

# Debugging the XNU Kernel with IDA Pro

## Table of Contents

1. Purpose .....	1
2. Debugging OSX with VMware .....	1
2.1. Quick Start .....	1
2.2. Using the KDK .....	3
2.3. Debugging a Development Kernel .....	4
2.4. Assembly-Level Debugging + DWARF .....	5
2.5. KEXT Debugging .....	6
2.6. Debugging a Prelinked Kernelcache .....	7
2.7. Debugging the OSX Kernel Entry Point .....	10
3. macOS11 Kernel Debugging .....	12
3.1. KernelCollections .....	12
3.2. Creating a macOS11 VM .....	12
3.3. Debugging: Quick Start .....	13
3.4. Symbolicating KernelCollections .....	14
3.5. Debugging KernelCollections .....	16
3.6. macOS11 + DWARF .....	17
4. UEFI Debugging .....	17
4.1. Debugging the OSX Bootloader .....	18
4.2. GetMemoryMap .....	19
4.3. UEFI Debugging + DWARF .....	19
5. Debugging iOS with Corellium .....	21
5.1. Quick Start .....	21
5.2. Creating a KDK for iOS .....	23
5.3. Debugging the iOS Kernel Entry Point .....	24
6. Known Issues and Limitations .....	25
6.1. iBoot debugging .....	25
6.2. 32-bit XNU .....	26
6.3. KDP .....	26

Last updated on November 1, 2020 – v1.1

## 1. Purpose

IDA 7.3 introduces the Remote XNU Debugger. It is designed to communicate with the GDB stub included with popular virtualization tools, namely VMware Fusion (for OSX) and Corellium (for iOS). The debugger allows you to observe the Darwin kernel as it is running, while at the same time utilising the full power of IDA's analysis capabilities. It works equally well on Mac, Windows, and Linux.

This writeup is intended to quickly get you familiar with debugger, as well as offer some hints to make the experience as smooth as possible.

## 2. Debugging OSX with VMware

### 2.1. Quick Start

To get started with debugging OSX, we will perform a simple experiment. This is the same experiment outlined in [this great writeup by GeoSn0w](#), but we will be performing the equivalent in IDA - which we hope you'll find is much simpler.

Begin with the following setup:

1. create an OSX virtual machine with [VMware Fusion](#). in this example the VM is OSX 10.13.6, but the experiment should work with any recent OSX version.
2. open Terminal in the VM and enable some basic XNU debugging options:

```
$ sudo nvram boot-args="slide=0 debug=0x100 keepsyms=1"
```

- shut down the VM and add the following line to the `.vmx` file:

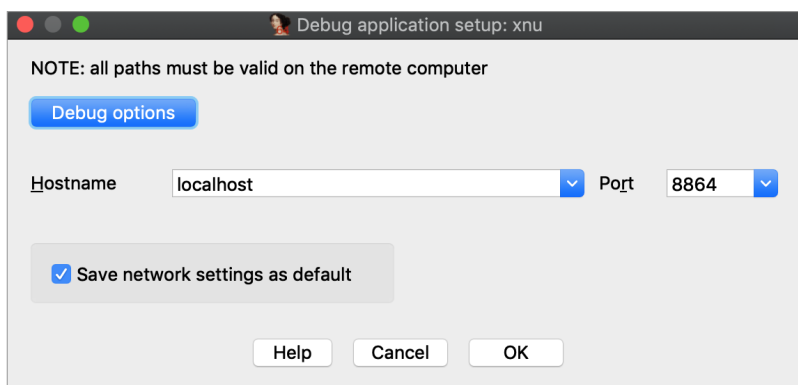
```
debugStub.listen.guest64 = "TRUE"
```

- power on the virtual machine, open Terminal, and run this command:

```
$ uname -v
Darwin Kernel Version 17.7.0 ... root:xnu-4570.71.17~1/RELEASE_X86_64
```

Let's use IDA to modify this version string.

Launch IDA, and when prompted with the window **IDA: Quick start**, choose **Go** to start with an empty database. Then go to menu **Debugger>Attach>Remote XNU Debugger** and set the following options:



Click OK, then select **<attach to the process started on target>**, and wait for IDA to attach. This step might take a few seconds (later we'll discuss how to speed things up). Once attached, the target is usually suspended in `machine_idle`:

IDA should have printed the message **FFFFFF8000200000: process kernel has started**, meaning it successfully detected the kernel image in memory. Now let's find the version string. Conveniently, the string appears in the kernel's symbol table, so we can simply use shortcut **G** and enter the name `_version` to jump right to it:

```

TEXT:  const:FFFFFF8000AF6A00  _version db 44h ; D
TEXT:  const:FFFFFF8000AF6A01  db 61h ; a
TEXT:  const:FFFFFF8000AF6A02  db 72h ; r
TEXT:  const:FFFFFF8000AF6A03  db 77h ; w
TEXT:  const:FFFFFF8000AF6A04  db 69h ; i
TEXT:  const:FFFFFF8000AF6A05  db 6Eh ; n
TEXT:  const:FFFFFF8000AF6A06  db 20h
TEXT:  const:FFFFFF8000AF6A07  db 4Bh ; K
TEXT:  const:FFFFFF8000AF6A08  db 65h ; e
TEXT:  const:FFFFFF8000AF6A09  db 72h ; r
TEXT:  const:FFFFFF8000AF6A0A  db 6Eh ; n
TEXT:  const:FFFFFF8000AF6A0B  db 65h ; e
TEXT:  const:FFFFFF8000AF6A0C  db 6Ch ; l

```

Use IDAPython to overwrite the bytes at this address:

```
idaapi.dbg_write_memory(0xFFFFF8000AF6A00, "IDAPRO".encode('utf-8'))
```

Resume the process and allow the VM to run freely. Go back to Terminal in the VM and run the same command as before:

```
$ uname -v
IDAPRO Kernel Version 17.7.0 ... root:xnu-4570.71.17~1/RELEASE_X86_64
```

The output should look almost the same, except **Darwin** has been replaced with **IDAPRO**. So, we have modified kernel memory without breaking anything! You can continue to explore memory, set breakpoints, pause and resume the OS as you desire.

## 2.2. Using the KDK

If you have installed a [Kernel Development Kit](#) from Apple, you can set **KDK\_PATH** in `dbg_xnu.cfg` to enable DWARF debugging:

```
KDK_PATH = "/Library/Developer/KDKs/KDK_10.13.6_17G4015.kdk";
```

Even if there is no KDK available for your OSX version, you can still utilise the `KDK_PATH` option in IDA to speed up debugging. For example, in the experiment above we could have done the following:

1. make your own KDK directory:

```
$ mkdir ~/MyKDK
```

2. copy the kernelcache from your VM:

```
$ scp user@vm:/System/Library/PrelinkedKernels/prelinkedkernel ~/MyKDK
```

3. decompress the kernelcache:

```
$ kextcache -c ~/MyKDK/prelinkedkernel -uncompressed
```

4. set `KDK_PATH` in `dbg_xnu.cfg`:

```
KDK_PATH = "~/MyKDK";
```

Now whenever IDA needs to extract information from the kernel or kexts, it will parse the kernelcache file on disk instead of parsing the images in memory. This should be noticeably faster.

## 2.3. Debugging a Development Kernel

Our next goal is to use the KDK to create a rich database that can be used to debug XNU in greater detail. In this example we will debug the development kernel included in the Apple KDK. Let's open this file in IDA:

```
$ export KDK=/Library/Developer/KDKs/KDK_10.13.6_17G4015.kdk
$ export KERNELS=$KDK/System/Library/Kernels
$ ida64 -okernel.i64 $KERNELS/kernel.development
```

Wait for IDA to load the DWARF info and complete the autoanalysis. This may take a few minutes, but we only need to do it once.

While we wait, we can prepare the virtual machine to use the development kernel instead of the release kernel that is shipped with OSX (Note: System Integrity Protection must now be disabled in the VM). Open Terminal in the VM and run the following commands:

1. copy the development kernel from the KDK:

```
$ sudo scp user@host:"\${KERNELS}/kernel.development" /System/Library/Kernels/
```

2. reconstruct the kernelcache:

```
$ sudo kextcache -i /
```

3. reboot:

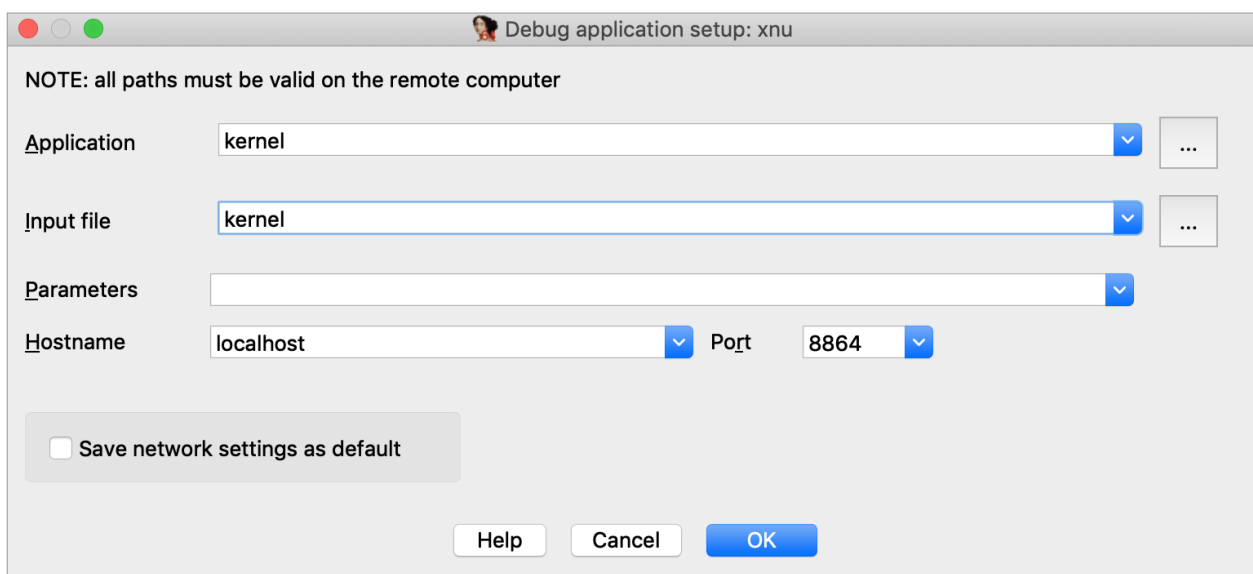
```
$ sudo shutdown -r now
```

4. after rebooting, check that the development kernel was properly installed:

```
$ uname -v
... root:xnu-4570.71.17~1/DEVELOPMENT_X86_64
```

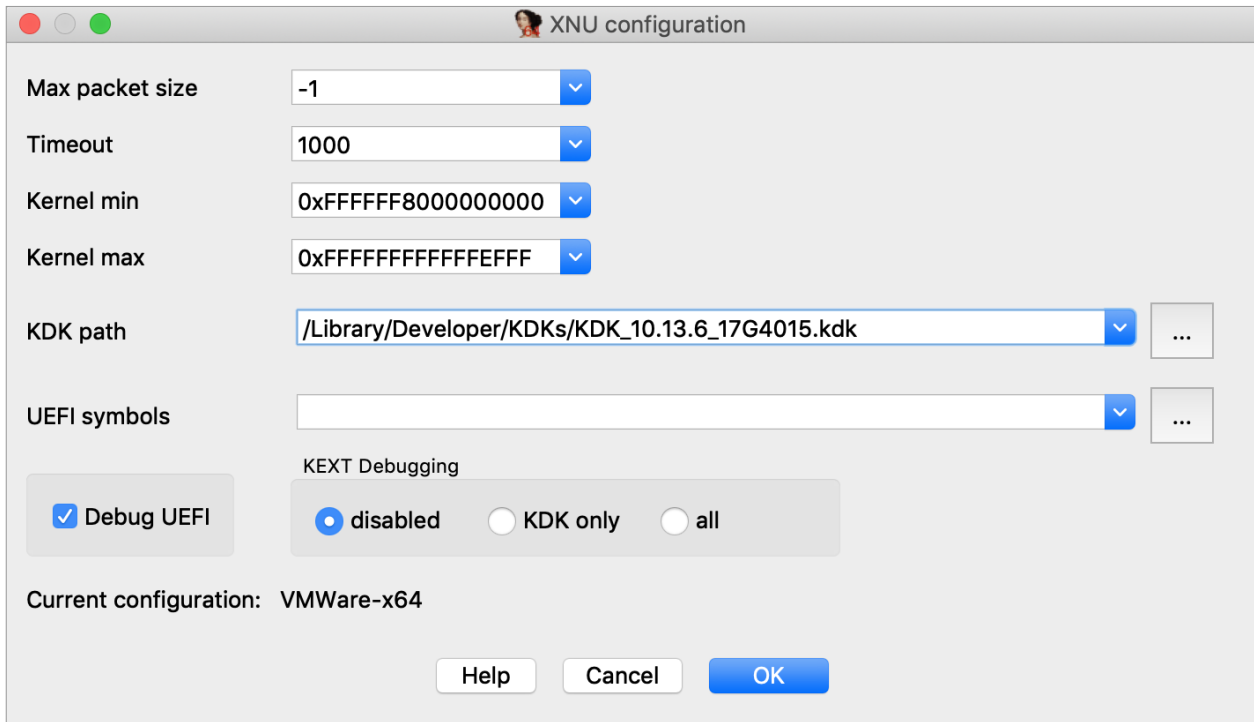
The VM is now ready for debugging.

Return to IDA and use **Debugger>Select debugger** to select **Remote XNU Debugger**. Then open **Debugger>Process options** and set the following fields:



Now go to **Debugger>Debugger options>Set specific options** and make sure the **KDK path** field is set:



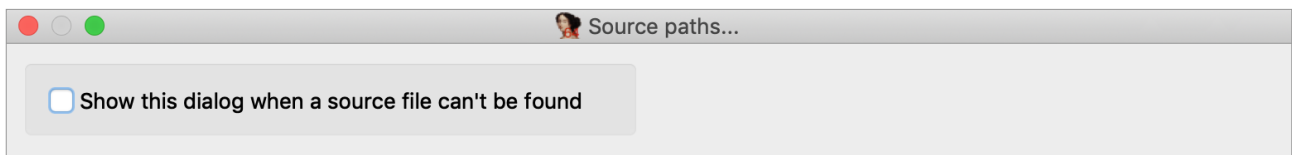


You can ignore the other options for now, and press OK.

## 2.4. Assembly-Level Debugging + DWARF

IDA supports source-level debugging for the XNU Kernel. However for demonstration purposes we will focus on assembly-level debugging, while taking advantage of source-level DWARF information like local variables. This is a bit more stable, and is still quite useful.

Before attaching the debugger, open **Options>Source paths...** and un-check the checkbox:



Then click **Apply**. This will prevent IDA from complaining when it can't find a source file.

Finally, select **Debugger>Attach to process>attach to the process started on target**. After attaching, jump to function **dofileread**, and use **F2** to set a breakpoint. Resume the debugger and wait for the breakpoint to be hit (typically it will be hit right away, if not try simply running a terminal command in the guest). Once XNU hits our breakpoint, open **Debugger>Debugger windows>Locals**:

The screenshot displays the IDA View-RIP window with assembly code for a subroutine named 'dofileread'. The assembly includes instructions like 'push rbp', 'mov rbp, rsp', 'push r15', etc. The 'General registers' window shows the state of registers like RAX, RBX, RCX, etc. The 'Locals' window shows variables like 'ctx', 'fp', 'f\_flags', 'f\_iocount', etc.

We can now perform detailed instruction-level debugging with the assistance of DWARF. You can continue to single step, set breakpoints, and inspect or modify local variables just like any other IDA debugger.

## 2.5. KEXT Debugging

IDA also supports debugging kext binaries. To demonstrate this, we will debug **IONetworkingFamily**, a submodule of IOKit that is typically shipped with the KDK. Begin by opening the binary in IDA:

```
$ export KEXTS=$KDK/System/Library/Extensions
$ ida64 -onet.i64 $KEXTS/IONetworkingFamily.kext/Contents/MacOS/IONetworkingFamily
```

Select **Remote XNU Debugger** from the debugger menu. Then in **Debugger>Process options**, set:

Application	kernel	...
Input file	com.apple.iokit.IONetworkingFamily	...
Parameters		
Hostname	localhost	Port 8864

Note that we provide the bundle ID of the kext (com.apple.iokit.IONetworkingFamily) as the **Input file** field. This allows the debugger to easily identify the target kext at runtime.

Also note that loading all kexts in kernel memory can be a slow operation, which is why it is disabled by default. Open **Debugger>Debugger options>Set specific options** and ensure the **KDK path** field is set, then set the **KEXT Debugging** option to **KDK only**:

KEXT Debugging
<input type="radio"/> disabled <input checked="" type="radio"/> KDK only <input type="radio"/> all

This tells the debugger to only load kexts that are present in the KDK. Since the KDK binaries are on the local filesystem, IDA can parse the kexts in a negligible amount of time - which is ideal since we're really only interested in IONetworkingFamily.

Now power on your VM and allow it to boot up. Once it is running idle, attach the debugger. Immediately IDA should

detect the kernel and all relevant kexts in memory, including IONetworkingFamily:

Path	Base
com.apple.iokit.IOGraphicsFamily	FFFFFFF7F80D9D000
com.apple.iokit.IOStorageFamily	FFFFFFF7F80EB7000
com.apple.iokit.IONetworkingFamily	FFFFFFF7F81079000
com.apple.iokit.IOHIDFamily	FFFFFFF7F81193000
com.apple.iokit.IOUSBHostFamily	FFFFFFF7F81253000

Double-click to bring up the debug names for this module, and search for **IONetworkInterface::if\_ioctl**:

Name	Address
IONetworkInterface::free (void)	FFFFFFF7F81084900
IONetworkInterface::isPrimaryInterface (void)	FFFFFFF7F81084AEC
IONetworkInterface::getController (void)	FFFFFFF7F81084B44
IONetworkInterface::initIfnet (ifnet *)	FFFFFFF7F81084B58
IONetworkInterface::initIfnetParams (ifnet_init_params *)	FFFFFFF7F81084B60
IONetworkInterface::if_output (_ifnet *, _mbuf *)	FFFFFFF7F81084BC4
IONetworkInterface::if_ioctl (_ifnet *,ulong,void *)	FFFFFFF7F81084D62
IONetworkInterface::if_set_bpf_tap (_ifnet *,uint,int (*) (...	FFFFFFF7F81084DE2
IONetworkInterface::if_detach (_ifnet *)	FFFFFFF7F81084F50
IONetworkInterface::configureOutputStartDelay (ushort,ushor...	FFFFFFF7F81084FA4
IONetworkInterface::lock (void)	FFFFFFF7F81085006
IONetworkInterface::unlock (void)	FFFFFFF7F8108501C

Line 514 of 1017

Now set a breakpoint at this function and resume the OS. Typically the breakpoint will be hit right away, but if it isn't try performing an action that requires a network interface (for instance, performing a google search). Once execution breaks in the kext we can use the database to debug it in detail:

IDA View-RIP

```

text:FFFFFFF7F81084D62 ; errno_t __fastcall IONetworkInterface::if_ioctl(ifnet_t ifp, unsigned __int64 cmd, void *data)
text:FFFFFFF7F81084D62 public _ZN18IONetworkInterface8if_ioctlEP7_ifnetmPv
text:FFFFFFF7F81084D62 _ZN18IONetworkInterface8if_ioctlEP7_ifnetmPv proc near
text:FFFFFFF7F81084D62 ; DATA XREF: IONetworkInterface::initIfnetParams(ifnet_init
text:FFFFFFF7F81084D62 ifp = rdi ; ifnet_t
text:FFFFFFF7F81084D62 cmd = rsi ; unsigned __int64
text:FFFFFFF7F81084D62 data = rdx ; void *
text:FFFFFFF7F81084D62 push rbp
text:FFFFFFF7F81084D63 mov rbp, rsp
text:FFFFFFF7F81084D66 push r15
text:FFFFFFF7F81084D68 push r14
text:FFFFFFF7F81084D6A push r13
text:FFFFFFF7F81084D6C push r12
text:FFFFFFF7F81084D6E push rbx
text:FFFFFFF7F81084D6F push rax
text:FFFFFFF7F81084D70 mov r14, data
text:FFFFFFF7F81084D73 data = r14 ; void *
text:FFFFFFF7F81084D73 mov r15, cmd
text:FFFFFFF7F81084D76 cmd = r15 ; unsigned __int64
text:FFFFFFF7F81084D76 mov r12, ifp
text:FFFFFFF7F81084D79 ifp = r12 ; ifnet_t
text:FFFFFFF7F81084D79 call near ptr ifnet_softc
0000BD62 FFFFFFFF7F81084D62: IONetworkInterface::if_ioctl(_ifnet *,ulong,void *) (Synchronized with RIP)

```

Call Stack

Address	Module	Function
FFFFFFF7F81084D73	com.app...	IONetworkInterface::if_ioctl
FFFFFFF8000650FE5	kernel	_ifnet_ioctl+245
FFFFFFF8000647A5F	kernel	_ifioctl+94F
FFFFFFF80008B0838	kernel	_soioctl+228
FFFFFFF80008519E0	kernel	_fo_ioctl+40
FFFFFFF80008AA8C7	kernel	_ioctl+527
FFFFFFF80009A4C17	kernel	_unix_syscall64+2D7
FFFFFFF800031EA30	kernel	_hndl_unix_scall64+10

Locals

Name	Value	Type
ifp	0xFFFFF800B12B618LL	ifnet_t
cmd	0xC02C6938LL	unsigned __int64
data	0xFFFFF8099BCBBE90LL	void *

Line 1 of 8

## 2.6. Debugging a Prelinked Kernelcache

For simplicity, all of the examples up until now have dealt with a subset of the kernel, but it is also possible to load a complete prelinked kernelcache in IDA and debug it. Naturally, we have some suggestions for this.

## 2.6.1. Extending the KDK

If you're interested in debugging the entire prelinked kernel, the biggest concern is speed. IDA must create a detailed and accurate depiction of kernel memory, which could contain hundreds of kext modules. If we're not careful, this can be slow.

Fortunately there is an easy solution. Try the following:

1. create a writable copy of Apple's KDK:

```
$ cp -r /Library/Developer/KDKs/KDK_10.13.6_17G4015.kdk ~/MyKDK
```

2. copy the kernelcache from your VM to the new KDK:

```
$ scp user@vm:/System/Library/PrelinkedKernels/prelinkedkernel ~/MyKDK
```


3. decompress the kernelcache:

```
$ kextcache -c ~/MyKDK/prelinkedkernel -uncompressed
```

Now IDA can use both the KDK and the kernelcache to extract debugging information for almost any kext at runtime. This should be fast.

## 2.6.2. Loading the Kernelcache

When loading a kernelcache, IDA now offers more load options:



In this example we want to load everything, so choose the **kernel + all kexts** option and wait for IDA to load all the subfiles and finish the autoanalysis. This will take a while but there's no way around it, it's a lot of code.


**IMPORTANT NOTE:** Try to avoid saving the IDA database file in the KDK directory. It is important to keep irrelevant files out of the KDK since they might slow down IDA's KDK parsing algorithm.

Now we might want to improve the static analysis by loading DWARF info from the KDK. In IDA 7.3 the dwarf plugin supports batch-loading all DWARF info from a KDK into a kernelcache database. Currently this feature must be invoked manually, so we have provided [this script](#) to make it easier.

Copy `kdk_utils.py` to the plugins directory of your IDA installation. This plugin will create a new menu **Edit>Other>KDK utils**, with two new menu actions:

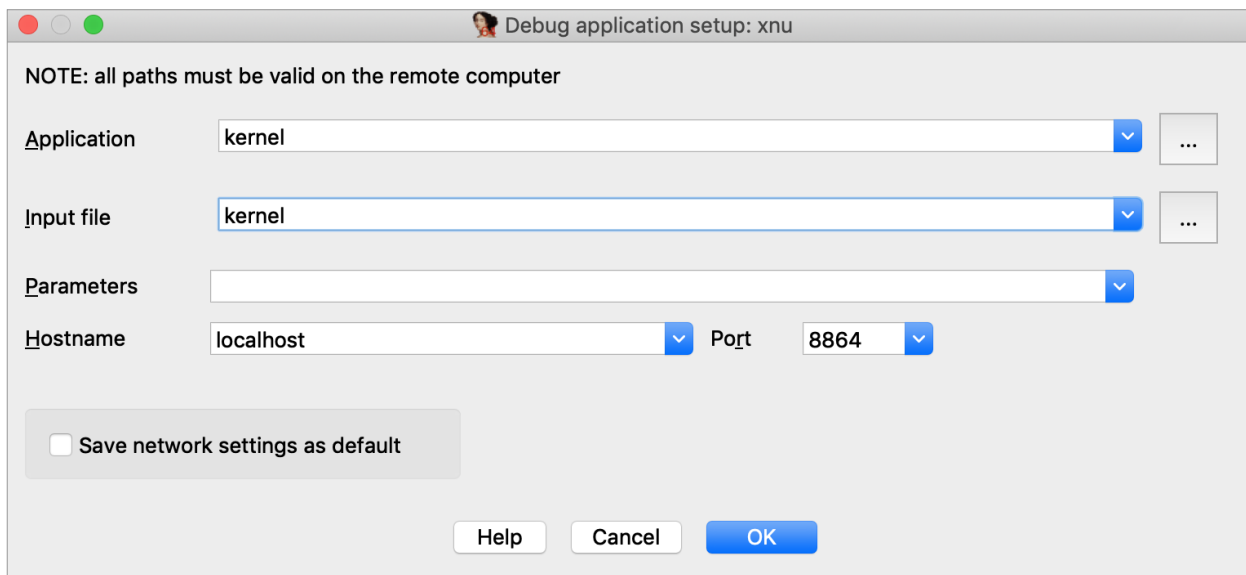
- **Load KDK:** This action will automatically detect all matching DWARF files in a given KDK, then apply the DWARF info to the subfiles in the database (including the kernel itself).
- **Load DWARF for a prelinked KEXT:** This action is useful if you have DWARF info for a prelinked kext that is not included in Apple's KDK. For a given DWARF file, the action will find a matching kext in the database and apply the DWARF info to this subfile.

Try opening **Edit>Other>KDK utils>Load KDK** and provide the KDK path:

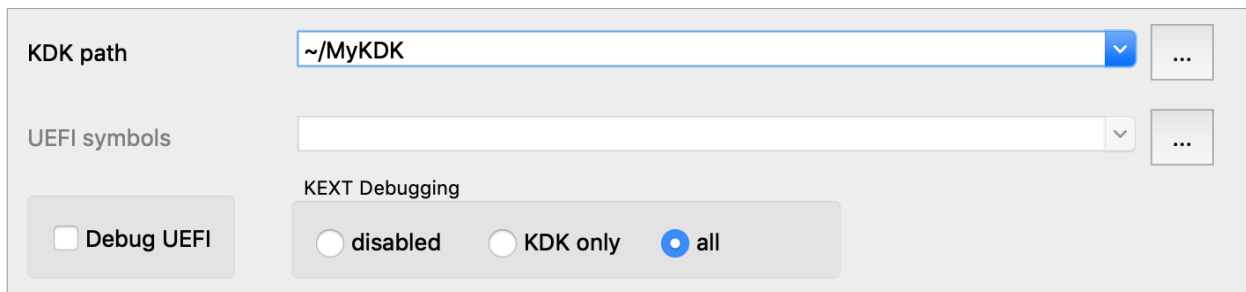


Wait for IDA to scan the KDK for matching DWARF files and load them. This operation can also take a while, but it's worth it for all the extra structures, prototypes, and names that are added to the database. In the end we have a very detailed database that we are ready to use for debugging.

Now open **Debugger>Process options** and set the following options:



Then open **Debugger>Debugger options>Set specific options** and set the following fields:



Note that we set the **KEXT Debugging** option to **all**. This tells the debugger to detect every kext that has been loaded into memory and add it to the Modules list, including any non-prelinked kexts (there are likely only a handful of them, so it doesn't hurt).

Finally, power on the VM and attach to it with **Debugger>Attach to process>attach to the process started on target**. IDA should be able to quickly generate modules for the kernel and all loaded kexts:

The screenshot shows the IDA View-RIP window with assembly code and the General registers window. The assembly code includes instructions like `test rax, rax`, `short loc_FFFFFFFF80004D5E80`, `mov rax, cs:earlyMaxIntDelay, 0`, `call rax`, `loc_FFFFFFFF80004D5E80:`, `cmp cs:pmInitDone, 0`, `jz short loc_FFFFFFFF80004D5E8C`, `mov rax, cs:pmDispatch`, `test rax, rax`, `short loc_FFFFFFFF80004D5E8C`, `mov rax, [rax+10h]`, `test rax, rax`, `short loc_FFFFFFFF80004D5E8C`, `mov rdi, 7FFFFFFFFFFFFFFFh`, `call rax`, `jmp short loc_FFFFFFFF80004D5EDF`, `loc_FFFFFFFF80004D5EDC:`, `sti`, `hlt`, `cli`, `loc_FFFFFFFF80004D5EDF:`, `pushfq`, `pop rax`, `test ah, 2`, `jnz loc_FFFFFFFF80004D6001`, `loc_FFFFFFFF80004D5E8A:`, `mov rax, qword ptr gs:unk_0`, `and qword ptr [rax+100h], 0`, `call _do_mfence`, `mov rax, qword ptr gs:unk_0`, `dword ptr [rax+108h], 0`, `jz short loc_FFFFFFFF80004D5F17`, `call process_pmap_updates`, `loc_FFFFFFFF80004D5F17:`. The General registers window shows values for RAX, RBX, RCX, RDX, RSI, RDI, RBP, RSP, R8, R9, R10, R11, R12, R13, R14, R15, RIP, and RFL. The Modules window shows a list of loaded modules with their paths and bases.

You are now free to explore the entire running kernel! Try performing any of the previous demos in this writeup. They should work about the same, but now they are all possible with one single database.

### 2.6.3. Kernel ASLR + Rebasing

It is worth noting that rebasing has been heavily improved in IDA 7.3. Even large databases like the one we just created can now be rebased in just a few seconds. Previous IDA versions would take quite a bit longer. Thus, IDA should be able to quickly handle kernel ASLR, even when working with prelinked kernelcaches.

## 2.7. Debugging the OSX Kernel Entry Point

In this example we demonstrate how to gain control of the OS as early as possible. This task requires very specific steps, and we document them here. Before we begin, we must make an important note about a limitation in VMware's GDB stub.

### 2.7.1. Limitation in VMware

Currently VMware's 64-bit GDB stub does not allow us to debug the kernel entry point in physical memory. According to VMware's support team, the correct approach is to use the 32-bit stub to debug the first few instructions of the kernel, then switch to a separate debugger connected to the 64-bit stub once the kernel switches to 64-bit addressing.

Since IDA's XNU debugger does not support 32-bit debugging, this approach is not really feasible (and it's not very practical anyway).

### 2.7.2. Workaround

Rather than add support for the 32-bit stub just to handle a few instructions, the official approach in IDA is to break at the first function executed in virtual memory (`i386_init`). This allows us to gain control of the OS while it is still in the early stages of initialization, which should be enough for most use cases.

Here's how you can do it:

1. Disable ALSR for the kernel. Open Terminal in the VM and run the following command:

```
sudo nvram boot-args="slide=0"
```

Then power off the VM.

2. Add this line to the `.vmx` file:



```
debugStub.hideBreakpoints = "TRUE"
```


This ensures that hardware breakpoints are enabled in the GDB stub. For most versions of VMware, TRUE is the default value, but it's better to be safe.

- Also add this line to the .vmx file:

```
monitor.debugOnStartGuest64 = "TRUE"
```


This will tell VMware to suspend the OS before it boots.

- Power on the VM. It will remain suspended until we attach the debugger.
- Load a kernel binary in IDA, and set the following XNU debugger options:



Note that we un-checked the **Debug UEFI** option. This option is explained in detail in the [UEFI Debugging](#) section, but for now just ensure it is disabled. This will prevent IDA from doing any unnecessary work.

- Attach the debugger. The VM will be suspended in the firmware before the boot sequence has begun:




- Now jump to the function `_i386_init` and set a hardware breakpoint at this location:

```
idaapi.add_bpt(here(), 1, BPT_EXEC)
```

We must use a hardware breakpoint because the kernel has not been loaded and the address is not yet valid. This is why steps 1 and 2 were important. It ensures the stub can set a breakpoint at a deterministic address, without trying to write to memory.

- Resume the OS, and wait for our breakpoint to be hit:



IDA should detect that execution has reached the kernel and load the kernel module on-the-fly. You can now continue to debug the kernel normally.

## 3. macOS11 Kernel Debugging

VMware Fusion 12 now supports macOS11 Big Sur, including support for the built-in gdb server.

There have been a lot of changes to the XNU kernel architecture in macOS11, so it might be interesting to use IDA to discover these changes and discuss how they affect our analysis.

**NOTE:** full support for macOS11 kernel debugging requires IDA 7.5 SP3

### 3.1. KernelCollections

The most notable change in macOS11 is the kernelcache format. In previous macOS versions a kernelcache was just the mach kernel binary, prelinked with some extra segments containing embedded KEXTs.

In macOS11, Apple introduced an entirely new file format (MH\_FILESET) specifically designed to be a generic container for the mach kernel and KEXTs. So a macOS11 kernelcache is a non-executable file that simply advertises a list of subfiles contained within it, via the LC\_FILESET\_ENTRY load commands.

Such files are now called "KernelCollections", and they are found in the /System/Library/KernelCollections/ directory of your macOS11 installation.

Let's try loading one of them in IDA:

```
$ ida64 -o/tmp/boot.i64 /System/Library/KernelCollections/BootKernelExtensions.kc
```

This cache contains the essential modules that macOS11 needs to boot, including the mach kernel itself (com.apple.kernel):

```
TEXT:HEADER:FFFFFF80002085C8 ; LC_FILESET_ENTRY
TEXT:HEADER:FFFFFF80002085C8 fileset_entry_command <80000035h, 38h, 0FFFFFF8000210000h, 10000h, <\
TEXT:HEADER:FFFFFF80002085C8 20h>, 0>
TEXT:HEADER:FFFFFF80002085E8 aComAppleKernel db 'com.apple.kernel',0 ; entry id
TEXT:HEADER:FFFFFF80002085F9 align 20h
TEXT:HEADER:FFFFFF8000208600 ; LC_FILESET_ENTRY
TEXT:HEADER:FFFFFF8000208600 fileset_entry_command <80000035h, 48h, 0FFFFFF8001268000h, 1108000h, <\
TEXT:HEADER:FFFFFF8000208600 20h>, 0>
TEXT:HEADER:FFFFFF8000208620 aComAppleNkeApp db 'com.apple.nke.applicationfirewall',0 ; entry id
TEXT:HEADER:FFFFFF8000208642 align 8
TEXT:HEADER:FFFFFF8000208648 ; LC_FILESET_ENTRY
TEXT:HEADER:FFFFFF8000208648 fileset_entry_command <80000035h, 48h, 0FFFFFF8001272000h, 1112000h, <\
TEXT:HEADER:FFFFFF8000208648 20h>, 0>
TEXT:HEADER:FFFFFF8000208668 aComAppleDriver db 'com.apple.driver.AppleACPIPlatform',0 ; entry id
TEXT:HEADER:FFFFFF800020868B align 10h
```

Also present is another KernelCollection that contains many important system KEXTs:

```
/System/Library/KernelCollections/SystemKernelExtensions.kc
```

And yet another one is used to manage third-party KEXTs, located at:

```
/Library/KernelCollections/AuxiliaryKernelExtensions.kc
```

For a detailed explanation of the various "flavors" of KernelCollections, see the doc for the **kmutil** command:

```
$ man kmutil
```

It seems that all of these KernelCollections will somehow be utilized by a running instance of macOS11. Is this obvious when debugging the kernel in IDA? Let's try it out.

### 3.2. Creating a macOS11 VM

To start debugging macOS11, we'll need to set up a macOS11 virtual machine with System Integrity Protection disabled and kernel debugging options enabled:

```
$ sudo nvram boot-args="slide=0 debug=0x100 keepsyms=1"
```



To do this you will need to boot the VM in recovery mode. As of this writing, macOS11 is still in beta and booting a macOS11 VM in recovery mode is very unstable. According to [this thread](#) VMware is working on fixing it, but for now you will likely have to research the latest workaround (it changes almost every beta version).

macOS11 has also broken the debug registers, which VMware Fusion uses to set breakpoints in the guest OS. To work around this you must enable software breakpoints in the vmx file:

```
debugStub.hideBreakpoints = "FALSE"
```

According to [this](#), VMware is aware of the issue and they are working resolve it. The workaround will do for now.

Also don't forget to enable the gdb stub in the vmx file:

```
debugStub.listen.guest64 = "TRUE"
debugStub.port.guest64 = "8864"
```

### 3.3. Debugging: Quick Start

Power on the macOS11 VM and launch IDA with an empty database:

```
$ ida64 -t
```

Use menu **Debugger>Select debugger** to select the **Remote XNU Debugger**, then use **Debugger>Process** options to set the **Hostname** and **Port** fields to localhost:8864, and finally **Debugger>Attach to process** to attach to the running VM:

The screenshot shows the IDA View-RIP window with assembly code for the `com.apple.kernel` process. The code includes instructions like `cli`, `mov rax, gs:qword_0`, `qword ptr [rax+100h], 0FFFFFFFFFFFFFFFFh`, `call near ptr do_mfence`, `mov rax, gs:qword_0`, `cmp dword ptr [rax+108h], 0`, `jz short loc_FFFFFFFF80003F05A4`, `xor edi, edi`, `mov esi, 1`, `xor edx, edx`, `mov rcx, 0FFFFFFFFFFFFFFFFh`, `call near ptr unk_FFFFFFFF80003C8320`, `loc_FFFFFFFF80003F05A4: ; CODE XREF: com.apple.kernel:__text:machine_idle+1ED↑j`, `dword ptr [rbx+11C8h], 5`, `mov rdi, _pal_rtc_nanotime_info`, `call near ptr _rtc_nanotime_read`, `mov [rbx+0FF8h], rax`, `inc qword ptr [rbx+1000h]`, `sub rax, r14`, `add [rbx+0FE8h], rax`, and `cmp cs_cpu_itime_bins, rax`. The Output window shows the debugger attached to the process `<GDB remote process>` (pid=4294967294) and lists loaded kernel extensions: `SystemKernelExtensions.kc`, `AuxiliaryKernelExtensions.kc`, `BootKernelExtensions.kc`, and `com.apple.kernel`.

From the Modules list we see that IDA was able to detect the kernel image, as well as all KernelCollections that have been loaded:

Path	Base	Size
SystemKernelExtensions.kc	FFFFFFFF7F80C11000	0000000020A98000
AuxiliaryKernelExtensions.kc	FFFFFFFF7FA16B2000	000000000164000
BootKernelExtensions.kc	FFFFFFFF8000200000	000000000331A000
com.apple.kernel	FFFFFFFF8000210000	0000000000A49000
<GDB remote process>		

Note that although `com.apple.kernel` is a subset of `BootKernelExtensions.kc`, IDA still created a separate module for it.

Internally IDA expects the debugger engine to identify the executable module for any running process. The **com.apple.kernel** subfile is a logical choice, since it is the only Mach-O image in memory with file type MH\_EXECUTE.

Now let's consider the **SystemKernelExtensions.kc** module. Use **Ctrl+S** to examine the debugger segments:

Name	Start	End
SystemKernelExtensions.kc:HEADER	FFFFFFF7F80C11000	FFFFFFF7F80C1D000
SystemKernelExtensions.kc:__BRANCH_STUBS	FFFFFFF7F80C1D000	FFFFFFF7F80C35000
SystemKernelExtensions.kc:__BRANCH_GOTS	FFFFFFF7F80C35000	FFFFFFF7F80C55000
SystemKernelExtensions.kc:__info	FFFFFFF7F80C55000	FFFFFFF7F80DCD000
SystemKernelExtensions.kc:__REGION0	FFFFFFF7F80DCD000	FFFFFFF7F80E32000
SystemKernelExtensions.kc:__REGION1	FFFFFFF7F80E32000	FFFFFFF7F80E50000
SystemKernelExtensions.kc:__REGION2	FFFFFFF7F80E50000	FFFFFFF7F80E9F000
SystemKernelExtensions.kc:__REGION3	FFFFFFF7F80E9F000	FFFFFFF7F80EAD000
SystemKernelExtensions.kc:__REGION4	FFFFFFF7F80EAD000	FFFFFFF7F80EFE000
SystemKernelExtensions.kc:__REGION5	FFFFFFF7F80EFE000	FFFFFFF7F80F0B000

The **\_\_REGION\*** segments contain the KEXT subfiles. By default the debugger does not recurse into KernelCollections to detect the loaded KEXTs, because it has the potential to be very slow. All KEXTs in the KernelCollection will share a common symbol table, which can be massive and takes long time to read from debugger memory.

To avoid slowdowns, create a rudimentary KDK for macOS11 by copying the KernelCollection files from the VM:

```
$ mkdir ~/macOS11.kdk
$ cd ~/macOS11.kdk
$ scp user@vm:/System/Library/KernelCollections/BootKernelExtensions.kc .
$ scp user@vm:/System/Library/KernelCollections/SystemKernelExtensions.kc .
$ scp user@vm:/Library/KernelCollections/AuxiliaryKernelExtensions.kc .
```

Detach from the VM with **Debugger>Detach from process**, then use **Debugger>Debugger options>Set specific options** to set the following fields:

KDK path:  ...

UEFI symbols:  ...

KEXT Debugging

Debug UEFI

disabled  KDK only  all

Now reattach to the VM. Since we enabled KEXT debugging and specified a KDK, the debugger will quickly create a debugging environment that can be explored in greater detail:

Path	Base	Name	Address
SystemKernelExtensions.kc	FFFFFFF7F80C11000	IOHDACodecDriver::start(IOService *)	FFFFFFF7F9E79E60E
com.apple.driver.AppleDiskImages2	FFFFFFF7F9E439000	IOHDACodecDriver::getCodecAddress(void)	FFFFFFF7F9E79E662
com.apple.fileutil	FFFFFFF7F9E4A7000	IOHDACodecDriver::executeCommand(uint, u...	FFFFFFF7F9E79E678
com.apple.AGDCPluginDisplayMetrics	FFFFFFF7F9E541000	IOHDACodecDriver::executeVerb(ushort, us...	FFFFFFF7F9E79E690
com.apple.AppleGPUWrangler	FFFFFFF7F9E544000	IOHDACodecDriver::getAudioController(vo...	FFFFFFF7F9E79E6A8
com.apple.AppleGraphicsDeviceControl	FFFFFFF7F9E54E000	IOHDACodecDriver::createFunctionGroupNu...	FFFFFFF7F9E79E6BE
com.apple.driver.AppleHDA	FFFFFFF7F9E5C7000	IOHDACodecDriver::handleOpen(IOServic...	FFFFFFF7F9E79E812
com.apple.driver.AppleHDAController	FFFFFFF7F9E67C000	IOHDACodecDriver::handleClose(IOServic...	FFFFFFF7F9E79E892
com.apple.driver.DspFuncLib	FFFFFFF7F9E6A9000	IOHDACodecDriver::handleIsOpen(IOServic...	FFFFFFF7F9E79E8EA
com.apple.iokit.IOHDAFamily	FFFFFFF7F9E79C000	IOHDACodecDriver::message(uint, IOServic...	FFFFFFF7F9E79E920
com.apple.driver.AppleHV	FFFFFFF7F9E7A4000	IOHDACodecDriver::registerAfgPowerState...	FFFFFFF7F9E79EA04

### 3.4. Symbolicating KernelCollections

Just like with previous macOS versions, IDA allows you to load a complete macOS11 kernelcache and debug it as a single input file. This can be done for any of the various KernelCollection types.

For example, let's try loading the "Auxiliary" KernelCollection in IDA:

```
$ ida64 -o/tmp/aux.i64 ~/macOS11.kdk/AuxiliaryKernelExtensions.kc
```

This cache contains various third-party KEXTs, many of which depend on functionality in the kernel. How do these KEXTs manage to invoke kernel functions? You may have noticed in previous versions of IDA that AuxiliaryKernelExtensions.kc contains stub functions in the \_\_BRANCH\_STUBS segment:

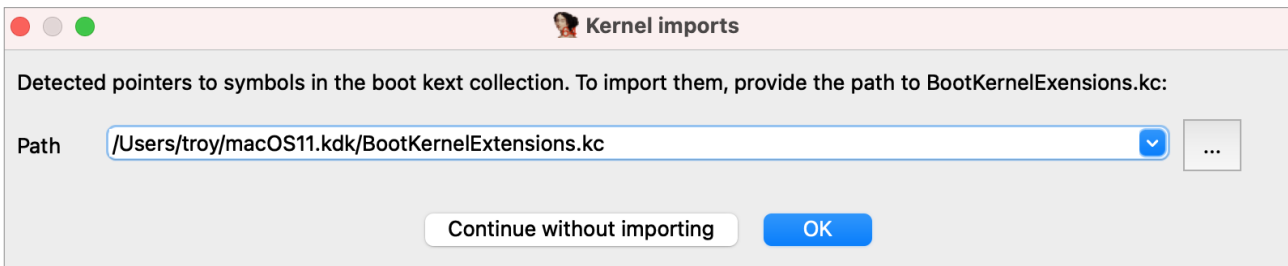
```
BRANCH_STUBS: stubs:000000000000C000 ; ===== S U B R O U T I N E =====
BRANCH_STUBS: stubs:000000000000C000
BRANCH_STUBS: stubs:000000000000C000 ; Attributes: thunk
BRANCH_STUBS: stubs:000000000000C000
BRANCH_STUBS: stubs:000000000000C000 sub_C000      proc near          ; CODE XREF: sub_1A2C4+22+p
BRANCH_STUBS: stubs:000000000000C000                                jmp          cs:qword_10000    ; HPTAbstractRAIDController::ReleaseResource
BRANCH_STUBS: stubs:000000000000C000 sub_C000      endp
BRANCH_STUBS: stubs:000000000000C000
BRANCH_STUBS: stubs:000000000000C006 ; ===== S U B R O U T I N E =====
BRANCH_STUBS: stubs:000000000000C006
BRANCH_STUBS: stubs:000000000000C006 ; Attributes: thunk
BRANCH_STUBS: stubs:000000000000C006
BRANCH_STUBS: stubs:000000000000C006 sub_C006      proc near          ; CODE XREF: sub_1A3E8+25+j
BRANCH_STUBS: stubs:000000000000C006                                jmp          cs:qword_10008
BRANCH_STUBS: stubs:000000000000C006 sub_C006      endp
BRANCH_STUBS: stubs:000000000000C006
```

The stubs read function pointers from a global offset table:

```
BRANCH_GOTS: __got:0000000000010000 ; Segment type: Pure data
BRANCH_GOTS: __got:0000000000010000 ; Segment permissions: Read/Write
BRANCH_GOTS: __got:0000000000010000 __BRANCH_GOTS__got segment byte public 'DATA' use64
BRANCH_GOTS: __got:0000000000010000 assume cs: __BRANCH_GOTS__got
BRANCH_GOTS: __got:0000000000010000 ;org 10000h
BRANCH_GOTS: __got:0000000000010000 qword_10000    dq 400000008D4D40h ; DATA XREF: sub_C000+r
BRANCH_GOTS: __got:0000000000010000                                ; Message_Send+I3Bio ...
BRANCH_GOTS: __got:0000000000010008 qword_10008    dq 400000008D4FD0h ; DATA XREF: sub_C006+r
BRANCH_GOTS: __got:0000000000010010 qword_10010    dq 400000008D5E30h ; DATA XREF: sub_C00C+r
BRANCH_GOTS: __got:0000000000010018 qword_10018    dq 400000008D4B80h ; DATA XREF: sub_C012+r
```

The values in the table (e.g. 400000008D4D40) are actually tagged offsets into BootKernelExtensions.kc. These are effectively "imported" symbols from the kernel that will be resolved once AuxiliaryKernelExtensions.kc is loaded in memory. Unfortunately these imports don't behave like imports in a normal Mach-O file, so IDA can't properly resolve the symbol names without doing some extra work.

IDA 7.5 SP3 provides a workaround. When loading AuxiliaryKernelExtensions.kc, SP3 will detect that the file contains pointers into the kernel and will offer to resolve them:



This leads to much cleaner analysis of the kernel stubs:

```
BRANCH_STUBS: __stubs:000000000000C000 ; ===== S U B R O U T I N E =====
BRANCH_STUBS: __stubs:000000000000C000
BRANCH_STUBS: __stubs:000000000000C000 ; Attributes: thunk
BRANCH_STUBS: __stubs:000000000000C000
BRANCH_STUBS: __stubs:000000000000C000 ; void __cdecl IOFree(void *address, vm_size_t size)
BRANCH_STUBS: __stubs:000000000000C000 _IOFree      proc near          ; CODE XREF: sub_1A2C4+22+p
BRANCH_STUBS: __stubs:000000000000C000                                jmp          cs:_IOFree_ptr   ; HPTAbstractRAIDController::ReleaseResource
BRANCH_STUBS: __stubs:000000000000C000 _IOFree      endp
BRANCH_STUBS: __stubs:000000000000C000
BRANCH_STUBS: __stubs:000000000000C006 ; ===== S U B R O U T I N E =====
BRANCH_STUBS: __stubs:000000000000C006
BRANCH_STUBS: __stubs:000000000000C006 ; Attributes: thunk
BRANCH_STUBS: __stubs:000000000000C006
BRANCH_STUBS: __stubs:000000000000C006 ; void __cdecl IOFreeAligned(void *address, vm_size_t size)
BRANCH_STUBS: __stubs:000000000000C006 _IOFreeAligned proc near        ; CODE XREF: sub_1A3E8+25+j
BRANCH_STUBS: __stubs:000000000000C006                                jmp          cs:_IOFreeAligned_ptr
BRANCH_STUBS: __stubs:000000000000C006 _IOFreeAligned endp
```

The values in \_\_got have been replaced by imported items, similar to what IDA does for regular Mach-O binaries:

```

BRANCH_GOTS: got:0000000000010000 ; Segment type: Pure data
BRANCH_GOTS: got:0000000000010000 ; Segment permissions: Read/Write
BRANCH_GOTS: got:0000000000010000 __BRANCH_GOTS__got segment byte public 'DATA' use64
BRANCH_GOTS: got:0000000000010000         assume cs: __BRANCH_GOTS__got
BRANCH_GOTS: got:0000000000010000         ;org 10000h
BRANCH_GOTS: got:0000000000010000 ; void (__cdecl *IOFree_ptr)(void *address, vm_size_t size)
BRANCH_GOTS: got:0000000000010000 _IOFree_ptr      dq offset __imp__IOFree ; DATA XREF: _IOFree@tr
BRANCH_GOTS: got:0000000000010000         ; _Message_Send+13B4o ...
BRANCH_GOTS: got:0000000000010000 ; void (__cdecl *IOFreeAligned_ptr)(void *address, vm_size_t size)
BRANCH_GOTS: got:0000000000010000 _IOFreeAligned_ptr dq offset __imp__IOFreeAligned
BRANCH_GOTS: got:0000000000010000         ; DATA XREF: _IOFreeAligned@tr
BRANCH_GOTS: got:0000000000010010 ; void (__cdecl *IOFreePageable_ptr)(void *address, vm_size_t size)
BRANCH_GOTS: got:0000000000010010 _IOFreePageable_ptr dq offset __imp__IOFreePageable
BRANCH_GOTS: got:0000000000010010         ; DATA XREF: _IOFreePageable@tr
BRANCH_GOTS: got:0000000000010018 ; void *(__cdecl *IOMalloc_ptr)(vm_size_t size)
BRANCH_GOTS: got:0000000000010018 _IOMalloc_ptr      dq offset __imp__IOMalloc
BRANCH_GOTS: got:0000000000010018         ; DATA XREF: _IOMalloc@tr

```

You can also consult the Imports view for a summary of the imported symbols:

Address	Ord Name
> com.apple.kernel	
> com.apple.iokit.IOPCIFamily	
000000000166A60	IOPCIDevice::extendedConfigRead8(ulong long)
000000000166A68	IOPCIDevice::extendedConfigRead16(ulong long)
000000000166A70	IOPCIDevice::extendedConfigRead32(ulong long)
000000000166A78	IOPCIDevice::extendedConfigWrite8(ulong long,uchar)
000000000166A80	IOPCIDevice::extendedConfigWrite16(ulong long,ushort)
000000000166A88	IOPCIDevice::extendedConfigWrite32(ulong long,uint)
000000000166A90	IOPCIDevice::metaClass
000000000166A98	IOPCIBridge::metaClass
> com.apple.iokit.IOOSCSIParallelFamily	
> com.apple.iokit.IOSTorageFamily	
> com.apple.iokit.IOGraphicsFamily	
000000000167260	IOFramebuffer::IOFramebuffer(OSMetaClass const*)
000000000167268	IOFramebuffer::~IOFramebuffer()
000000000167270	IOFramebuffer::gMetaClass
000000000167278	`vtable for'IOFramebuffer
000000000167280	IOFramebuffer::free(void)
000000000167288	IOFramebuffer::setProperties(OSObject *)

Also note that kernelcache symbolication can be done automatically using the `BOOT_KC_PATH` option in `macho.cfg`, which can be useful when doing analysis in batch mode.

Now that the static analysis is sufficiently robust, let's try debugging this database.

### 3.5. Debugging KernelCollections

To debug our kernelcache database, use the same xnu debugging options as before:

KDK path  ...

UEFI symbols  ...

KEXT Debugging

Debug UEFI
 
 disabled
  KDK only
  all

Then use **Debugger>Process options** to set the following fields:

Debug application setup: xnu

NOTE: all paths must be valid on the remote computer

Application:  ...

Input file:  ...

Parameters:

Hostname:  Port:

IDA will identify the load address of the input file after attaching, and rebase the database accordingly:

```

FFFFFFFF8000AFC9C0: detected Darwin Kernel Version 20.1.0: Thu Sep 24 20:22:06 PDT 2020; root:xnu-7195.40.89.100.3-1/RELEASE_ARM_T8040
XNU platform version: macOS 11.0.0
found kernelcache: /Users/troy/macOS11.kdk/BootKernelExtensions.kc
found kernelcache: /Users/troy/macOS11.kdk/SystemKernelExtensions.kc
found kernelcache: /Users/troy/macOS11.kdk/AuxiliaryKernelExtensions.kc
FFFFFFFF80011ED840: gLoadedKextSummaries
FFFFFFFF80009AB6C0: bpt hook for _OSKextLoadedKextSummariesUpdated: 1
FFFFFFFF8000200000: loaded BootKernelExtensions.kc
FFFFFFFF7F80C11000: loaded SystemKernelExtensions.kc
Rebasing program to 0xFFFFFFFF7FA16AA000... ←
FFFFFFFF7FA16B2000: loaded AuxiliaryKernelExtensions.kc

```

From the modules window we can see that although the entire AuxiliaryKernelExtensions.kc file was loaded into memory, only two of the KEXT subfiles are currently active:

Path	Base	Size
com.apple.iokit.IOUserEthernet	FFFFFFFF7F9FF59000	0000000000003000
com.apple.driver.LuaHardwareAccess	FFFFFFFF7F9FF73000	0000000000003000
com.apple.kext.OSvKernDSPLib	FFFFFFFF7FA0598000	000000000000F000
com.apple.filesystems.autofs	FFFFFFFF7FA0746000	0000000000008000
com.apple.kext.triggers	FFFFFFFF7FA086F000	0000000000003000
AuxiliaryKernelExtensions.kc	FFFFFFFF7FA16B2000	0000000000164000
com.vmware.kext.vmhgfs	FFFFFFFF7FA16C2000	000000000000B000
com.vmware.kext.VMwareGfx	FFFFFFFF7FA17AF000	0000000000007000
BootKernelExtensions.kc	FFFFFFFF8000200000	000000000331A000
com.apple.kernel	FFFFFFFF8000210000	0000000000A49000
com.apple.nke.applicationfirewall	FFFFFFFF8001268000	0000000000099000
com.apple.driver.AppleACPIPlatform	FFFFFFFF8001272000	0000000000077000

Hopefully we've shown that IDA allows you to debug a macOS11 KernelCollection almost as easy as any other Mach-O file. Also note that this same approach can be used to debug the "Boot" and "System" KernelCollections.

### 3.6. macOS11 + DWARF

So far Apple has published Kernel Debug Kits for the macOS11 beta versions, but they are of limited use.

Apparently when the mach kernel binary is inserted into BootKernelExtensions.kc, its segment mappings are modified. This means that the running kernel image in memory will differ from the kernel/DWARF binaries shipped with Apple's KDK. This will inevitably cause problems during debugging/analysis.

Apple seems to be doing something nonsensical here, so DWARF-enabled debugging with macOS11 is currently a TODO. Hopefully the situation will be resolved once macOS11 is officially released, and we will update this writeup.

## 4. UEFI Debugging

It is possible to debug the EFI firmware of a VMware Fusion guest. This section discusses some interesting examples.



## 4.1. Debugging the OSX Bootloader

Firmware debugging gives us the unique opportunity to debug the OSX bootloader. Here's how it can be easily done in IDA:

1. First copy the bootloader executable from your VM:

```
$ scp user@vm:/System/Library/CoreServices/boot.efi .
```

2. Now shut down the VM and add this line to the .vmx file:

```
monitor.debugOnStartGuest64 = "TRUE"
```

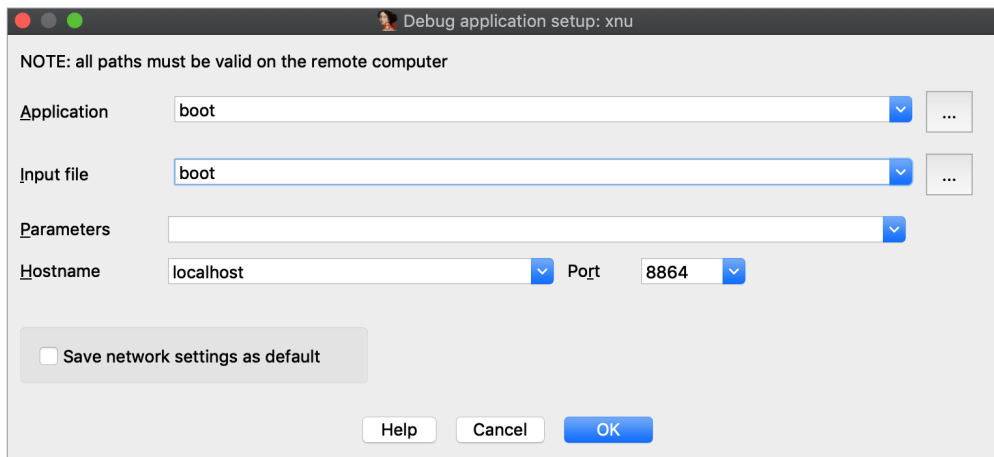
3. Load the boot.efi binary in IDA, open **Debugger>Debugger options**, check **Suspend on library load/unload**, and set **Event condition** to:

```
get_event_id() == LIB_LOADED && get_event_module_name() == "boot"
```

This will suspend the OS just before the bootloader entry point is invoked.

**Note:** For some older versions of OSX, the bootloader will be named "boot.sys". You can check the name under the **.debug** section of the executable.

4. Now select **Remote XNU Debugger** from the Debugger menu, and set the following fields in **Debugger>Process options**:



5. We're now ready to start debugging the bootloader.

Power on the VM (note that the VM is unresponsive since it is suspended), and attach to it with **Debugger>Attach to process**. After attaching IDA will try to detect the **EFI\_BOOT\_SERVICES** table. You should see the debugger print something like this to the console:

```
7FFD7430: EFI_BOOT_SERVICES
```

6. Now resume the process.

You should see many UEFI drivers being loaded, until eventually boot.efi is loaded and IDA suspends the process:

The screenshot shows the IDA View-RIP window with assembly code for the DxeCore module. The code starts with a `StartImage:` label and a `push rbp` instruction. The `Modules` window on the right lists various modules loaded, including `boot`, `apfs`, `FirmwareUpdate`, `TestDriver`, `VgaMiniPort`, `VgaClassDxe`, `Mtftp6Dxe`, `Dhcp6Dxe`, `Udp6Dxe`, `Ip6Dxe`, `UefiPxeBcDxe`, and `Mtftp4Dxe`.

```

DxeCore:.text:000000007FFC3448 db 0C9h ; E
DxeCore:.text:000000007FFC3449 db 0C3h ; A
DxeCore:.text:000000007FFC344A ; -----
DxeCore:.text:000000007FFC344A StartImage:
RIP DxeCore:.text:000000007FFC344A push rbp
DxeCore:.text:000000007FFC344B mov rbp, rsp
DxeCore:.text:000000007FFC344E push r15
DxeCore:.text:000000007FFC3450 push r14
DxeCore:.text:000000007FFC3452 push r13
DxeCore:.text:000000007FFC3454 mov r13, rcx
DxeCore:.text:000000007FFC3457 push r12
DxeCore:.text:000000007FFC3459 mov r12, rdx
DxeCore:.text:000000007FFC345C push rdi
DxeCore:.text:000000007FFC345D mov rdi, r8
DxeCore:.text:000000007FFC3460 push rsi
DxeCore:.text:000000007FFC3461 mov rsi, 800000000000002h
DxeCore:.text:000000007FFC346B push rbx
DxeCore:.text:000000007FFC346C sub rsp, 28h
DxeCore:.text:000000007FFC3470 call near ptr unk_7FFC2AEB
DxeCore:.text:000000007FFC3475 test rax, rax
UNKNOWN 000000007FFC344A: DxeCore: (Synchronized with RIP)
  
```

Output window:

```

7EA77000: loaded Dhcp6Dxe
7EA6B000: loaded Mtftp6Dxe
7EA65000: loaded VgaClassDxe
7EA60000: loaded VgaMiniPort
7EA48000: loaded TestDriver
7EA43000: loaded FirmwareUpdate
7E9B9000: loaded apfs
7FF78B98: EFI_RUNTIME_SERVICES
Rebasing program to 0x000000007E89F000...
7E89E000: loaded boot
  
```

7. At this point, the bootloader entry point function is about to be invoked.

Jump to `_ModuleEntryPoint` in `boot.efi` and press F4. We can now step through `boot.efi`:

The screenshot shows the IDA View-RIP window with the `ModuleEntryPoint` function. The function starts with a `push rbp` instruction. The `General registers` window on the right shows the state of the registers, including `RAX`, `RBX`, `RCX`, `RDY`, `RSI`, `RDI`, `RBP`, `RSP`, `R8`, `R9`, `R10`, `R11`, `R12`, `R13`, `R14`, `R15`, `RIP`, and `EFL`.

```

.text:000000007E8A3818 ModuleEntryPoint proc near
.text:000000007E8A3818
.text:000000007E8A3818 var_160= qword ptr -160h
.text:000000007E8A3818 var_150= qword ptr -150h
.text:000000007E8A3818 var_D0= byte ptr -0D0h
.text:000000007E8A3818 var_80= qword ptr -80h
.text:000000007E8A3818 var_78= qword ptr -78h
.text:000000007E8A3818 var_70= qword ptr -70h
.text:000000007E8A3818 var_68= qword ptr -68h
.text:000000007E8A3818 var_60= qword ptr -60h
.text:000000007E8A3818 var_58= qword ptr -58h
.text:000000007E8A3818 var_50= qword ptr -50h
.text:000000007E8A3818 var_48= dword ptr -48h
.text:000000007E8A3818 var_40= qword ptr -40h
.text:000000007E8A3818 var_38= dword ptr -38h
.text:000000007E8A3818 var_34= word ptr -34h
RIP .text:000000007E8A3818 push rbp
.text:000000007E8A3819 mov rbp, rsp
.text:000000007E8A381C push r15
.text:000000007E8A381E push r14
.text:000000007E8A3820 push r12
.text:000000007E8A3822 push rsi
.text:000000007E8A3823 push rdi
  
```

General registers:

```

RAX 8000000000000003 MEMORY:8000000000000003 ID 0
RBX 000000007FC0D018 FIRMWARE:00000007FC0D018 VIP 0
RCX 000000007FC21498 FIRMWARE:00000007FC21498 AC 0
RDY 000000007FF78018 FIRMWARE:SystemTable VM 0
RSI 000000007FFD7960 "ldri" RF 0
RDI 000000007FFBFD70 FIRMWARE:ConfigurationTable+470D8 NT 0
RBP 000000007FFBFA80 FIRMWARE:ConfigurationTable+46E18 OF 0
RSP 000000007FFBFA48 FIRMWARE:ConfigurationTable+46DB0 DF 0
R8 0000000000000004 FIRMWARE:0000000000000004 IP 1
R9 000000007FC13418 FIRMWARE:00000007FC13418 TF 0
R10 0000000000000000 FIRMWARE:0000000000000000 SF 0
R11 0000000000000009 FIRMWARE:0000000000000009 ZF 1
R12 000000007FFBFD78 FIRMWARE:ConfigurationTable+470E0 AF 0
R13 000000007FC21498 FIRMWARE:00000007FC21498 PF 1
R14 0000000000000000 FIRMWARE:0000000000000000 CF 0
R15 0000000000000153 FIRMWARE:0000000000000153
RIP 000000007E8A3818 _ModuleEntryPoint
EFL 00000246
  
```

## 4.2. GetMemoryMap

To facilitate UEFI debugging, IDA provides an IDC helper function: `xnu_get_efi_memory_map`.

This function will invoke the `GetMemoryMap` function in the `EFI_BOOT_SERVICES` table and return an array of `EFI_MEMORY_DESCRIPTOR` objects:

```

IDC>extern map;
IDC>map = xnu_get_efi_memory_map();
IDC>map.size
    35.    23h    43o
IDC>map[27]
object
  Attribute: 0x80000000000000Fi64
  NumberOfPages: 0x20i64
  PhysicalStart: 0x7FF09000i64
  Type: "EfiRuntimeServicesCode"
  VirtualStart: 0x0i64
  
```

This function can be invoked at any point during firmware debugging.

## 4.3. UEFI Debugging + DWARF

If you build your own EFI apps or drivers on OSX, you can use IDA to debug the source code.

In this example we will debug a sample EFI application. On OSX the convention is to build EFI apps in the Mach-O format, then convert the file to PE .efi with the **mtoc** utility. In this example, assume we have an EFI build on our OSX virtual machine that contains the following files in the ~/TestApp directory:

- **TestApp.efi** - the EFI application that will be run
- **TestApp.dll** - the original Mach-O binary
- **TestApp.dll.dSYM** - DWARF info for the app
- **TestApp.c** - source code for the app

Here's how we can debug this application in IDA:

1. On your host machine, create a directory that will mirror the directory on the VM:

```
mkdir ~/TestApp
```

2. Copy the efi, macho, dSYM, and c files from your VM:

```
scp -r vmuser@vm:TestApp/TestApp.* ~/TestApp
```

3. Open the TestApp.efi binary in IDA, and wait for IDA to analyze it.

Note that you can improve the disassembly by loading the DWARF file from TestApp.dll.dSYM. You can do this with **Edit>Plugins>Load DWARF file**, or you can load it programatically from IDAPython:

```
path = "~/TestApp/TestApp.dll.dSYM/Contents/Resources/DWARF/TestApp.dll"
node = idaapi.netnode()
node.create("$ dwarf_params")
node.supset(1, os.path.expanduser(path))
idaapi.load_and_run_plugin("dwarf", 3)
```

4. Select **Remote XNU Debugger** from the debugger menu, and set the following fields in **Debugger>Process options**:

Application	TestApp	...
Input file	TestApp	...
Parameters		
Hostname	localhost	Port 8864

5. In **Debugger>Debugger options**, enable **Suspend on library load/unload** and set the **Event condition** field to:

```
get_event_id() == LIB_LOADED && get_event_module_name() == "TestApp"
```

6. In **Debugger>Debugger options>Set specific options**, set the following fields:

UEFI symbols	~/TestApp	...
KEXT Debugging		
<input checked="" type="checkbox"/> Debug UEFI	<input checked="" type="radio"/> disabled	<input type="radio"/> KDK only
	<input type="radio"/> all	

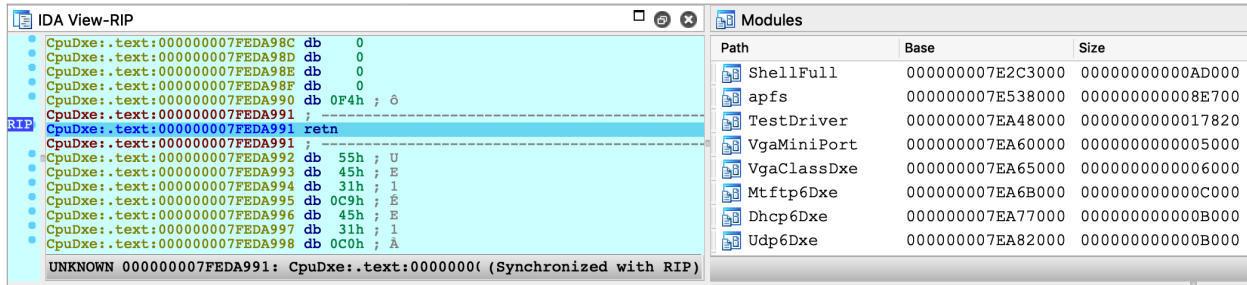
Note that we must enable the **Debug UEFI** option, and set the **UEFI symbols** option so the debugger can find DWARF info for the EFI app at runtime.

7. If the usernames on the host and VM are different, we will need a source path mapping:



```
idaapi.add_path_mapping("/Users/vmuser/TestApp", "/Users/hostuser/TestApp")
```

- Reboot the VM and enter the EFI Shell
- Attach the debugger. After attaching IDA will detect the firmware images that have already been loaded:

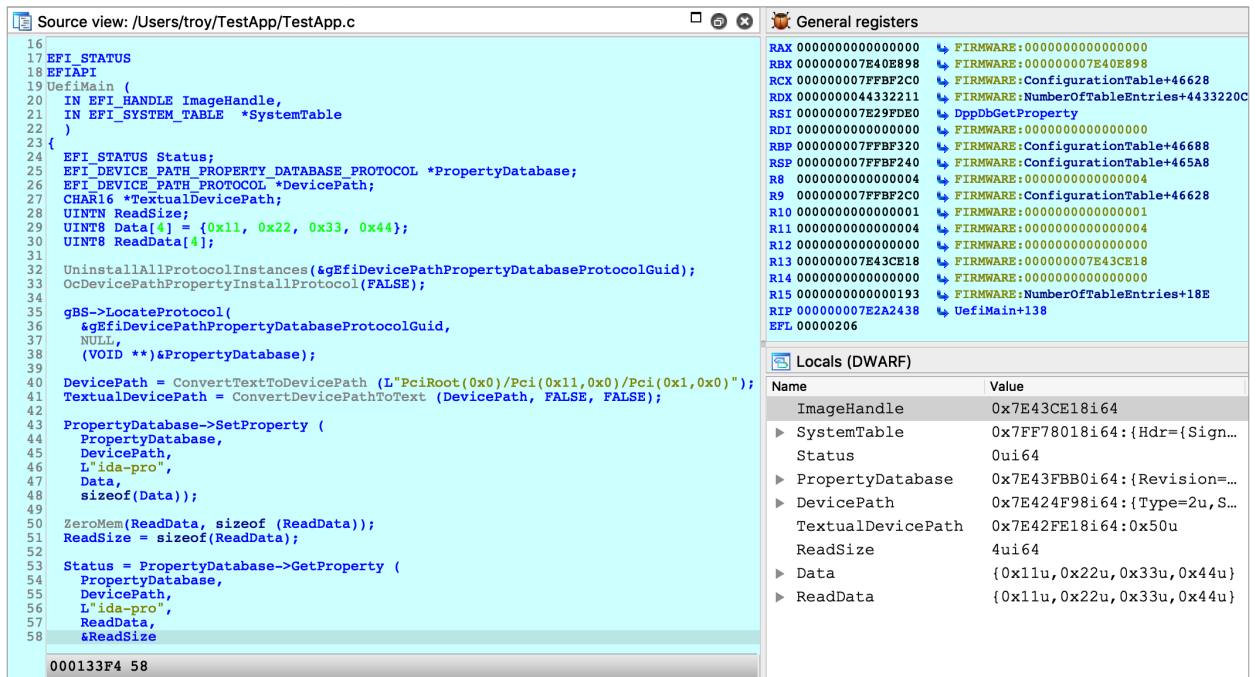


- Resume the OS and launch TestApp from the EFI Shell prompt:

```
Shell>fs1:\Users\vmuser\TestApp\TestApp.efi
```

At this point IDA will detect that the target app has been loaded, and suspend the process just before the entry point of TestApp.efi (because of step 5).

- Now we can set a breakpoint somewhere in TestApp.efi and resume the OS. The debugger will be able to load source file and local variable information from TestApp.dll.dSYM:



**IMPORTANT NOTE:** You must wait until TestApp has been loaded into memory before setting any breakpoints. If you add a breakpoint in the database before attaching the debugger, IDA might not set the breakpoint at the correct address. This is a limitation in IDA that we must work around for now.

## 5. Debugging iOS with Corellium

IDA can also debug the iOS kernel, provided you have access to a virtual iOS device from Corellium.

### 5.1. Quick Start

To get started with debugging iOS, we will perform a simple experiment to patch kernel memory.

The device used in this example is a virtual iPhone XS with iOS 12.1.4, but it should work with any model or iOS version

that Corellium supports. Begin by powering on your device and allow it to boot up. In the Corellium UI, look for the line labeled **SSH** under **Advanced options**:

```
Advanced Options
To connect to an SSH daemon running on the device over USB/lockdown:
SSH ssh root@10.11.1.3
```

Ensure you can connect to the device by running this command over ssh:

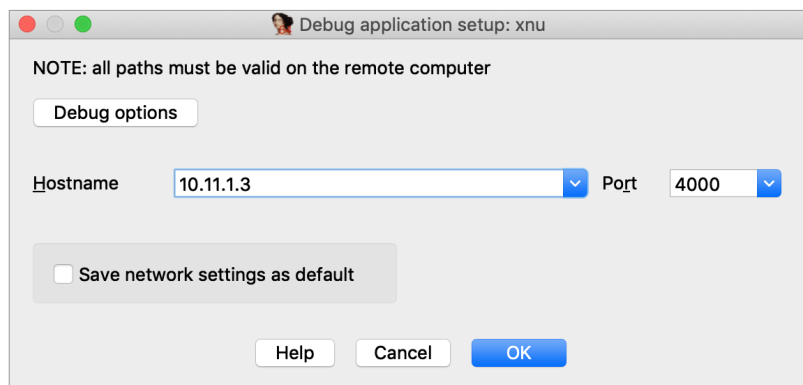
```
$ ssh root@10.11.1.3 uname -v
Darwin Kernel Version 18.2.0 ... root:xnu-4903.242.2~1/RELEASE_ARM64_T8020
```

We will use IDA to patch this version string.


Now launch IDA, and when prompted with the window **IDA: Quick start**, choose **Go** to start with an empty database and open **Debugger>Attach>Remote XNU Debugger**. In the Corellium UI, find the hostname:port used by the kernel GDB stub. It should be specified in the line labeled **kernel gdb**:

```
To attach to the iOS kernel (download here) using a debugger with the gdb-remote protocol:
kernel gdb lldb --one-line "gdb-remote 10.11.1.3:4000"
```

And set the **Hostname** and **Port** fields in IDA's application setup window:



Now click on **Debug options>Set specific options**, and for the **Configuration** dropdown menu, be sure to select **Corellium-ARM64**:



You can ignore the other config options for now, and click OK.

Click OK again, and wait for IDA to establish a connection to Corellium's GDB stub (this may take a few seconds). Then select **<attach to the process started on target>** and wait for IDA to attach. This might take several seconds (we will address this later), but for now simply wait for IDA to perform the initial setup.

If IDA could detect the kernel, it should appear in the Modules list:

IDA View-PC

```

TEXT_EXEC: text:FFFFFFF0079ED031 DCB 0xF
TEXT_EXEC: text:FFFFFFF0079ED032 DCB 0x5F ;
TEXT_EXEC: text:FFFFFFF0079ED033 DCB 0xD6 ;
TEXT_EXEC: text:FFFFFFF0079ED034 DCB 0x7F ;
TEXT_EXEC: text:FFFFFFF0079ED035 DCB 0x20
TEXT_EXEC: text:FFFFFFF0079ED036 DCB 3
TEXT_EXEC: text:FFFFFFF0079ED037 DCB 0xD5
TEXT_EXEC: text:FFFFFFF0079ED038 ;
TEXT_EXEC: text:FFFFFFF0079ED038 CBZ X30, loc_FFFFFFFF0079ED040
TEXT_EXEC: text:FFFFFFF0079ED03C RET
TEXT_EXEC: text:FFFFFFF0079ED040 ;
TEXT_EXEC: text:FFFFFFF0079ED040 loc_FFFFFFFF0079ED040 ;
TEXT_EXEC: text:FFFFFFF0079ED040 MOV X0, #1
TEXT_EXEC: text:FFFFFFF0079ED044 BL unk_FFFFFFFF007B963B0
TEXT_EXEC: text:FFFFFFF0079ED048 MOV X0, #0
TEXT_EXEC: text:FFFFFFF0079ED04C BL unk_FFFFFFFF007B8179C
TEXT_EXEC: text:FFFFFFF0079ED050 ;
TEXT_EXEC: text:FFFFFFF0079ED050 loc_FFFFFFFF0079ED050 ;
TEXT_EXEC: text:FFFFFFF0079ED050 B loc_FFFFFFFF0079ED050 ;
TEXT_EXEC: text:FFFFFFF0079ED054 DCB 0x1F
TEXT_EXEC: text:FFFFFFF0079ED055 DCB 0x20
TEXT_EXEC: text:FFFFFFF0079ED056 DCB 3
TEXT_EXEC: text:FFFFFFF0079ED057 DCB 0xD5
TEXT_EXEC: text:FFFFFFF0079ED058 DCB 0xFD
TEXT_EXEC: text:FFFFFFF0079ED059 DCB 0x7B ; {
TEXT_EXEC: text:FFFFFFF0079ED05A DCB 0xBF
UNKNOWN_FFFFFFFF0079ED038: __TEXT_EXEC: __text:FFFF (Synchronized with PC)

```

General Registers

X18	0000000000000000	MEMORY:0000000000000000	N	0
X19	FFFFFFF008F85000	__DATA:__data:FFFFFFF008F85000	Z	1
X20	FFFFFFF00625A90	MEMORY:FFFFFFF00625A90	C	1
X21	0000008DB723C252	MEMORY:0000008DB723C252	V	0
X22	FFFFFFF009121220	__DATA:__common:FFFFFFF009121220	Q	0
X23	0000000000000001	MEMORY:0000000000000001	IT2	0
X24	FFFFFFF00913D000	__DATA:__common:FFFFFFF00913D000	J	0
X25	FFFFFFF0091208CC	__DATA:__common:FFFFFFF0091208CC	GE	0
X26	00000000FFFFFFF	MEMORY:00000000FFFFFFF	IT	0
X27	FFFFFFF0090F8000	__DATA:__common:FFFFFFF0090F8000	E	1
X28	FFFFFFF00625B10	MEMORY:FFFFFFF00625B10	A	1
X29	FFFFFFF06BD77FF0	MEMORY:FFFFFFF06BD77FF0	I	1
X30	FFFFFFF007B8178C	__TEXT_EXEC: __text:FFFFFFF007B8178C	F	1
SP	FFFFFFF06BD77FC0	MEMORY:FFFFFFF06BD77FC0	T	0
PC	FFFFFFF0079ED038	__TEXT_EXEC: __text:FFFFFFF0079ED038	MODE	4
PSR	604003C4			

Modules

Path	Base	Size
kernel	FFFFFFF007004000	000000001F40000
<GDB remote pr...		

Line 1 of 2

and the kernel version will be printed to the console:

```
FFFFFFF007029FD7: detected Darwin Kernel Version 18.2.0 ...
```

Navigate to this address and use IDAPython to overwrite the string:

```

TEXT: const:FFFFFFF007029FD6 DCB 0
TEXT: const:FFFFFFF007029FD7 DCB 0x44 ; D
TEXT: const:FFFFFFF007029FD8 DCB 0x61 ; a
TEXT: const:FFFFFFF007029FD9 DCB 0x72 ; r
TEXT: const:FFFFFFF007029FDA DCB 0x77 ; w
TEXT: const:FFFFFFF007029FDB DCB 0x69 ; i
TEXT: const:FFFFFFF007029FDC DCB 0x6E ; n
TEXT: const:FFFFFFF007029FDD DCB 0x20
TEXT: const:FFFFFFF007029FDE DCB 0x4B ; K
TEXT: const:FFFFFFF007029FDF DCB 0x65 ; e
TEXT: const:FFFFFFF007029FE0 DCB 0x72 ; r
TEXT: const:FFFFFFF007029FE1 DCB 0x6E ; n
TEXT: const:FFFFFFF007029FE2 DCB 0x65 ; e
TEXT: const:FFFFFFF007029FE3 DCB 0x6C ; l
TEXT: const:FFFFFFF007029FE4 DCB 0x20

```

```
idaapi.dbg_write_memory(0xFFFFF007029FD7, "IDAPRO".encode('utf-8'))
```

Resume the OS, and try running the same command as before:

```
$ ssh root@10.11.1.3 uname -v
IDAPRO Kernel Version 18.2.0 ... root:xnu-4903.242.2~1/RELEASE_ARM64_T8020
```

If we could successfully write to kernel memory, **IDAPRO** should appear in the output.

## 5.2. Creating a KDK for iOS

Typically a Kernel Development Kit is not available for iOS devices, but we can still utilise the KDK\_PATH option in IDA to achieve faster debugging. In the example above, the initial attach can be slow because IDA must parse the kernel image in memory (which can be especially slow if the kernel has a symbol table).

Here's how you can speed things up:

1. create the KDK directory:

```
$ mkdir ~/iPhoneKDK
```

2. copy the kernelcache from the virtual device:

```
$ scp root@ip:/System/Library/Caches/com.apple.kernelcaches/kernelcache /tmp
```

3. uncompress the kernelcache with `lzssdec`:

```
$ lzssdec -o OFF < /tmp/kernelcache > ~/iPhoneKDK/kernelcache
```

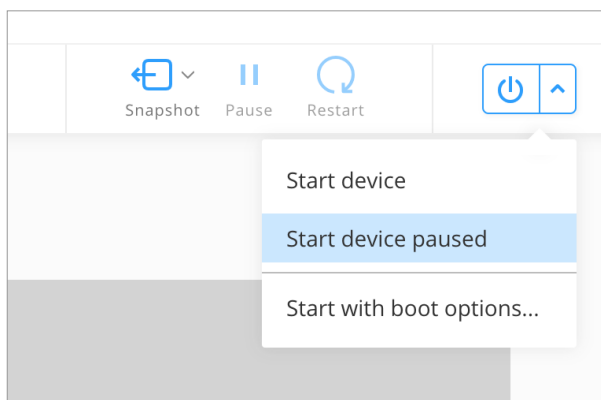
4. set `KDK_PATH` in `dbg_xnu.cfg`:

```
KDK_PATH = "~/iPhoneKDK";
```

Now whenever the debugger must extract information from the kernel, it will parse the local file on disk. This should be noticeably faster, especially if the device is hosted by Corellium's web service.

### 5.3. Debugging the iOS Kernel Entry Point

Corellium allows us to debug the first few instructions of kernel initialization. This can be very useful if we want to gain control of the OS as early as possible. In the Corellium UI, power on your device with the **Start device paused** option:



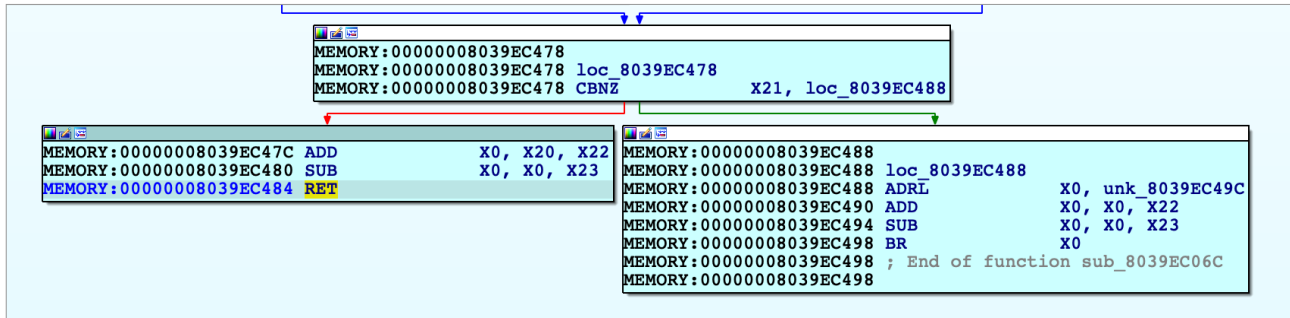
Now start IDA with an empty database and attach to the suspended VM:

```
PC MEMORY:0000008039E818F DCB 0
MEMORY:0000008039E8190 ; -----
MEMORY:0000008039E8190 B loc_8039EC06C
MEMORY:0000008039E8190 ; -----
MEMORY:0000008039E8194 DCB 0x1F
MEMORY:0000008039E8195 DCB 0x20
MEMORY:0000008039E8196 DCB 3
MEMORY:0000008039E8197 DCB 0xD5
MEMORY:0000008039E8198 DCB 0x1F
MEMORY:0000008039E8199 DCB 0x20
MEMORY:0000008039E819A DCB 3
```

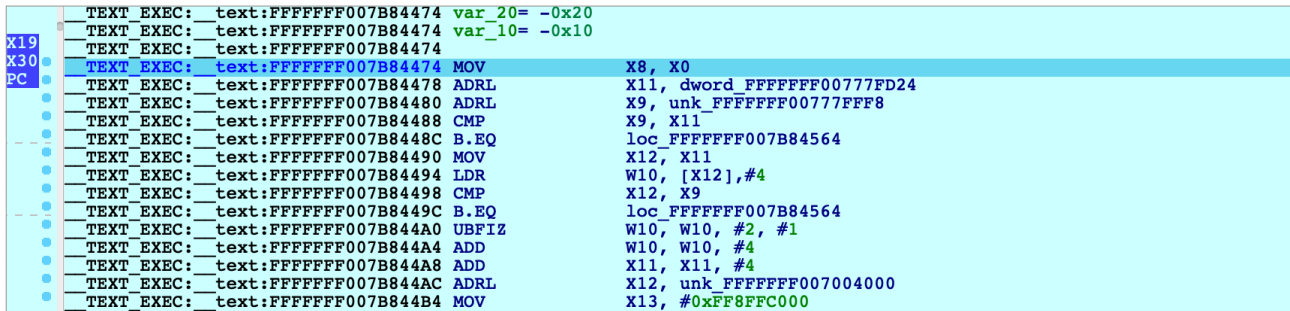
From the XNU source, this is likely the `_start` symbol in `osfmk/arm64/start.s`, which simply branches to `start_first_cpu`. After stepping over this branch:

```
PC MEMORY:0000008039EC06B DCB 0x14
MEMORY:0000008039EC06C ; -----
MEMORY:0000008039EC06C loc_8039EC06C ; CODE XREF: MEMORY:0000008039E8190↑j
MEMORY:0000008039EC06C MSR #0, c1, c0, #4
MEMORY:0000008039EC070 MSR #6, #0xF
MEMORY:0000008039EC074 MOV X20, X0
MEMORY:0000008039EC078 MOV X21, #0
MEMORY:0000008039EC07C ADRL X0, unk_8039E9000
MEMORY:0000008039EC084 MSR #0, c12, c0, #0, X0
MEMORY:0000008039EC088 MOV X1, X30
MEMORY:0000008039EC08C MOVK X0, #0x4455, LSL#48
MEMORY:0000008039EC090 MOVK X0, #0x4455, LSL#32
MEMORY:0000008039EC094 MOVK X0, #0x6466, LSL#16
MEMORY:0000008039EC098 MOVK X0, #0x6677
MEMORY:0000008039EC09C BL unk_80411903C
```

Press shortcut **P** to analyze `start_first_cpu`. This is where the kernel performs its initial setup (note that the value in `X0` is a pointer to the `boot_args` structure). This function is interesting because it is responsible for switching the kernel to 64-bit virtual addressing. Typically the switch happens when this function sets `X30` to a virtual address, then performs a `RET`:



Use F4 to run to this RET instruction. In this example X30 will now point to virtual address **0xFFFFFFFF007B84474**. After single stepping once more, we end up in **arm\_init** in virtual memory:



After this single step, IDA detected that execution reached the kernel's virtual address space and automatically initialized the debugging environment. In this case a message will be printed to the console:

```
FFFFFFFF007004000: process kernel has started
```

This signifies that IDA successfully detected the kernel base and created a new module in the Modules list. If the kernel has a symbol table, debug names will be available. Also note that PC now points inside the segment **\_\_TEXT\_EXEC:\_\_text** instead of **MEMORY**, because the debugger parsed the kernel's load commands to generate proper debug segments.

Now that we know the address of **arm\_init**, we can streamline this task:

1. power on the device with **Start device paused**
2. attach to the paused VM
3. set a hardware breakpoint at **arm\_init**:

```
idaapi.add_bpt(0xFFFFFFFF007B84474, 1, BPT_EXEC)
```

4. resume, and wait for the breakpoint to be hit

This gives us a quick way to break at the first instruction executed in virtual memory. You can continue debugging iOS as normal.

## 6. Known Issues and Limitations

Here is a list of known shortcomings in the XNU Debugger. Eventually we will address all of them, but it is unlikely we will resolve all of them by the next release. If any of the following topics are important to you, please let us know by sending an email to [support@hex-rays.com](mailto:support@hex-rays.com). Issues with vocal support from our users are automatically prioritised.

### 6.1. iBoot debugging

Debugging the iOS firmware/bootloader is not yet supported. An On-Premise Corellium box is required for this functionality, so we will only implement it if there is significant demand.

## 6.2. 32-bit XNU

The XNU Debugger does not support debugging 32-bit XNU. Since pure 32-bit OSes are quite outdated it is unlikely we will support them unless there is exceptional demand.

## 6.3. KDP

The XNU Debugger relies on the Remote GDB Protocol, and currently Apple's Kernel Debugging Protocol (KDP) is not supported. It is possible to add KDP support to IDA in the future.