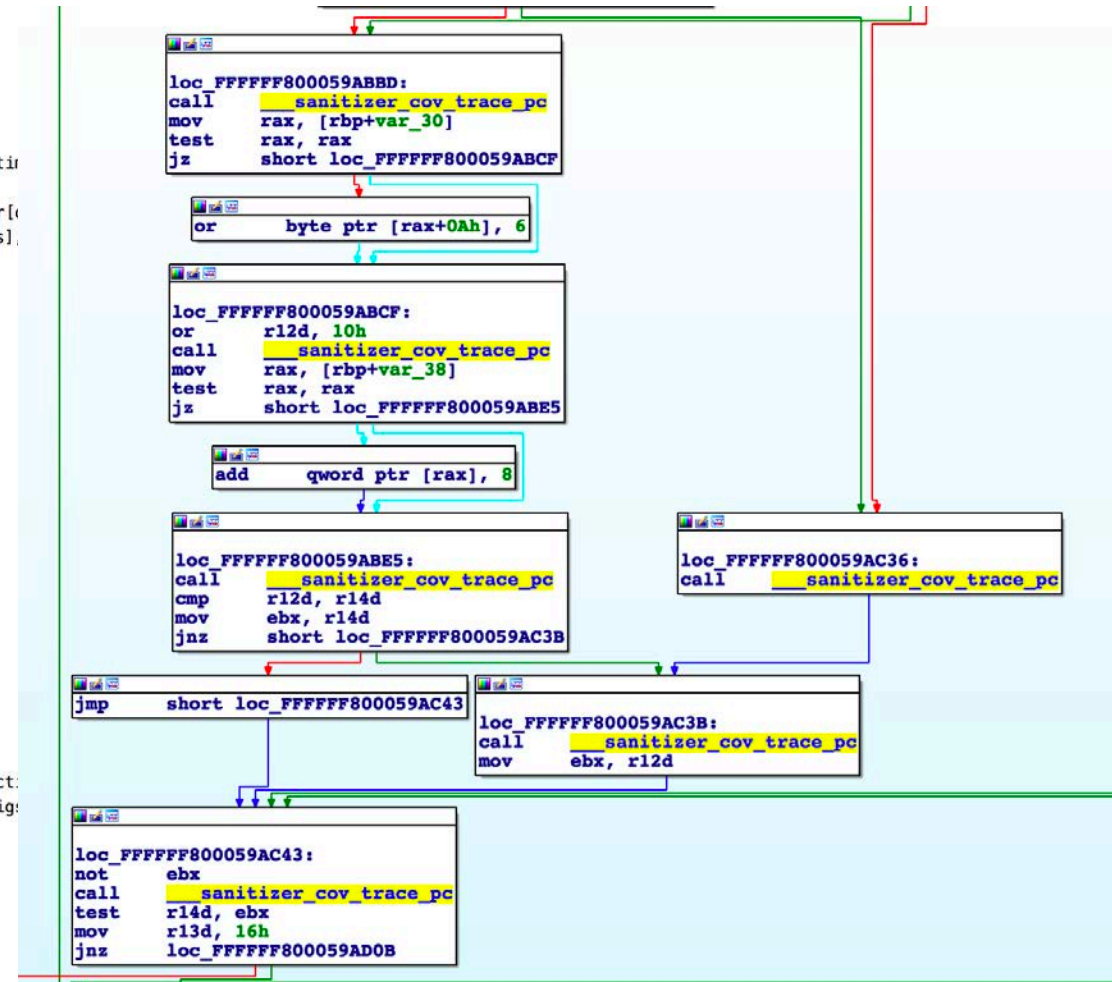
```

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_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], itimerval)
97 exit(code intptr)
98 wait4(pid pid, status ptr[out], int32, opt], options flags[wait_options], ru ptr[in], rlim ptr[in], rlim ptr[in])
99 wait4_nocancel(pid pid, status ptr[out], int32, opt], options flags[wait_options], ru ptr[in], rlim ptr[in], rlim ptr[in])
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(uid ptr[out], uid, gid 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, fhp 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], sigset)
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

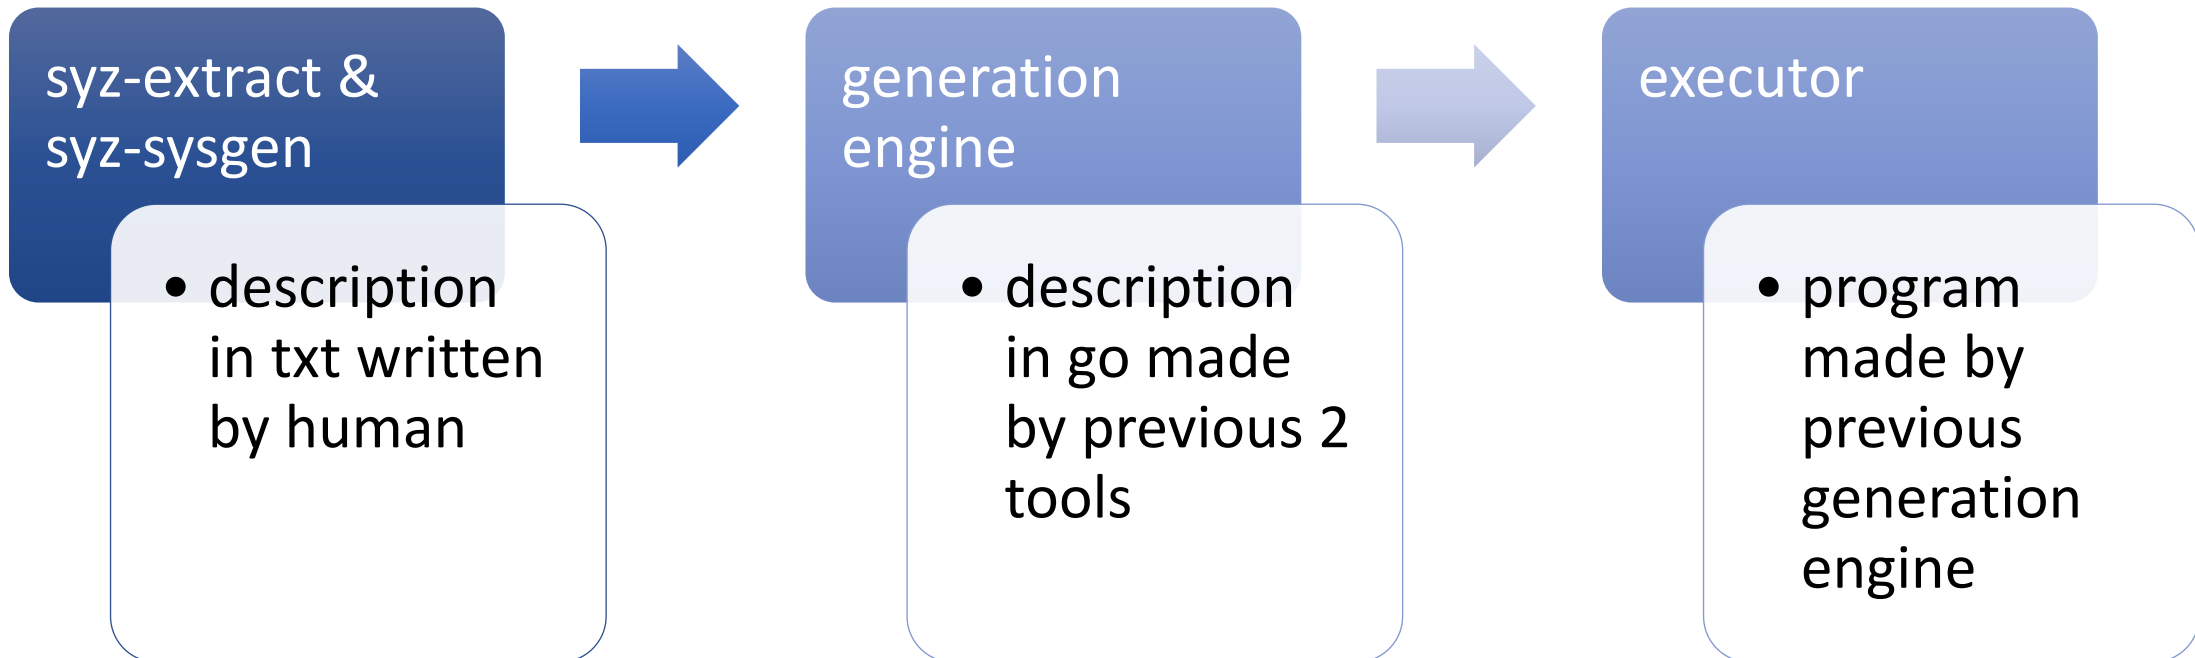
## My Efforts

- 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](https://github.com/google/syzkaller/blob/master/docs/syscall_descriptions_syntax.md)
- Refer to sample txt files in syzkaller project

# Sanitizers



# Basic Concepts 1: User Mode Sanitizers<sup>1</sup>

Name	Features	Comments
AddressSanitizer <sup>2</sup>	<ul style="list-style-type: none"><li>• Out-of-bounds accesses to heap, stack and globals</li><li>• Use-after-free</li><li>• Use-after-return</li><li>• Use-after-scope</li><li>• Double-free, invalid free</li></ul>	<ul style="list-style-type: none"><li>• compiler instrumentation module</li><li>• run-time library</li></ul>
MemorySanitizer <sup>3</sup>	<ul style="list-style-type: none"><li>• uninitialized reads</li></ul>	
SanitizerCoverage <sup>4</sup>	<ul style="list-style-type: none"><li>• get function/block/edge coverage</li></ul>	<ul style="list-style-type: none"><li>• Instrumentations</li><li>• Default callbacks provided</li></ul>
...		<ul style="list-style-type: none"><li>• ThreadSanitizer<sup>5</sup></li><li>• UndefinedBehaviorSanitizer<sup>6</sup></li><li>• DataflowSanitizer<sup>7</sup></li><li>• LeakSanitizer<sup>8</sup></li></ul>



## Basic Concepts 2: Kernel Mode Sanitizers

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

# 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





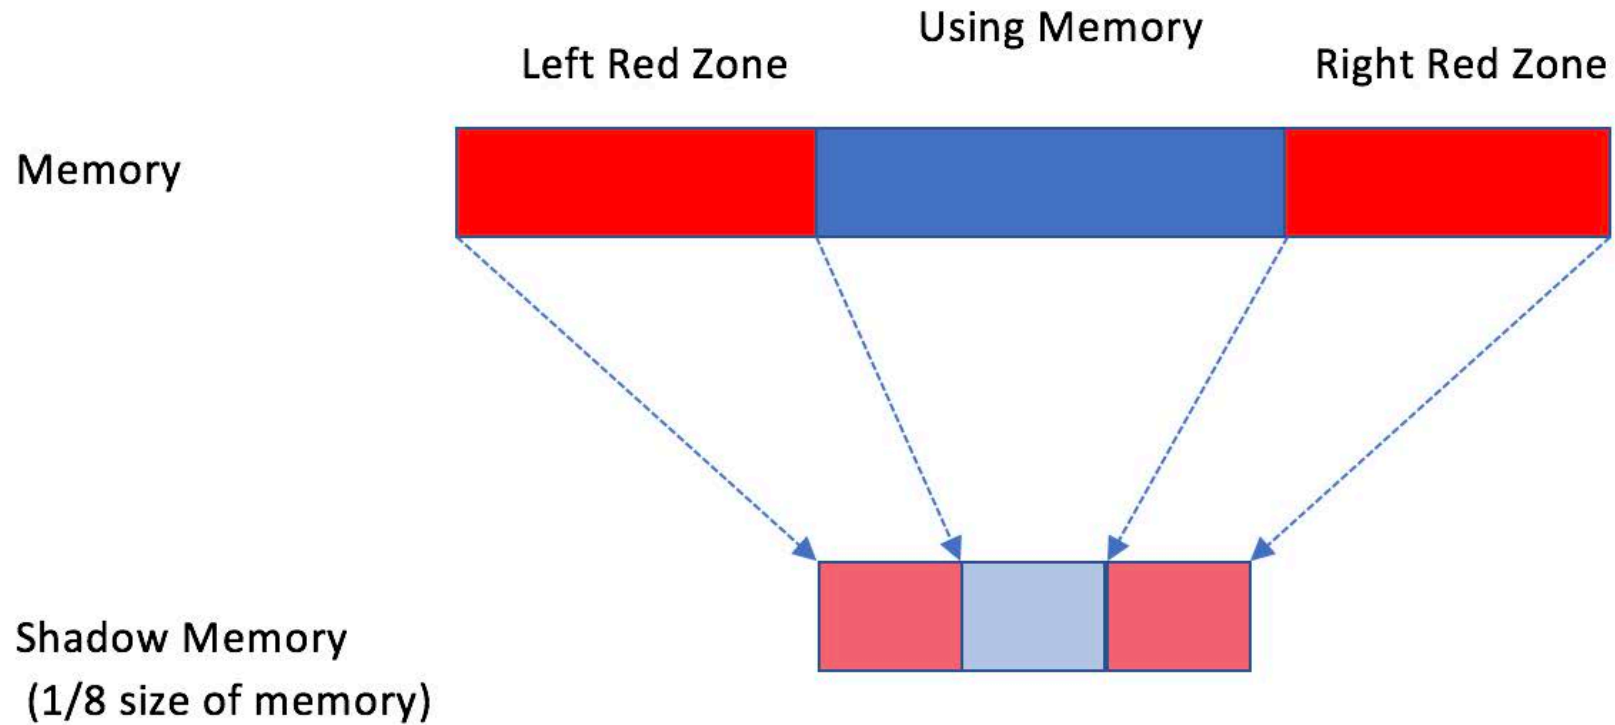
# KASAN

- 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`

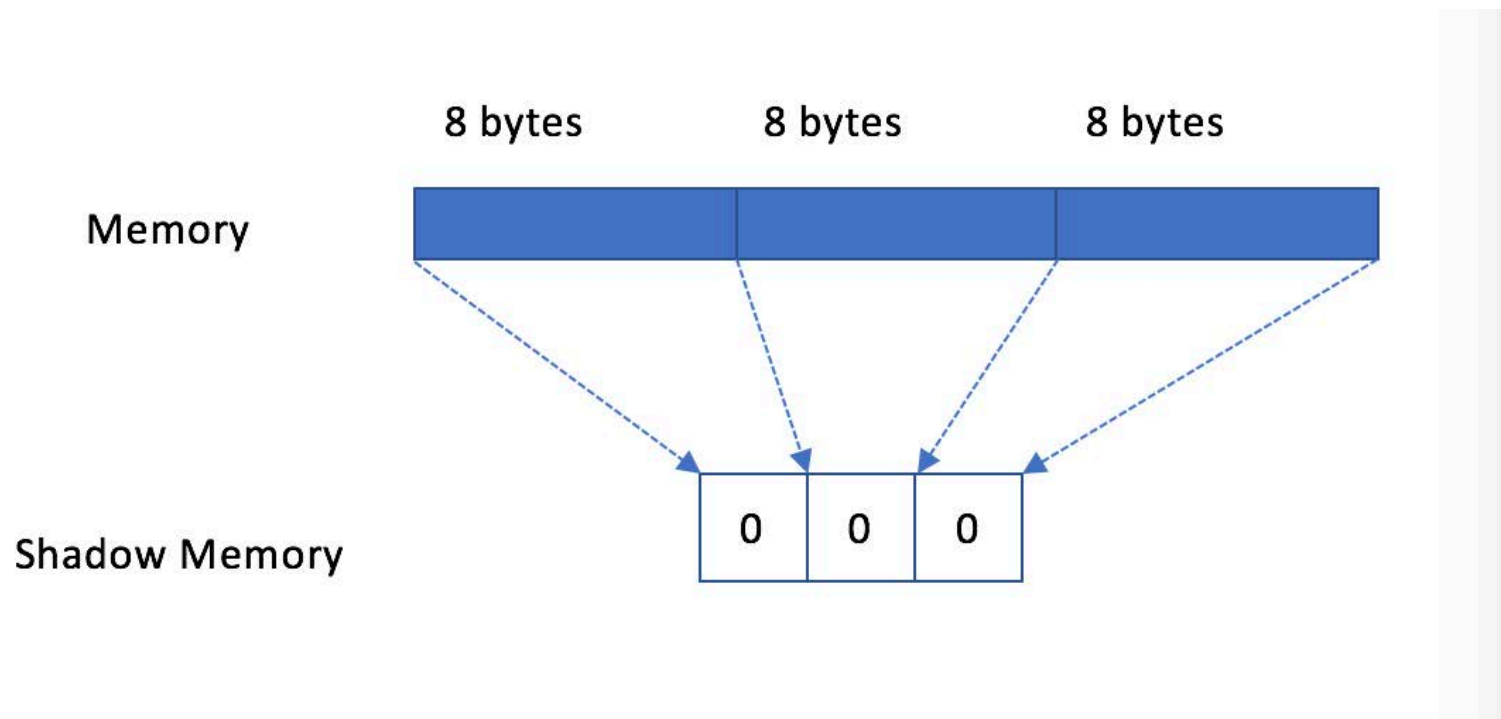
`-fsanitizer=address`

buildin calls in xnu source code



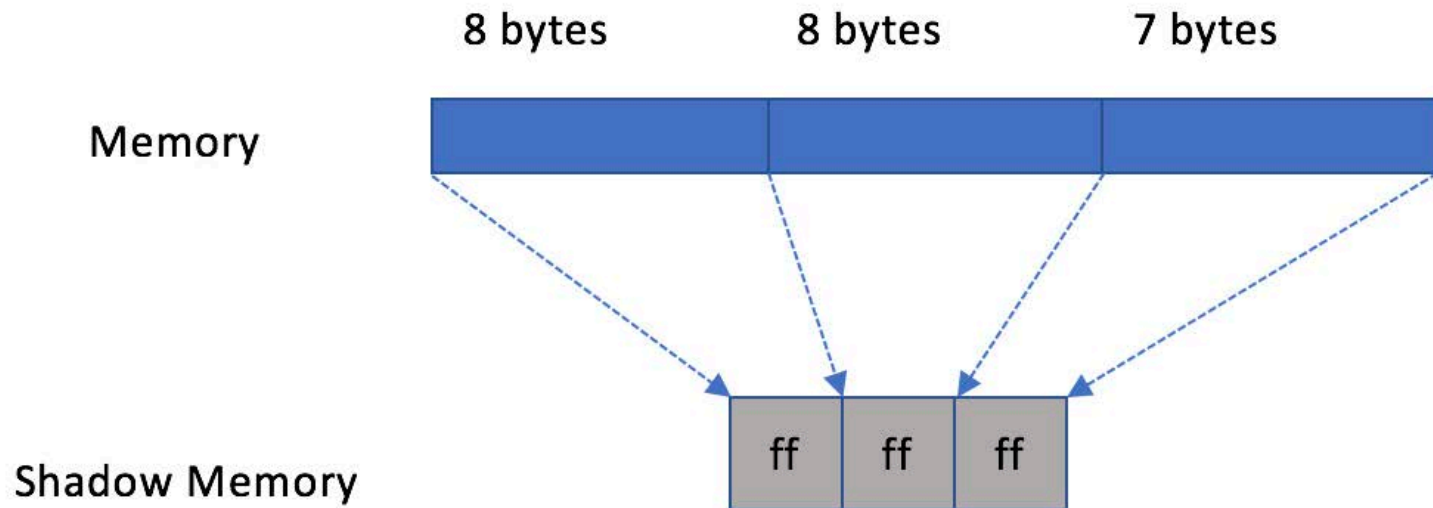
## Example: Detect UAF

- 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
- 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

```
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)
}
}
```

## 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):

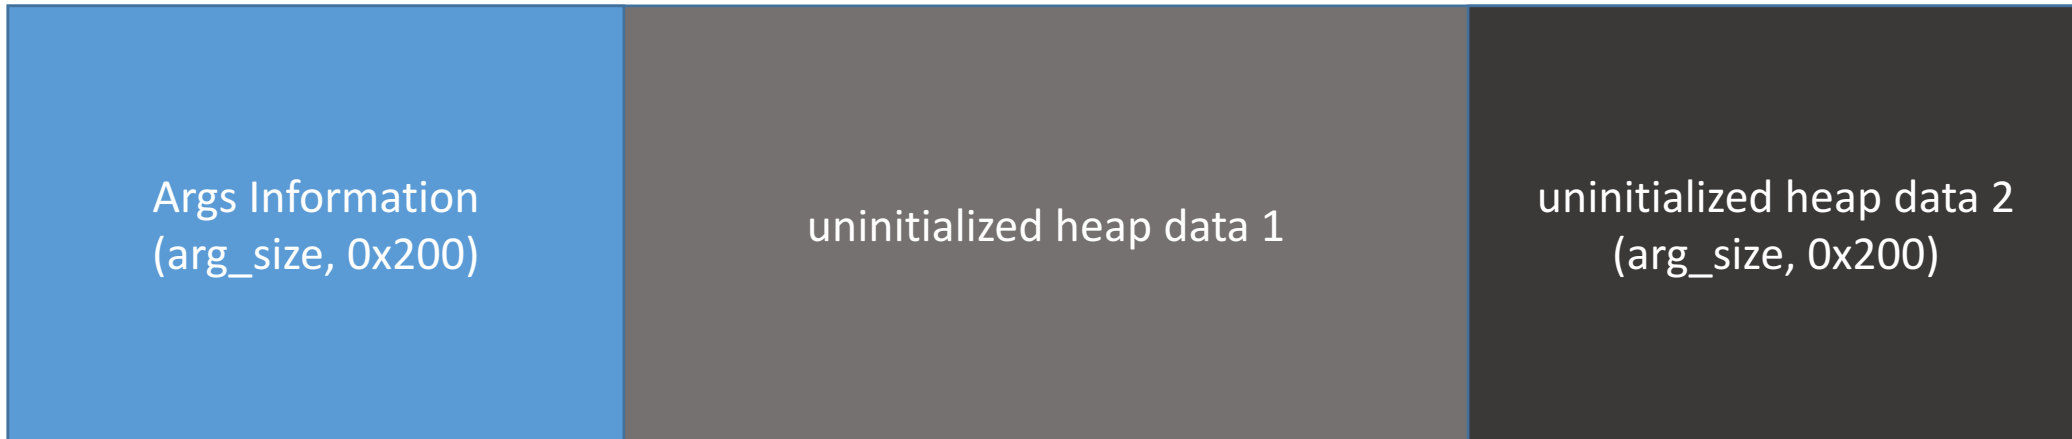
- copyout the 0x200 bytes leaked heap information to user buffer



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  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8  
0xffffffff80256eb1b8 0xffffffff80256eb1b8
```



Destroy the ports memory:

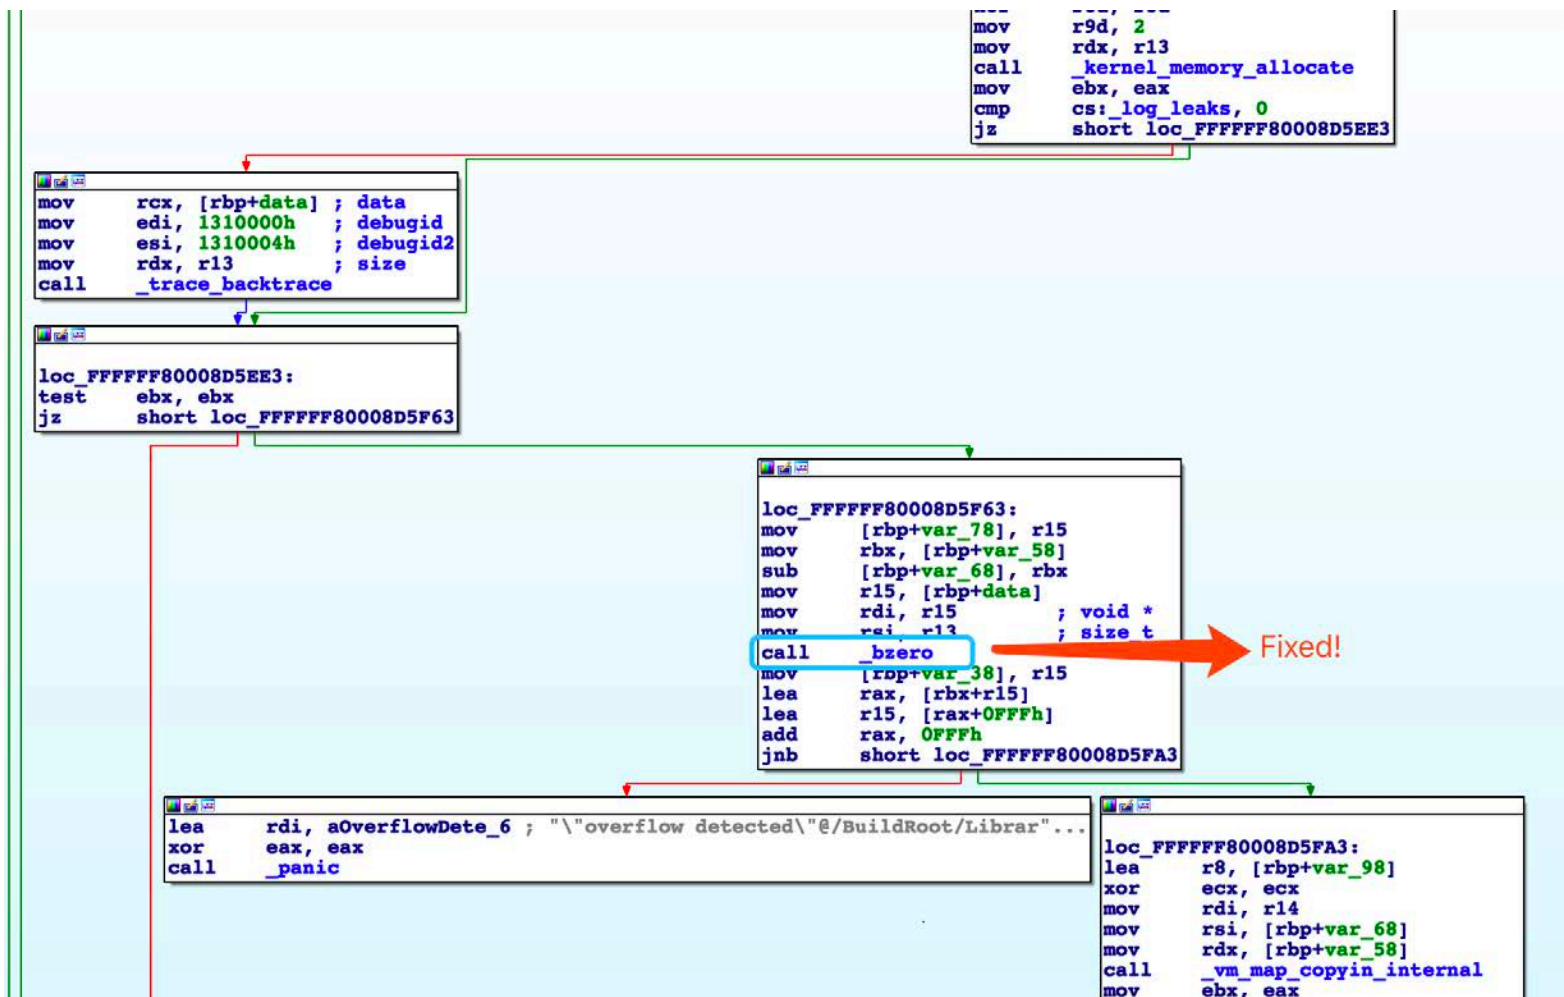
```
mach_port_destroy(mach_task_self(), q);
```



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


# CVE-2018-4413

Apple fixed it by calling bzero.



# CVE-2018-4425

## NECP Attack Surface 1



- necp\_client\_add
- necp\_client\_remove
- necp\_client\_copy
- necp\_client\_list
- necp\_client\_agent\_action
- necp\_client\_copy\_agent
- necp\_client\_agent\_use
- necp\_client\_copy\_interface
- necp\_client\_copy\_route\_statistics
- necp\_client\_update\_cache
- necp\_client\_copy\_client\_update

```
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\_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) {
        NECPLLOG(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;
            }
        }
    }
}

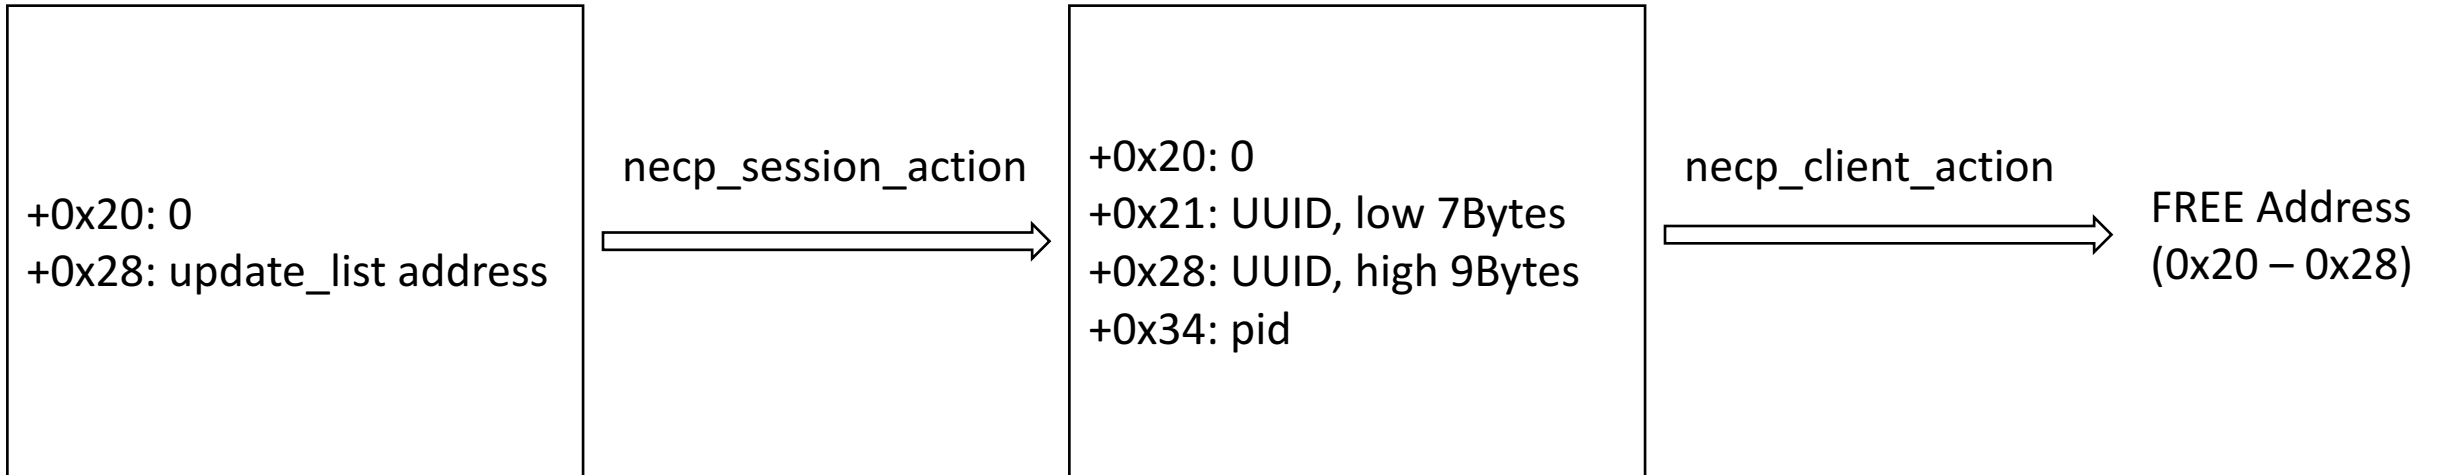
FREE(client_update, M_NECP); ---(f)
client_update = NULL;
} else {
    error = ENOENT;
}

return (error);
}
```



## CVE-2018-4425

Exploit : arbitrary address free



**For Example**, we set MachO UUID(16 bytes) as 41414141414141414141414141414141, here we get 0x4141414141414100 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
0          mov [rbp-50h], rdx
1          mov r13, rsi
2          lea rax, __stack_chk_guard
3          mov rax, [rax]
4          mov [rbp-30h], rax
5          movsxd rbx, dword ptr [r13+0]
6          call _current_proc
7          mov r15, rax
8          lea r12, [r15+0C0h]
9          mov rdi, r12
0          call _lck_mtx_lock_spin_always
1          mov r14d, 9
2          test rbx, rbx
3          js loc_FFFFFFFF80006BD34B
4          mov rax, [r15+0E8h]
5          test rax, rax
6          jz loc_FFFFFFFF80006BD34B
7          cmp [rax+48h], ebx
8          jle loc_FFFFFFFF80006BD34B
9          mov rcx, [rax]
0          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
0          jnz loc_FFFFFFFF80006BD336
1          mov rbx, [rax+38h] ; fg_data
2          mov r14d, 16h
3          cmp byte ptr [rbx], 1 ; sub_type check Fixed!
4          jnz loc_FFFFFFFF80006BD34B
5          mov rdi, r12
6          call _lck_mtx_unlock
7          lea r15, [rbx+18h]
    
```

BD235: \_necp\_session\_action+A5 (Synchronized with Hex View-1)

```

0      public _necp_client_action
0      _necp_client_action proc near
0          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
0          lea rax, __stack_chk_guard
1          mov rax, [rax]
2          mov [rbp-30h], rax
3          mov [rbp-418h], rsi
4          movsxd r14, dword ptr [rsi]
5          call _current_proc
6          mov r12, rax
7          lea r15, [r12+0C0h]
8          mov rdi, r15
9          call _lck_mtx_lock_spin_always
0          mov ebx, 9
1          test r14, r14
2          js loc_FFFFFFFF80006DE5E3
3          mov rax, [r12+0E8h]
4          test rax, rax
5          jz loc_FFFFFFFF80006DE5E3
6          cmp [rax+48h], r14d
7          jle loc_FFFFFFFF80006DE5E3
8          mov rcx, [rax]
9          mov rdx, [rcx+r14*8]
0          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
0          mov r14, [rax+38h] ; fg_data
1          mov ebx, 16h
2          cmp byte ptr [r14], 2 ; sub_type check Fixed!
3          jnz loc_FFFFFFFF80006DE5E3
4          mov rdi, r15
5          call _lck_mtx_unlock
6          mov r12, [rbp-418h]
    
```

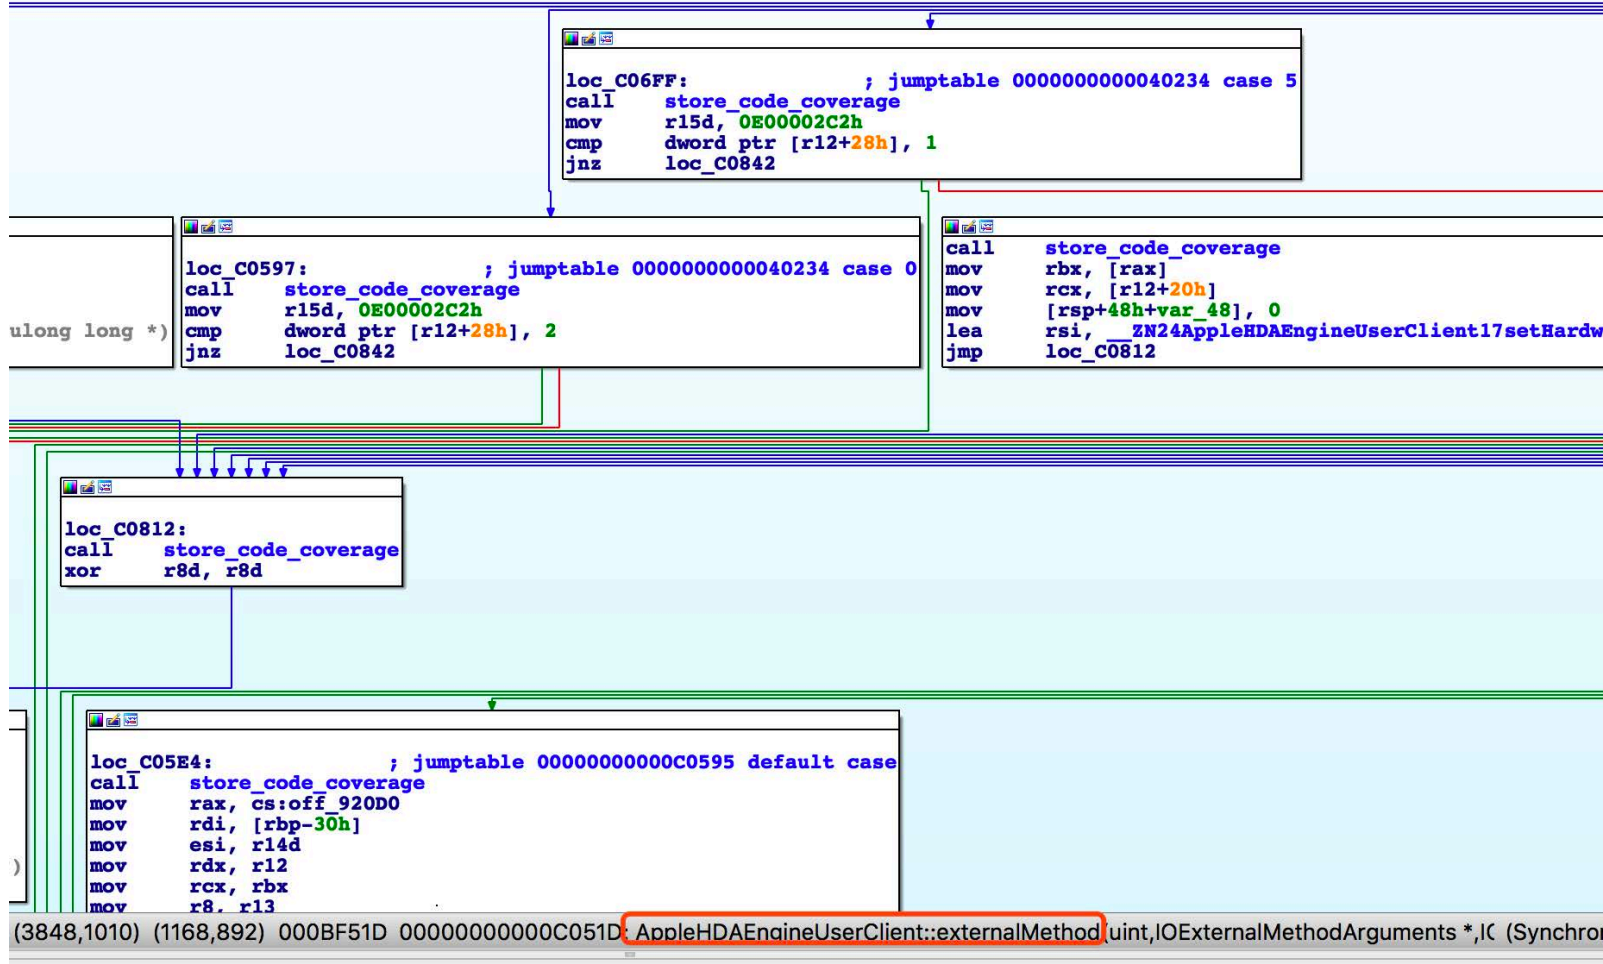
DE4E3: \_necp\_client\_action+B3 (Synchronized with Hex View-1)

# Future Plan of Our Fuzzing Tool

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



## IOKit Code Coverage Example





## 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)



## More Information

- follow me on twitter: @panicall



# Acknowledge

- Google Project Syzkaller<sup>1</sup>

ANY QUESTIONS?