

Understanding Syscalls: Direct, Indirect, and Cobalt Strike Implementation

By Mohamed Adel

Published: 2023-08-18 · Archived: 2026-04-05 17:01:27 UTC

In case images fail to load, it might be due to jsDelivr CDN ban in Egypt. To resolve this, consider using a VPN. :)

- To Bypass user-mood hooks. why?
 - For Hiding a code inside a legitimate process (Process Injection)
 - Avoiding EDR alerts!

Hooking user-mode functions by placing a jump to another code section. EDRs use hooks to check the function parameters. For example, if you are trying to change the memory protections of some data to add executable protections. This is a very suspicious activity so EDRs will be alert to that. Most Hooks are on the lowest level of the user-mode interface in **ntdll.dll** which are the system calls.

Windows has a defined schema of how `syscalls` are used. Most of the documented windows APIs are just a wrapper of a lower-level Functions in `ntdll.dll` which are compiled to a `syscall` with the right SSN (System Service Number). To look at how `Nt*` version of the higher-level API is implemented.

```

1      0:018> uf NtOpenProcess
2      ntdll!NtOpenProcess:
3      00007ffa`4874d4c0 4c8bd1      mov     r10,rcx
4      00007ffa`4874d4c3 b826000000    mov     eax,26h
5      00007ffa`4874d4c8 f604250803fe7f01 test   byte ptr [SharedUserData+0x308 (00000000`7ffe0308)],1
6      00007ffa`4874d4d0 7503          jne     ntdll!NtOpenProcess+0x15 (00007ffa`4874d4d5) Branch
7
8      ntdll!NtOpenProcess+0x12:
9      00007ffa`4874d4d2 0f05          syscall
10     00007ffa`4874d4d4 c3           ret
11
12     ntdll!NtOpenProcess+0x15:
13     00007ffa`4874d4d5 cd2e          int     2Eh
14     00007ffa`4874d4d7 c3           ret
    
```

At address `00007ffa~4874d4d2` there `syscall` instruction. This instruction transfers the execution to the system-handler at the kernel. The handler is specified using pre-defined SSN number loaded into `EAX` Register (In this case `EAX = 0x26` at address `00007ffa~4874d4c3`). So, to make a `syscall` The SSN associated. The code stub of the `syscalls` is simple.

```

1      mov r10, rcx
2      mov eax, <syscall_number>
3      syscall
4      ret
    
```

Now, the missing thing is the `syscall_number` . These numbers are changing based on the Build version of windows. There are some techniques to get these numbers.

1. SysWhispers

[SysWhispers](#) That generate the table of these numbers in the form of a header file and assembly file that can be embedded in the code. The generated code contains `syscall` number for multiple versions, The right windows build version is detected at runtime using PEB structure.

<pre> 1 2 3 4 5 </pre>	<pre> ... +0x118 OSMajorVersion : Uint4B +0x11c OSMinorVersion : Uint4B +0x120 OSBuildNumber : Uint2B ... </pre>
------------------------	---

The assembly code generated (Full document at [example-output](#))

<pre> 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 </pre>	<pre> ... NtOpenProcess PROC mov rax, gs:[60h] ; Load PEB into RAX. NtOpenProcess_Check_X_X_XXXX: ; Check major version. cmp dword ptr [rax+118h], 5 je NtOpenProcess_SystemCall_5_X_XXXX cmp dword ptr [rax+118h], 6 je NtOpenProcess_Check_6_X_XXXX cmp dword ptr [rax+118h], 10 je NtOpenProcess_Check_10_0_XXXX jmp NtOpenProcess_SystemCall_Unknown NtOpenProcess_Check_6_X_XXXX: ; Check minor version for Windows Vista/7/8. cmp dword ptr [rax+11ch], 0 je NtOpenProcess_Check_6_0_XXXX cmp dword ptr [rax+11ch], 1 je NtOpenProcess_Check_6_1_XXXX cmp dword ptr [rax+11ch], 2 je NtOpenProcess_SystemCall_6_2_XXXX cmp dword ptr [rax+11ch], 2 je NtOpenProcess_SystemCall_6_3_XXXX jmp NtOpenProcess_SystemCall_Unknown NtOpenProcess_Check_6_0_XXXX: ; Check build number for Windows Vista. cmp dword ptr [rax+120h], 6000 je NtOpenProcess_SystemCall_6_0_6000 cmp dword ptr [rax+120h], 6001 je NtOpenProcess_SystemCall_6_0_6001 cmp dword ptr [rax+120h], 6002 je NtOpenProcess_SystemCall_6_0_6002 jmp NtOpenProcess_SystemCall_Unknown NtOpenProcess_Check_6_1_XXXX: ; Check build number for Windows 7. cmp dword ptr [rax+120h], 7600 je NtOpenProcess_SystemCall_6_1_7600 cmp dword ptr [rax+120h], 7601 je NtOpenProcess_SystemCall_6_1_7601 jmp NtOpenProcess_SystemCall_Unknown NtOpenProcess_Check_10_0_XXXX: ; Check build number for Windows 10. cmp dword ptr [rax+120h], 10240 je NtOpenProcess_SystemCall_10_0_10240 cmp dword ptr [rax+120h], 10586 je NtOpenProcess_SystemCall_10_0_10586 ... </pre>
--	---

- SSN code stub** This technique doesn't look for SSN number, instead it gets the code stub of the required API. This can be done by opening the PE file and parsing the Export table of `ntdll`.
- Extract SSN** It extracts the SSN from `ntdll` by parsing the Export table. The difference between it and the previous one is that it only extracts the `syscall` number. Both methods load `ntdll.dll` from the disk first using `win32 API OpenFile` which might be hooked. [hell's gate](#) for more.
- Syscalls' number sequence** This method takes advantage of the SSNs are in a sequence for example if a syscall number is 0x26 the following will be 0x27 and so on. This relies also on the fact that not all the system calls are hooked! So, to get the SSN of a function, you need to find the nearest unhooked syscall. This was presented by [halos gate](#). But this is not valid in newer versions of Windows as the SSNs sequence is no longer valid.

4. **Parallel loading** This is an interesting technique explained in this [blog](#). It uses windows feature introduced in windows 10 to load DLLs through multiple threads instead of one in older versions of windows. It was found that the syscall stub of native Functions

`NtOpenFile()`, `NtCreateSection()`, `ZwQueryAttributeFile()`, `ZwOpenSection()` and `ZwMapViewOfFile()` -

There is a lot of things happens between the two actions, detailed explanation in the previously mentioned [blog](#) -are copied into `LdrpThunkSignature` array. This is done to check the integrity of the functions' code. These APIs' syscall numbers can be used to load a new version of ntdll.dll from the disk and avoid any user-mood hooks.

5. **Sorting by system call address** This technique uses the relation between the address of the system call stub and the SSN. It is known as [FreshyCalls](#). In simple words, it walks the Export Address Table of `ntdll` and saves the Name -or a hash of the name- and Address of each entry in a table. Then, it sorts the entries by the addresses in ascending order. It was found that the first function `NtAccessCheck` (by address) has an SSN = 0

```

1      0:007> uf NtAccessCheck
2      ntdll!NtAccessCheck:
3      00007ffa`4874d000 4c8bd1      mov     r10,rcx
4      00007ffa`4874d003 b800000000      mov     eax,0
5      00007ffa`4874d008 f604250803fe7f01 test   byte ptr [SharedUserData+0x308 (00000000`7ffe0308)],1
6      00007ffa`4874d010 7503      jne    ntdll!NtAccessCheck+0x15 (00007ffa`4874d015) Branch
7
8      ntdll!NtAccessCheck+0x12:
9      00007ffa`4874d012 0f05      syscall
10     00007ffa`4874d014 c3      ret
11
12     ntdll!NtAccessCheck+0x15:
13     00007ffa`4874d015 cd2e      int    2Eh
14     00007ffa`4874d017 c3      ret

```

and if we unassembled the next function by adding one to the last address (as ret opcode is one byte) we will get that the next function's SSN is 1!

```

1      0:007> uf 00007ffa`4874d017+1
2      ntdll!NtAccessCheck+0x18:
3      00007ffa`4874d018 0f1f840000000000 nop    dword ptr [rax+rax]
4      00007ffa`4874d020 4c8bd1      mov     r10,rcx
5      00007ffa`4874d023 b801000000      mov     eax,1
6      00007ffa`4874d028 f604250803fe7f01 test   byte ptr [SharedUserData+0x308 (00000000`7ffe0308)],1
7      00007ffa`4874d030 7503      jne    ntdll!NtWorkerFactoryWorkerReady+0x15 (00007ffa`4874d035) Branch
8
9      ntdll!NtWorkerFactoryWorkerReady+0x12:
10     00007ffa`4874d032 0f05      syscall
11     00007ffa`4874d034 c3      ret
12
13     ntdll!NtWorkerFactoryWorkerReady+0x15:
14     00007ffa`4874d035 cd2e      int    2Eh
15     00007ffa`4874d037 c3      ret

```

So, by sorting the functions by the addresses, we have the SSN. for the code, look at [MDSec](#) (8. Sorting by System Call Address) blog or see [FreshlyCalls](#) implementation. The execution of the system call is not direct by calling `syscall` instruction. Instead, it uses the method explained below. Briefly, it uses the `syscall` instructions from `ntdll`.

All the methods described are workarounds to get the system call number without getting caught. `syscall` instruction reveals that some suspicious activity is going on. This is done using `KPROCESS!InstrumentationCallback` in windows.

```

1      0:030> dt _kprocess
2      ntdll!_KPROCESS
3      +0x000 Header : _DISPATCHER_HEADER
4      ...
5      +0x3d8 InstrumentationCallback : Ptr64 Void

```

6	...
7	+0x3f8 EndPadding : [8] UInt8B

Any time the windows is done with a syscall and returns to user-mode, it checks this member it is not `NULL`, the execution will be transferred to that pointer. To check if the syscall is legit, the return address after finishing the syscall is checked to see if it is not from a valid place. If the address is in the address space of the process running, it's not a legitimate place to make a syscall. This check was done by ScyllaHide to detect manual syscalls, the source code can be found [here](#).

1	<code>if (InterlockedOr(TlsGetInstrumentationCallbackDisabled(), 0x1) == 0x1)</code>
2	<code>return ReturnVal; // Do not recurse</code>
3	
4	<code>const PVOID ImageBase = NtCurrentPeb()->ImageBaseAddress;</code>
5	<code>const PIMAGE_NT_HEADERS NtHeaders = RtlImageNtHeader(ImageBase);</code>
6	<code>if (NtHeaders != nullptr && ReturnAddress >= (ULONG_PTR)ImageBase &&</code>
7	<code>ReturnAddress < (ULONG_PTR)ImageBase + NtHeaders->OptionalHeader.SizeOfImage)</code>
8	<code>{</code>
9	<code> // Syscall return address within the exe file</code>
10	<code>ReturnVal = (ULONG_PTR)(ULONG)STATUS_PORT_NOT_SET;</code>
11	
12	<code> // Uninstall ourselves after we have completed the sequence { NtQIP, NtQIP }. More NtSITs will follow but we can't do anyt</code>
13	<code>NumManualSyscalls++;</code>
14	<code>if (NumManualSyscalls >= 2)</code>
15	<code>{</code>
16	<code> InstallInstrumentationCallbackHook(NtCurrentProcess, TRUE);</code>
17	<code>}</code>
18	<code>}</code>
19	
20	<code>InterlockedAnd(TlsGetInstrumentationCallbackDisabled(), 0);</code>
21	
22	<code>return ReturnVal;</code>
23	<code>}</code>

It checks the return address of the successful system call. If it resides on the address space of the binary we are running, it is an indication of manual system call.

The Solution The solution to this hooking method is done by [Bouncy Gate](#) and [Recycled Gate](#) method. The idea is quite simple, it is an adjusted version of [Hell's Gate](#). Instead of directly executing `syscall` instruction and getting caught by static signatures and system call callbacks described above, the author replaces the `syscall` instruction with a trampoline jump (`JMP`) to a `syscall` instruction address from `ntdll.dll`. now there is no direct `syscall` instruction and the system call originated from a legitimate place `ntdll`. This is also implemented in [SysWhispers3](#). To get the address of the syscall instruction in `ntdll` we can parse the export table and search for `syscall`, `ret` opcodes `0F 05 0C` or the constant pattern of syscalls in `ntdll` can be used to get the syscall address. If the function is not hooked, the syscall instruction is on offset `0x12` from the function's address, we can verify that by comparing the opcodes.

The sample from [Dodo's blog](#) Where he already analyzed how indirect syscalls implemented in Cobalt Strike. for easy access, here is [UnpacMe Results 020b20098f808301cad6025fe7e2f93fa9f3d0cc5d3d0190f27cf0cd374bcf04](#). The sample is packed. The unpacking process is easy. Just put a breakpoint on `VirtualProtect` and get the base address (First Argument). Function `sub_18001B6B0` contains the important part, system call SSN retrieving and execution methods. You can get to this function by following the `call` instruction to `rax` which contains a `qword` memory area or a call to the `qword` directly. These locations are populated with addresses of the required APIs in this function. We can see multiple calls to `sub_18001A73C` with arguments: `qword_*`, a hash (such as `0B12B7A69h`), variable passed to the function `sub_18001A7F4` and another allocated memory which is also passed to `sub_18001A7F4`.

```

.text:000000018001B6E5      and     [rbp+1E70h+arg_0], 0
.text:000000018001B6ED      mov     edi, 1F4h
.text:000000018001B6F2      lea    r8, [rbp+1E70h+arg_0]
.text:000000018001B6F9      lea    rcx, [rsp+1F70h+var_1F40] ; void *
.text:000000018001B6FE      mov     edx, edi
.text:000000018001B700      call   sub_18001A7F4
.text:000000018001B705      mov     rbx, [rbp+1E70h+arg_0]
.text:000000018001B70C      lea    rax, qword_18004FFA0
.text:000000018001B713      lea    rcx, [rsp+1F70h+var_1F40]
.text:000000018001B718      mov     r8, rbx
.text:000000018001B71B      mov     r9d, 0B12B7A69h
.text:000000018001B721      mov     edx, edi
.text:000000018001B723      mov     [rsp+1F70h+var_1F50], rax
.text:000000018001B728      call   sub_18001A73C
.text:000000018001B72D      lea    rax, qword_18004FFB8
.text:000000018001B734      lea    rcx, [rsp+1F70h+var_1F40]
.text:000000018001B739      mov     r9d, 0C508CF8Bh
.text:000000018001B73F      mov     r8, rbx
.text:000000018001B742      mov     edx, edi
.text:000000018001B744      mov     [rsp+1F70h+var_1F50], rax
.text:000000018001B749      call   sub_18001A73C
.text:000000018001B74E      lea    rax, qword_18004FFD0
.text:000000018001B755      lea    rcx, [rsp+1F70h+var_1F40]
.text:000000018001B75A      mov     r9d, 35AF2123h
.text:000000018001B760      mov     r8, rbx
.text:000000018001B763      mov     edx, edi
.text:000000018001B765      mov     [rsp+1F70h+var_1F50], rax
.text:000000018001B76A      call   sub_18001A73C
.text:000000018001B76F      lea    rax, qword_18004FFE8
.text:000000018001B776      lea    rcx, [rsp+1F70h+var_1F40]
.text:000000018001B77B      mov     r9d, 0B598B908h
.text:000000018001B781      mov     r8, rbx
.text:000000018001B784      mov     edx, edi
.text:000000018001B786      mov     [rsp+1F70h+var_1F50], rax
.text:000000018001B78B      call   sub_18001A73C
.text:000000018001B790      lea    rax, qword_180050000
.text:000000018001B797      lea    rcx, [rsp+1F70h+var_1F40]
.text:000000018001B79C      mov     r9d, 765EF075h
.text:000000018001B7A2      mov     r8, rbx
.text:000000018001B7A5      mov     edx, edi
.text:000000018001B7A7      mov     [rsp+1F70h+var_1F50], rax
.text:000000018001B7AC      call   sub_18001A73C

```

Function `sub_18001A73C` is to resolve the function address (`syscall stub` address) by the hash. And function `sub_18001A7F4` used to populate the list with the system call SSN and system call stub. So, `sub_18001A7F4` is our target. In the following picture is the beginning of the function.

```

.text:000000018001A80D      sub     rsp, 30h
.text:000000018001A811      mov     rax, gs:30h
.text:000000018001A81A      mov     r13, rcx
.text:000000018001A81D      mov     r8d, 20202020h
.text:000000018001A823      mov     r9, [rax+_TEB.ProcessEnvironmentBlock]
.text:000000018001A827      mov     r10, [r9+_PEB.Ldr]
.text:000000018001A82B      add     r10, _PEB_LDR_DATA.InLoadOrderModuleList
.text:000000018001A82F      mov     r9, [r10]
.text:000000018001A832      loc_18001A832: ; CODE XREF: sub_18001A7F4+5B↓j
.text:000000018001A832      ; sub_18001A7F4+65↓j ...
.text:000000018001A832      cmp     r9, r10
.text:000000018001A835      jz     loc_18001A9E4

```

The function starts with getting a pointer to the first entry in `InLoadOrderModuleList` structure by going through reading the Process Environment Block (PEB). here in the picture, `r10` is holding the current entry of the structure and `r9` is like a variable to get each entry, this is the breaking condition of the loop as the `_LIST_ENTRY` structure wrap around itself (doubly linked list).

The next step is to get the Export directory of `ntdll.dll` but first, get `ntdll` address in memory.

```

.text:00000018001A83B      mov     rdi, [r9+LDR_DATA_TABLE_ENTRY.DllBase]
.text:00000018001A83F      mov     r9, [r9]
.text:00000018001A842      movsxd rax, [rdi+IMAGE_DOS_HEADER.e_lfanew]
.text:00000018001A846      mov     ecx, [rax+rdi+IMAGE_NT_HEADERS.OptionalHeader.DataDirectory.VirtualAddress]
.text:00000018001A84D      test   ecx, ecx
.text:00000018001A84F      jz     short loc_18001A832
.text:00000018001A851      lea    rsi, [rdi+rcx] ; data_dir_RVA + BaseAddress = address of IMAGE_DATA_DIRECTORY
.text:00000018001A851      ; First member is IMAGE_EXPORT_DIRECTORY
.text:00000018001A855      cmp    [rsi+IMAGE_EXPORT_DIRECTORY.NumberOfNames], 0
.text:00000018001A859      jz     short loc_18001A832
.text:00000018001A85B      mov     ecx, [rsi+IMAGE_EXPORT_DIRECTORY.Name]
.text:00000018001A85E      mov     eax, [rcx+rdi] ; DLL name
.text:00000018001A861      or     eax, r8d
.text:00000018001A864      cmp    eax, 'ldtn'
.text:00000018001A869      jnz   short loc_18001A832
.text:00000018001A86B      mov     eax, [rcx+rdi+4]
.text:00000018001A86F      or     eax, r8d
.text:00000018001A872      cmp    eax, 'ld.l'
.text:00000018001A877      jnz   short loc_18001A832
.text:00000018001A879      movzx  eax, word ptr [rcx+rdi+8]
.text:00000018001A87E      or     ax, 20h ; ' '
.text:00000018001A882      cmp    ax, 'l'
.text:00000018001A886      jnz   short loc_18001A832
.text:00000018001A888      mov     ebx, [rsi+IMAGE_EXPORT_DIRECTORY.AddressOfFunctions]
.text:00000018001A88B      mov     ebp, [rsi+IMAGE_EXPORT_DIRECTORY.AddressOfNames]
.text:00000018001A88E      mov     eax, [rsi+IMAGE_EXPORT_DIRECTORY.AddressOfNameOrdinals]
.text:00000018001A891      mov     r8d, edx
.text:00000018001A894      add    rax, rdi
.text:00000018001A897      add    rbx, rdi
.text:00000018001A89A      add    rbp, rdi
.text:00000018001A89D      mov    rcx, r13 ; void *
.text:00000018001A8A0      shl   r8, 4 ; Size
.text:00000018001A8A4      xor   edx, edx ; Val
.text:00000018001A8A6      mov   [rsp+68H+AddressOfFunctions], rbx
.text:00000018001A8AB      mov   [rsp+68H+AddressOfNameOrdinals], rax
.text:00000018001A8B3      mov   [rsp+68H+AddressOfNames], rbp
.text:00000018001A8B8      xor   r15d, r15d
.text:00000018001A8BB      call  memset

```

It is looking for the right module in the `InLoadOrderModuleList` by going through each entry, the `fLink` is a pointer to `LDR_DATA_TABLE_ENTRY` where we can get a pointer to the module. By parsing the module (going through PE file headers) to get the name of the DLL which resides in the Export directory (First member) which is the first member of `IMAGE_DATA_DIRECTORY` structure. It is then tested to see if it is the target module (`ntdll`). If the module is `ntdll` , it saves a pointer to `AddressOfFunctions` , `AddressOfNames` and `AddressOfNameOrdinals` . A memory region of size `0x1f40` is then zeroed as it will hold the structures of the system call information needed. The next part is checking the function prefix `Ki` and `Zw` . It looks for only one function prefixed by `Ki` with the hash `80CD4499h` , but I couldn't find function with this hash (using debugger). Then, a call to a hashing function is made. The hashing function is simple.

```

.text:00000018001A7C8      sub_18001A7C8      proc near ; CODE XREF: sub_18001A7F4+F5↓p
.text:00000018001A7C8      ; sub_18001A7F4+138↓p
.text:00000018001A7C8      xor     r9d, r9d
.text:00000018001A7C8      mov     r8, rcx
.text:00000018001A7CE      mov     eax, 52964EE9h
.text:00000018001A7D3      cmp    [rcx], r9b
.text:00000018001A7D6      jz     short locret_18001A7F2
.text:00000018001A7D8      loc_18001A7D8: ; CODE XREF: sub_18001A7C8+28↓j
.text:00000018001A7D8      movzx  ecx, word ptr [rcx]
.text:00000018001A7DB      mov     edx, eax
.text:00000018001A7DD      inc    r9d
.text:00000018001A7E0      ror    edx, 8
.text:00000018001A7E3      add    edx, ecx
.text:00000018001A7E5      mov     ecx, r9d
.text:00000018001A7E8      add    rcx, r8
.text:00000018001A7EB      xor    eax, edx
.text:00000018001A7ED      cmp    byte ptr [rcx], 0
.text:00000018001A7F0      jnz   short loc_18001A7D8
.text:00000018001A7F2      locret_18001A7F2: ; CODE XREF: sub_18001A7C8+E1↓j
.text:00000018001A7F2      rep retn
.text:00000018001A7F2      sub_18001A7C8      endp
.text:00000018001A7F2

```

It uses `0x52964EE9` as an initial key value to start the process then:

- Get 2-bytes of the Function name (little endian).
- Rotate the key by 8 (2 characters).
- Add the key and the 2-bytes of the name.
- Increment the counter by 1 (Resulting that all the chars in between the start and end taken two times in the calculation for example `ZwOpenProcess` will take `Wz` in the first iteration and `Ow` in the second and so on).
- The result of the addition is XORed with the key to produce the new key. The hash value returned is the last result of the XOR operation.

The resulting value is stored in the following form, in the pre-allocated space.

Address	Hex	ASCII
000000000013D6C0	62 7F CD 64 40 19 05 00	b . i d e . . . @ . b w . . .
000000000013D6D0	A1 0F 9C 9C 50 19 05 00	j . . . P . . . P . b w . . .
000000000013D6E0	CE 26 A2 DC A0 15 05 00	t & e U b w . . .
000000000013D6F0	00 00 00 00 00 00 00 00
000000000013D700	00 00 00 00 00 00 00 00
000000000013D710	00 00 00 00 00 00 00 00
000000000013D720	00 00 00 00 00 00 00 00

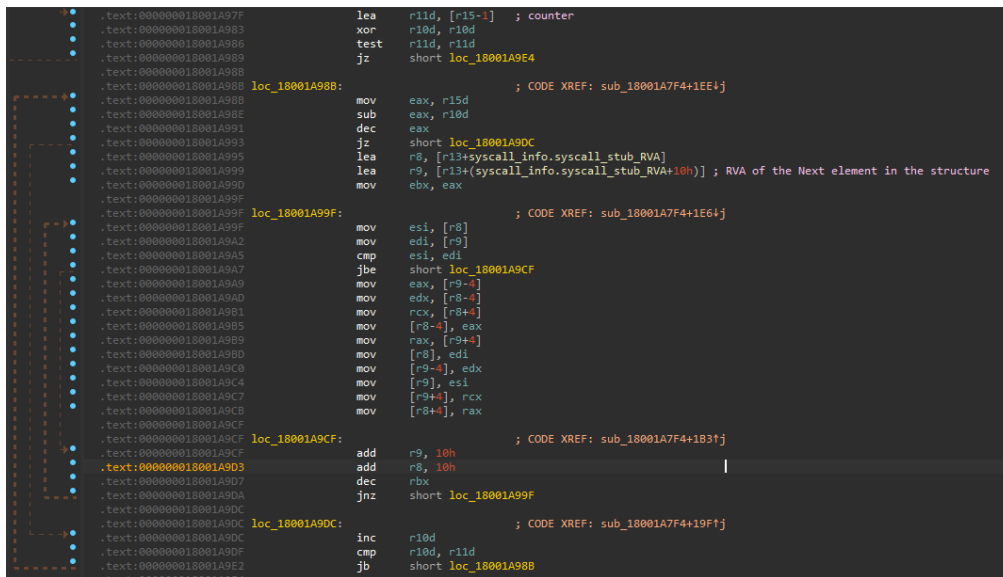
- The first **DWORD** is the hash.
- The second **DWORD** is the Relative Virtual Address (RVA) of the system call0.
- The third **QWORD** is the Virtual Address (VA) of the system call stub (RVA + ntdll Base Address).

So, it can be written as:

```

1      struct syscall_info {
2          DWORD API_hash;
3          DWORD syscall_stub_RVA;
4          QWORD syscall_stub_address;
5      };
    
```

After populating the structure with the addresses. The structure elements are being sorted by the RVA of the system call stub (second entry in the structure).



After the sorting algorithm is done, the memory structure look like the following:

Address	Hex	ASCII
000000000012D300	77 15 AA 3B 40 13 05 00	w . * ; @ . . . @ . N w . . .
000000000012D310	FD D6 91 2E 50 13 05 00	y O . . P . . . P . N w . . .
000000000012D320	B3 D8 6C F1 60 13 05 00	* 0 1 n N w . . .
000000000012D330	FB FD 5E C5 70 13 05 00	u y ^ A p . . . p . N w . . .
000000000012D340	8A 36 BD A7 80 13 05 00	. 6 % S N w . . .
000000000012D350	CB EF 59 D1 90 13 05 00	E T Y N N w . . .
000000000012D360	96 44 C7 9C A0 13 05 00	. D C N w . . .
000000000012D370	D1 EC 6F 3D B0 13 05 00	N i o = N w . . .
000000000012D380	DC C9 B3 2A C0 13 05 00	U E * = A A . N w . . .
000000000012D390	AA 8F 35 A4 D0 13 05 00	* . 5 s p D . N w . . .
000000000012D3A0	EF A4 49 6E E0 13 05 00	i r i n a a . N w . . .
000000000012D3B0	7A E9 E8 C0 F0 13 05 00	z e e A O o . N w . . .
000000000012D3C0	E1 1F 52 00 00 14 05 00	. a . R N w . . .
000000000012D3D0	33 FC BF 84 10 14 05 00	3 u j N w . . .
000000000012D3E0	CB C6 DF 49 20 14 05 00	E a t i N w . . .
000000000012D3F0	91 62 0D 79 30 14 05 00	. b . y O N w . . .
000000000012D400	A9 D7 1B EA 40 14 05 00	@ x . e e @ . N w . . .
000000000012D410	A8 DF 70 FD 50 14 05 00	e n n o p P . N w . . .

The first address is the address to the Lowest address **ZwMapUserPhysicalPagesScatter** (Could be different at newer versions of windows) at address **00000000774E1340** If we see the system call SSN of it:

```

00000000774E1340 4C:8BD1 mov r10,rcx
00000000774E1343 EB:00000000 mov eax,0
00000000774E1348 0F05          syscall
00000000774E134A C3           ret
    
```

system call number is zero. This is how it gets the SSN for any function, by iterating the structure to get the right hash, the counter will be used to get the SSN (SSN = counter). So far, this is remarkably like [MDSec](#) (8. Sorting by System Call

Address) implementation of the technique known as `FreshlyCalls` . We could rewrite the technique using MDSEc implementation as follows:

```

1  #define RVA2VA(Type, DllBase, Rva) (Type)((ULONG_PTR) DllBase + Rva)
2
3  static
4  void
5  GetSyscallList(PYSYSCALL_LIST List) {
6      PPEB_LDR_DATA      Ldr;
7      PLDR_DATA_TABLE_ENTRY LdrEntry;
8      PIMAGE_DOS_HEADER  DosHeader;
9      PIMAGE_NT_HEADERS  NtHeaders;
10     DWORD               i, j, NumberOfNames, VirtualAddress, Entries=0;
11     PIMAGE_DATA_DIRECTORY DataDirectory;
12     PIMAGE_EXPORT_DIRECTORY ExportDirectory;
13     PDWORD              Functions;
14     PDWORD              Names;
15     PWORD               Ordinals;
16     PCHAR               DllName, FunctionName;
17     PVOID               DllBase;
18     PSYSCALL_ENTRY     Table;
19     SYSCALL_ENTRY      Entry;
20
21     //
22     // Get the DllBase address of NTDLL.dll
23     // NTDLL is not guaranteed to be the second in the list.
24     // so it's safer to loop through the full list and find it.
25     Ldr = (PPEB_LDR_DATA)NtCurrentTeb()->ProcessEnvironmentBlock->Ldr;
26
27     // For each DLL loaded
28     for (LdrEntry=(PLDR_DATA_TABLE_ENTRY)Ldr->Reserved2[1];
29          LdrEntry->DllBase != NULL;
30          LdrEntry=(PLDR_DATA_TABLE_ENTRY)LdrEntry->Reserved1[0])
31     {
32         DllBase = LdrEntry->DllBase;
33         DosHeader = (PIMAGE_DOS_HEADER)DllBase;
34         NtHeaders = RVA2VA(PIMAGE_NT_HEADERS, DllBase, DosHeader->e_lfanew);
35         DataDirectory = (PIMAGE_DATA_DIRECTORY)NtHeaders->OptionalHeader.DataDirectory;
36         VirtualAddress = DataDirectory[IMAGE_DIRECTORY_ENTRY_EXPORT].VirtualAddress;
37         if(VirtualAddress == 0) continue;
38
39         ExportDirectory = (PIMAGE_EXPORT_DIRECTORY) RVA2VA(ULONG_PTR, DllBase, VirtualAddress);
40         DllName = RVA2VA(PCHAR, DllBase, ExportDirectory->Name);
41
42         if(*(ULONG*)DllName | 0x20202020) != 'ldtn') continue;
43         if(*(ULONG*)(DllName + 4) | 0x20202020) == 'ld.l') break;
44     }
45
46     NumberOfNames = ExportDirectory->NumberOfNames;
47
48     Functions = RVA2VA(PDWORD, DllBase, ExportDirectory->AddressOfFunctions);
49     Names      = RVA2VA(PDWORD, DllBase, ExportDirectory->AddressOfNames);
50     Ordinals   = RVA2VA(PWORD, DllBase, ExportDirectory->AddressOfNameOrdinals);
51
52     Table      = List->Table;
53
54     do {
55         FunctionName = RVA2VA(PCHAR, DllBase, Names[NumberOfNames-1]);
56         if(*(USHORT*)FunctionName == 'iK' && HashSyscall(FunctionName) == 0x8DCD4499) {
57             Table[Entries].API_Hash = HashSyscall(&FunctionName);
58             Table[Entries].syscall_stub_RVA = Functions[Ordinals[NumberOfNames-1]];
59             Table[Entries].syscall_stub_address = RVA2VA(void, DllBase, Functions[Ordinals[NumberOfNames-1]]);
60
61             Entries++;

```

```

62         if(Entries == MAX_SYSCALLS) break;
63
64     }
65     if(*(USHORT*)FunctionName == 'wZ') {
66         Table[Entries].API_Hash = HashSyscall(&FunctionName);
67         Table[Entries].syscall_stub_RVA = Functions[Ordinals[NumberOfNames-1]];
68         Table[Entries].syscall_stub_address = RVA2VA(void, DLLBase, Functions[Ordinals[NumberOfNames-1]]);
69
70         Entries++;
71         if(Entries == MAX_SYSCALLS) break;
72     }
73 } while (--NumberOfNames);
74
75 //
76 // Save total number of system calls found.
77 //
78 List->Entries = Entries;
79
80 //
81 // Sort the list by address in ascending order.
82 //
83 for(i=0; i<Entries - 1; i++) {
84     for(j=0; j<Entries - i - 1; j++) {
85         if(Table[j].syscall_stub_RVA > Table[j+1].syscall_stub_RVA) {
86             //
87             // Swap entries.
88             //
89             Entry.Hash = Table[j].Hash;
90             Entry.Address = Table[j].Address;
91
92             Table[j].API_Hash = Table[j+1].API_Hash;
93             Table[j].syscall_stub_RVA = Table[j+1].syscall_stub_RVA;
94             Table[j].syscall_stub_address = Table[j+1].syscall_stub_address;
95
96             Table[j+1].API_Hash = Entry.API_Hash;
97             Table[j+1].syscall_stub_RVA = Entry.syscall_stub_RVA;
98             Table[j+1].syscall_stub_address = Entry.syscall_stub_address;
99
100        }
101    }
102 }
103 }

```

The next thing is to use the structure to get the SSN, and `syscall` instruction to call. This is done by function `sub_18001A73C`.

```

.text:000000018001A73C      test     edx, edx
.text:000000018001A73E      jz      short locret_18001A78B
.text:000000018001A740      mov     [rsp+counter], rbx
.text:000000018001A745      push   rdi
.text:000000018001A746      sub     rsp, 20h
.text:000000018001A74A      xor     edi, edi
.text:000000018001A74C      mov     rax, rcx
.text:000000018001A74F      loc_18001A74F:                                ; CODE XREF: get_api+20↓j
.text:000000018001A74F      cmp     [rax], r9d
.text:000000018001A752      jz      short loc_18001A760
.text:000000018001A754      inc     edi
.text:000000018001A756      add     rax, 10h
.text:000000018001A75A      cmp     edi, edx
.text:000000018001A75C      jb     short loc_18001A74F
.text:000000018001A75E      jmp     short loc_18001A781
.text:000000018001A760      ; -----
.text:000000018001A760      loc_18001A760:                                ; CODE XREF: get_api+16↑j
.text:000000018001A760      mov     eax, edi
.text:000000018001A762      add     rax, rax
.text:000000018001A765      mov     rbx, [rcx+rax*8+8] ; API address
.text:000000018001A76A      mov     rcx, rbx
.text:000000018001A76D      call   get_syscall_ret_address

```

The function takes the following parameters:

- The array of structures that has the system call info (called `syscall_info` above)
- constant value `0x1F4` the maximum length of the structure members (structure size = `0x1F4 * 0x10`).
- Pre-Allocated memory
- The function hash.
- Global variable to get the system call SSN and stub. The function is simple, it searches the populated structure to find the given hash. If it's found, the counter value is taken and to get the Address of the system call stub. To get the address, the base address of the structure is added to the offset multiplied by `0x10` (struct size) and add 8 to get the last QWORD.

```

1      API_Address = *(STRUCT_BASE_ADDR + COUNTER * 0x10 + 8)

```

The address the passed to `get_syscall_ret_address` to get the `syscall ret` addresses to use it to execute the system call to bypass the callback mentioned before (call stack tracing is be used to detect this trick).

```

.text:000000018001A78C
.text:000000018001A78C      mov     [rsp+arg_8], 50Fh ; syscall
.text:000000018001A793      movzx  r8d, [rsp+arg_8]
.text:000000018001A799      mov     r9b, 0C3h ; ret
.text:000000018001A79C      xor     eax, eax
.text:000000018001A79E      loc_18001A79E: ; CODE XREF: get_syscall_ret_address+28↓j
.text:000000018001A79E      movsxd rdx, eax
.text:000000018001A7A1      cmp     r8w, [rdx+rcx]
.text:000000018001A7A6      jnz    short loc_18001A7AF
.text:000000018001A7A8      cmp     r9b, [rdx+rcx+2]
.text:000000018001A7AD      jz     short loc_18001A7B9
.text:000000018001A7AF      loc_18001A7AF: ; CODE XREF: get_syscall_ret_address+1A↑j
.text:000000018001A7AF      inc     eax
.text:000000018001A7B1      cmp     eax, 20h ;
.text:000000018001A7B4      jl     short loc_18001A79E
.text:000000018001A7B6      xor     eax, eax
.text:000000018001A7B8      retn
.text:000000018001A7B9 ; -----
.text:000000018001A7B9      loc_18001A7B9: ; CODE XREF: get_syscall_ret_address+21↑j
.text:000000018001A7B9      cdqe
.text:000000018001A7BB      add     rax, rcx
.text:000000018001A7BE      retn
.text:000000018001A7BE      get_syscall_ret_address endp

```

The global variable is used to store:

- QWORD to store System call address (function address at `ntdll`)
- QWORD to store `syscall` , `ret` instruction sequence address.
- DWORD to store system call number SSN. We can rewrite it as follows:

```

1      struct syscall_required_addresses {
2      QWORD syscall_stub_address;
3      QWORD syscall_intruction_address;
4      DWORD syscall_number;
5      };

```

(Creative names I know :)

There are some choices to call the required function. This is done based on the value at a global variable (`0x18004BC6C`):

- 1 : Direct call using the first member of the structure (Address of the function in `ntdll`)
- 2 : Indirect system call using trampoline jump using the system call number and the `syscall` address stored before.

```
.text:000000018001D0BF ; __int64 sub_18001D0BF()  
.text:000000018001D0BF sub_18001D0BF proc near ; CODE XREF: sub_18001B11C+5C↑p  
.text:000000018001D0BF mov r11, cs:syscall_inst_addr  
.text:000000018001D0C6 mov eax, cs:SSN  
.text:000000018001D0CC mov r10, rcx  
.text:000000018001D0CF jmp r11  
.text:000000018001D0CF sub_18001D0BF endp  
.text:000000018001D0CF
```

- anything else: Direct call to Win32 API.

System calls can be used to bypass user mood hooks but there are other methods to detect Direct and Indirect syscalls. To detect Direct system calls, Windows provides a large set of callback functions, one of them is

`KPROCESS!InstrumentationCallback`. This callback is triggered whenever the system returns from the kernel mode to user mode. This could be used to check the return address of the `syscall` which reveals the location of `syscall` instruction execution. This location should be `ntdll` but in case of the direct system calls, it will be from the `.text` section of the PE file. This was used by [ScyllaHide](#). Indirect system calls solved this problem by getting the address of `syscall` instruction in `ntdll` and jump to it. To detect indirect syscalls the call stack tracing method can be used to check from where the system call originated -before jumping to `ntdll` -. This also can be bypassed by creating a new thread to get a new call stack using callback functions like `TpAllocWork` and `RtlQueueWorkItem` . If you want to know more about this, you can read [Hiding In Plain Sight 1&2](#)

Note: This was personal notes I wrote when I was learning about syscalls, if there's anything not accurate, please let me know

<https://www.mdsec.co.uk/2020/12/bypassing-user-mode-hooks-and-direct-invocation-of-system-calls-for-red-teams/>

https://www.youtube.com/watch?v=eIA_eiqWefw&t=3176s

<https://offensivedefence.co.uk/posts/dinvoke-syscalls/>

<https://www.felixcloutier.com/x86/syscall.html>

<https://www.mdsec.co.uk/2022/04/resolving-system-service-numbers-using-the-exception-directory/>

<https://github.com/j00ru/windows-syscalls/>

<https://cocomelonc.github.io/malware/2023/06/07/syscalls-1.html>

<https://www.crummie5.club/freshycalls/>

<https://github.com/x64dbg/ScyllaHide/blob/master/HookLibrary/HookedFunctions.cpp>

<https://eversinc33.com/posts/avoiding-direct-syscall-instructions/>

<https://redops.at/en/blog/direct-syscalls-a-journey-from-high-to-low>

[https://github.com/dodo-sec/Malware-Analysis/blob/main/Cobalt Strike/Indirect Syscalls.md](https://github.com/dodo-sec/Malware-Analysis/blob/main/Cobalt%20Strike/Indirect%20Syscalls.md)

<https://github.com/crummie5/FreshyCalls/blob/112bdf0db63a5f7104ba5243af6a672bc098a1ad/syscall.cpp#L65>

<https://0xdarkvortex.dev/hiding-in-plainsight/>

<https://0xdarkvortex.dev/proxying-dll-loads-for-hiding-etwti-stack-tracing/>

Source: <https://d01a.github.io/syscalls/>