



Vulnerable Driver Manipulation

Abstract—This research paper documents the process of using a vulnerable Windows kernel driver exposing a physical memory read and write¹ primitive to call any function inside of the Windows kernel; while also teaching you the basics of paging and physical memory.

Keywords—Vulnerable Driver, Physical Memory, Virtual Address, Relative Virtual Address (RVA), Inline Hooking, Arbitrary Physical Memory Access

I. INTRODUCTION

Exploiting vulnerable Windows drivers to leverage kernel execution is not a new concept. Although software that exploits vulnerable drivers has been around for a long time, there has yet to be a highly modular library of code that can be used to exploit multiple drivers exposing the same vulnerability. Windows drivers exposing arbitrary physical memory read and write primitives are the most abundant form of vulnerable drivers. These drivers are used for many things ranging from reading CPU fan speeds to flashing BIOS. Although there are thousands of drivers that expose this primitive; doing anything useful with these drivers is not necessarily a straightforward task. In this research paper, I will be describing the steps on how to obtain kernel execution with an arbitrary physical memory read and write primitive. Furthermore, I will be demonstrating how simple it is to find and exploit such drivers whilst providing example code along the way.

II. LOCATING A VULNERABLE DRIVER

Finding a driver that exposes arbitrary physical memory read and write is as easy as googling the phrases: BIOS flashing utility for Windows, CPU fan speed utility for Windows, or ASUS overlocking utility for Windows. There are hundreds if not thousands of these drivers which allow for arbitrary physical memory read and write. In this research paper I'm going to specifically speak about phymem.sys; a Supermicro BIOS flashing Windows utility which I discovered during the process of making the introduction to this paper. Although there is an abundance of vulnerable drivers that expose physical memory read and write; typically the ranges of physical memory that can be manipulated are restricted. In the case of phymem.sys, only the first 4GB of physical memory can be arbitrarily read and written to. Be aware of these potential memory range restrictions when hunting for a vulnerable driver yourself.

Once you think you have found a vulnerable driver determining that it is in fact vulnerable can be done by concluding that user controlled data is passed to either: MmMapIoSpace², ZwMapViewOfSection, or MmCopyMemory³. This user controlled data is delivered to the driver's device control major function by calling DeviceIoControl. In the case of phymem.sys, user controlled data is passed to MmMapIoSpace.

III. INTERFACING WITH A VULNERABLE DRIVER

After determining that a driver is vulnerable the next step is to figure out how to interface with said vulnerable driver. The three most important values one should look for when reverse engineering the IRP_MJ_DEVICE_CONTROL function is: I/O control codes, IOCTL input and output buffer lengths, and finally input and output buffer(s)⁴. By observing how the user controlled data is used; a structure can be constructed.

```
1 InputBufferLength = StackLocation->Parameters
2   .DeviceIoControl
3   .InputBufferLength;
4
5 OutPutBufferLength = StackLocation->Parameters
6   .DeviceIoControl
7   .OutputBufferLength;
8
9 // IRP_MJ_DEVICE_CONTROL...
10 if (StackLocation->MajorFunction == 0xE)
11 {
12     v22 = StackLocation->Parameters
13         .Read
14         .ByteOffset
15         .LowPart;
16
17     // 0x80002000 (MAP_PHYSICAL_MEMORY)
18     switch (v22 + 0x7FFFE000)
19     {
20     case 0u:
21         // mind lengths for DeviceIoControl...
22         if (InputBufferLength != 16i64 ||
23             OutPutBufferLength != 8i64)
24             return {}; // invalid lengths...
25         else
26         {
27             // 4gb of physical memory limit
28             UserControlledPhysicalAddress.QuadPart =
29                 *((_QWORD*)SystemBuffer + 1)
30                 & 0xFFFFFFFFi64;
31
32             VirtualAddress = MmMapIoSpace
33                 (UserControlledPhysicalAddress,
34                 *((_QWORD*)SystemBuffer, NULL);
35             // IoAllocateMdl
36             // MmBuildMdlForNonPagedPool...
37             // MmMapLockedPagesSpecifyCache...
```

Listing 1. IRP_MJ_DEVICE_CONTROL Function of Phymem.sys

In the case of phymem.sys, the input buffer length is 16 bytes, and the output buffer length is 8 bytes. Looking at how SystemBuffer is used in Listing 1 you can see that it is a structure containing two QWORD sized fields. Further inspection concludes that the first QWORD field contains the size value in bytes of how much physical memory to map, and the second QWORD field is the physical address of memory to be mapped. As you can see in Listing 1, line 33, the top 32bits of the physical address is ignored. This limits the physical address to 32bits in size and thus the driver only allows us to map physical memory which is located in the first 4GB of physical memory.

```

1 #define MAP_PHYSICAL_MEMORY 0x80002000
2 #define UNMAP_PHYSICAL_MEMORY 0x80002004
3
4 // 16 bytes
5 typedef struct _map_phys_t
6 {
7     union
8     {
9         std::uintptr_t map_size; // + 0x0
10        std::uintptr_t virt_addr; // + 0x0
11    }
12    std::uintptr_t phys_addr; // + 0x8
13 } map_phys_t, *pmmap_phys_t;

```

Listing 2. Structure Passed via DeviceIoControl

Once a structure has been defined; interfacing with the vulnerable driver is just a matter of loading the driver into the kernel using NtLoadDriver⁵, and then controlling the driver with DeviceIoControl.

IV. SCANNING PHYSICAL MEMORY

Although physical memory may seem ambiguous, it is organized into fixed sized chunks called pages. Each page on a 64-bit system using a four layer paging table configuration is 4kB⁶. Within this chunk size memory is contiguous. The last 12bits of every 64-bit virtual address is called the page offset. Knowing this one can scan every single page at a specific offset for specific bytes.

```

1 PAGE:00000001C01265F0 ; ===== S U B R O U T I N E =====
2 PAGE:00000001C01265F0 public NtGdiDdDDICreateContext
3 PAGE:00000001C01265F0 NtGdiDdDDICreateContext proc near
4 PAGE:00000001C01265F0     mov     [rsp+arg_8], rbx
5 PAGE:00000001C01265F5     mov     [rsp+arg_10], rdi
6 PAGE:00000001C01265FA     mov     [rsp+arg_18], r12
7 PAGE:00000001C01265FF     push   r13
8 PAGE:00000001C0126601     push   r14
9 PAGE:00000001C0126603     push   r15
10 PAGE:00000001C0126605     sub     rsp, 200h
11 PAGE:00000001C012660C     mov     rax, cs:__security_cookie
12 PAGE:00000001C0126613     xor     rax, rsp
13 PAGE:00000001C0126616     mov     [rsp+218h+var_28], rax
14 PAGE:00000001C012661E     mov     r14, rcx
15 PAGE:00000001C0126621     mov     [rsp+218h+var_170], rcx
16 PAGE:00000001C0126629     mov     [rsp+218h+var_1A0], rcx
17 PAGE:00000001C012662E     or     rdi, 0FFFFFFFFFFFFFFFFh
18 PAGE:00000001C0126632     mov     [rsp+218h+var_1C0], edi

```

Listing 3. NtGdiDdDDICreateContext First Few Instructions

In Listing 3, the page offset for the system routine NtGdiDdDDICreateContext is 0x5F0. Simply scanning every single physical page at offset 0x5F0 for opcodes in NtGdiDdDDICreateContext is enough to get a handful of results. Testing each occurrence is required in order to determine that we have found the real⁷ NtGdiDdDDICreateContext in physical

memory. Scanning each page one at a time is quite slow so to expedite the process VDM creates a new thread for each physical memory range⁸.

V. ELEVATING TO KERNEL EXECUTION

Everytime an occurrence of NtGdiDdDDICreateContext's bytes is found in physical memory, a test is conducted to determine if the correct memory has been located. This test places some assembly code over the first few instructions of NtGdiDdDDICreateContext. NtGdiDdDDICreateContext is then called to see if the desired instructions were executed. Finally regardless of the situation the original bytes are restored.

```

1 bool vdm_ctx::valid_syscall(void* syscall_addr) const
2 {
3     static std::mutex syscall_mutex;
4     syscall_mutex.lock();
5
6     static const auto proc =
7         GetProcAddress(
8             LoadLibraryA(syscall_hook.second),
9             syscall_hook.first
10        );
11
12    // 0: 48 31 c0     xor rax, rax
13    // 3: c3         ret
14    std::uint8_t shellcode[] = { 0x48, 0x31, 0xC0, 0xC3 };
15    std::uint8_t orig_bytes[sizeof shellcode];
16
17    // save original bytes and install shellcode...
18    vdm::read_phys(
19        syscall_addr,
20        orig_bytes,
21        sizeof orig_bytes);
22
23    vdm::write_phys(
24        syscall_addr,
25        shellcode,
26        sizeof shellcode);
27
28    auto result = reinterpret_cast<
29        NTSTATUS(__fastcall*)(void)>(proc)();
30
31    vdm::write_phys(
32        syscall_addr,
33        orig_bytes,
34        sizeof orig_bytes);
35
36    syscall_mutex.unlock();
37    return result == STATUS_SUCCESS;
38 }

```

Listing 4. Checking Physical Page to Ensure its the Right One

Now that we know the correct location of NtGdiDdDDICreateContext's routine in physical memory; we can install an inline hook at the beginning of the function everytime we want to call a specific function in the kernel, and then restore the original bytes once the syscall has finished. Locating specific routines in the kernel can be done with simple arithmetic. The location of kernel module base addresses can be obtained simply with NtQuerySystemInformation using SystemModuleInformation. This allows us to calculate the absolute virtual address of any kernel function we want. Simply by loading the driver⁹ which contains the function desired and subtracting the address of it from the base address of the loaded driver a relative virtual address is produced. Subsequently the inverse operation (addition) can be applied to the kernel modules base address to produce the absolute kernel virtual address of the desired function. In conjunction with the ability to inline hook

NtGdiDdDDICreateContext this allows a VDM user to call any kernel function they desire.

8. <https://www.remkoweitjnen.nl/blog/2009/03/20/reading-physical-memory-size-from-the-registry/>

9. LoadLibraryEx using DONT_RESOLVE_DLL_REFERENCES for its dwFlag parameter.

VI. USING A VULNERABLE DRIVER WITH VDM

VDM allows a programmer to easily integrate a vulnerable driver into the project simply by coding four functions used by the rest of the project. The four functions that are required for VDM to work are: vdm::load_drv, vdm::unload_drv, vdm::read_phys, and vdm::write_phys. Once these functions have been programmed appropriately the library will take care of the rest. Most drivers map and unmap physical memory, so when programming vdm::read_phys and vdm::write_phys map the physical memory, use memcpy, then unmap the physical memory.

VII. LIMITATIONS

- VDM will not work on HVCI systems.
- Inline hook on syscall is not thread safe and can cause system instability.

VIII. CONCLUSION

VDM abstracts the concept of a vulnerable driver that exposes physical memory read and write to a method in which you can call into any kernel function you desire. The overabundance of vulnerable drivers exposing this primitive allows VDM to be much more modular and thus much more attractive than other public options.

IX. EXAMPLES

- VDM: <https://github.com/xerox/vdm>
- TDL: <https://github.com/hfiref0x/TDL>
- KDU: <https://github.com/hfiref0x/KDU>
- nasa-mapper: <https://github.com/xerox/nasa-mapper>
- kdmapper: <https://github.com/z175/kdmapper>
- gdrv-loader: <https://github.com/alxbrn/gdrv-loader>
- drvmap: <https://github.com/not-wlan/drvmap>

NOTES

1. Although I say read and write physical memory, this typically is manifested by mapping and unmapping physical memory via MmMapIoSpace and ZwMapViewOfSection. There are some drivers that directly expose reading and writing to physical memory via MmCopyMemory.
2. MmMapIoSpace by itself is not enough to map physical memory into a user-mode process. IoAllocateMdl, MmMapLockedPages, and a few other functions are required.
3. MmCopyMemory can be used to read and write physical memory when Flags parameter is MM_COPY_MEMORY_PHYSICAL.
4. In the case of phymem.sys "SystemBuffer" is used for both the input and output buffer.
5. I have made a C++17 header-only library for this by the name of loadup. It can be found here: <https://github.com/xerox/loadup>.
6. Although pages are 4kB by default, a four layer paging table system supports 2MB pages and even 1GB pages, when a page is large, the part of the address used to describe the index into a paging table is instead used to describe an offset into the physical page.
7. These bytes are the ones executed when NtGdiDdDDICreateContext is invoked.