Interpreter Exploitation

Dionysus Blazakis

Independent Security Evaluators
Baltimore, MD 21209

`dion@securityevaluators.com`

Abstract

As remote exploits further dwindle and perimeter defenses become the standard, remote client-side attacks are becoming the standard vector for attackers. Modern operating systems have quelled the explosion of client-side vulnerabilities using mitigation techniques such as data execution prevention (DEP) and address space layout randomization (ASLR). This work illustrates two novel techniques to bypass these mitigations. The two techniques leverage the attack surface exposed by the script interpreters commonly accessible within the browser. The first technique, pointer inference, is used to find the memory address of a string of shellcode within the Adobe Flash Player's ActionScript interpreter despite ASLR. The second technique, JIT spraying, is used to write shellcode to executable memory, bypassing DEP protections, by leveraging predictable behaviors of the ActionScript JIT compiler. Previous attacks are examined and future research directions are discussed.

1 Introduction

The difficulty in finding and exploiting a remote vulnerability has motivated attackers to devote their resources to finding and exploiting client side vulnerabilities. This influx of client side attacks has pushed Microsoft to implement mitigation techniques to make exploiting these vulnerabilities much harder. Sotirov and Dowd [13] have described in detail each of the mitigation techniques and their default configurations on versions of Windows through Windows 7 RC. Their work showed some of the bypass techniques available at the time and how Microsoft's design choices have influenced the details of these bypasses. After reading this work, one thing that stands out is how ripe a target the browser is for exploitation — the attacker can use multiple plug-ins, picking and choosing specific exploitable features, to set-up a reliable exploit scenario. The main contribution of this paper is to further prove how much power an attacker gains through access to an interpreter environment. Within a programmable environment, the complex interactions between the interpreter and the operating systems memory subsystem are highly influenced by the attacker and enable attacks that are difficult to predict at design time.

It would be difficult to design a more exploit friendly environment than the classic web browser. Bursting at the seams with plug-ins, it requires a robust parser to salvage what could be any of 6+ versions of mark-up taking into account possible legacy work-arounds. With the advent of "Web 2.0", a browser must also include a high performance scripting environment with the ability to rewrite those parsed pages dynamically. The document object model exposed to the scripting runtime continues to grow. Finally, most browsers are now taking advantage of recent JIT and garbage collection techniques to speed up Javascript execution. All this attack surface without even discussing the details of common plug-ins.

Rich internet applications (RIAs) are not going away and Adobe currently maintains a hold over the market with Flash; the Flash Player is on 99% of hosts with web browsers installed. Suns Java Runtime Environment is another interpreter commonly installed. Microsoft Silverlight is an RIA framework based upon the .NET runtime and tools. Each of these plug-ins require a complex parser and expose more attack surface through a surplus of attacker reachable features. For example, Adobe Flash Player implements features including a large GUI library, a JIT-ing 3D shader language, an RMI system, an ECMAScript based JIT-ing virtual machine, embeddable PDF support, and multiple audio and video embedding or streaming options. All of this is available by default to the lucky attacker.

Considering this, it is worth putting in time and effort to develop application specific techniques that will help exploit vulnerabilities within this browser ecosystem. DEP and ASLR are very real thorns in the side of an exploit developer. DEP makes locating shellcode difficult; the attacker must find a page with executable permission and find a way to write to it when it has writable permission and figure out the location. Alternatively, the attacker could find some code to reuse as in the return-oriented programming based attacks [7, 11]. ASLR further complicates an attack by obfuscating the base address of the loaded images. See Sotirov and Dowd [13] for a detailed explanation on the implementation of ASLR and DEP in various Windows operating systems. For this, we will ignore the possibility of an ASLR design or implementation flaw; we will assume

the attacker cannot guess the location of a loaded image, the heap, or the stack by leveraging a flaw in the operating system.

This paper focuses on two general techniques to bypass these mitigations: pointer inference and JIT spraying. Pointer inference is the act of recovering the memory address of an internal object via "normal" (i.e. not through an exploit) interactions with the target software. Adobe Flash Player is used as an example. We will present and walk through a script to determine the address of a string object within Flash Player's ActionScript interpreter. Next, we will present JIT spraying; a technique similar to heap spraying. The attacker will allocate many executable pages of attacker influenced data. We will show how to construct an ActionScript script that, when JIT compiled by the ActionScript engine, lays out stage-0 shellcode to load and execute the next stage of shellcode (located in an ActionScript string). Next, we discuss the possibility of using the two techniques together to attack a Windows Vista target. Finally, future research directions are discussed.

2 Pointer Inference

Getting from a software flaw to control of the instruction pointer to exploit is much more difficult when the address space is randomized. Even getting from a proof-of-concept crash to control of EIP may be difficult. Proving a vulnerability is sometimes easily accomplished using a heap spray to place an attacker constructed structure at a known address. Unfortunately for the exploit developer, heap sprays are not always reliable (and are certainly not always possible if the attacker cannot allocate large heap objects). Additionally, heap spray mitigation techniques have already appeared in the academic world [10] and Microsoft has a simple mitigation (mapping the commonly used heap spray pages at process start-up) in their EMET tool [8]. Information leak vulnerbilities will soon be highly prized for their use in reliable exploitation. The pointer inference technique we describe below is one example of an information leakage flaw. The ability to derive the heap address of an object at runtime provides the attacker with an address upon which to anchor an attack.

Scripting environments are the perfect target for pointer inference; the objects usually live on the heap and the runtime language and libraries provide multiple ways to manipulate and inspect objects. Additionally, scripting languages tend to be dynamically typed, so the built in container objects are heterogeneous objects stored by reference are often stored in the same type as those stored by value. The goal is to find a way to determine the memory address of a script object in the interpreter/virtual machine. Extracting the location of a Javascript string containing shellcode, while not the only

use of the technique, would be a good exploit building block. Target VMs include (ECMA or Java or Action)-Script, Java, Python, Ruby, PHP, and the .NET CLR. The Javascript engine in the browser, the Javascript engine in Adobe Reader, and the ActionScript engine in the Adobe Flash Player are all available from within most browser installations (if the target has the Adobe plug-ins installed and don't have Javascript disabled).

For this paper, we will show how to derive the memory address of an object in the ActionScript virtual machine using the ordering of objects when iterating over an ActionScript Dictionary object. We've chosen to show this proof of concept using the Flash Player for multiple reasons: we've spent more time reversing it than the others, it is cross platform – regardless of browser it is the same code base for the plug-in, and, last but not least, Adobe has released the source to a fork of their engine [1]. To understand the details of the technique, we will first describe how the objects are stored internally by the interpreter and then how the built-in Dictionary container is implemented by the interpreter.

2.1 ActionScript Objects in Memory

Prior to understanding the technique for disclosing the address of an ActionScript objects, we must understand how the objects are stored internally by the interpreter. When an ActionScript object is instantiated by the executing script, the interpreter provides a word sized entry to contain this object. It fills the entry based upon the type of the object. If the object is a small primitive object, such as an integer or boolean, it is stored by value. For other Objects, such as doubles, strings, or class instances, the object will be stored by reference; the interpreter will allocate a buffer and store the address of the buffer. The same entry can store either type (a value or reference).

ActionScript is a dynamically typed language. Dynamically typed languages do not assign types to variables at compile time. Instead of building type annotations into the language and enforcing typing constraints at compile time, these languages provide runtime functions for examining and comparing object types. The interpreter then ensures that all operations are valid for the operands used at the time of the operation. When the interpreter detects a type mismatch (e.g. an instance is queried for a method it does not implement), ActionScript will either throw an exception or perform an implicit coercion. This runtime typing requires the object heap to store both the type information and the value.

To handle this runtime typing requirement, the ActionScript interpreter represents internal objects using tagged pointers; internally, this object is called an "atom". Tagged pointers are a common implementation technique to differentiate between those objects stored

by value and those stored by reference using the same word sized memory cell. A tagged pointer stores type information in the least significant bits and stores a type specific value in the most significant bits. As shown in Figure 1 and Table 1, the ActionScript atom is 32 bits wide; it allocates 3 bits to store the type information and uses 29 bits for the value.

Some examples might help illustrate how this works. Take the following ActionScript declarations:

```
var x = 42; // Integer atom
var y = WOOT 2010; // String atom
var z = new Dictionary(); // Object atom
var b = true; // Boolean atom
```

The variable x is a local variable holding the value 42. The interpreter will create a mapping between the local variable within the current scope and the value 42. As described above, the value will be stored internally as an atom. An integer atom, which can hold values between -2^{28} and $2^{28} - 1$, is created by shifting the value left 3 bits to make room for the type tag. The integer atom is tagged with a 6 as shown in Figure 1. This process is shown in the Python session below:

```
>>> def atomFromInteger(n):
return (n << 3) | 6

>>> '0x%08x' % (atomFromInteger(42), )
'0x00000156'
```

The String and Object atoms are "reference" atoms. They store pointers to garbage collected memory on the interpreter heap. Converting the y and z variables to atoms requires first allocating a block of memory to store the value and then creating the atoms using the memory address of the actual value. Below is an example of doing this for z in Python – the extra calls are mocked up.

```
>>> def atomFromObject(obj):
return (obj & ~7) | 1

>>> a = AS.HeapAlloc(size_of_Dictionary)
>>> Dictionary.initialize(a)
>>> '0x%08x' % (atomFromObject(a), )
'0x00c8b301'
```

The goal of illustrating this internal representation is to be able to explain that both values and references (i.e. memory addresses or pointers) are used as atoms by the interpreter. Next, we will explain the use and implementation of the ActionScript Dictionary class.

2.2 ActionScript Dictionary Objects

The built-in ActionScript Dictionary class exposes an associative map data structure. When used from within an ActionScript script, it provides an interface to associate any ActionScript object with any other Action-Script object as a key/value relation. The Dictionary object can then be queried using square brackets, similar to a Python `dict`. Additionally, the user can iterate over the dictionary to operate on each key/value pair. The order of the iteration is not specified by the definition of the API and is an implementation detail not to be relied upon. Example use:

```
var dict = new Dictionary();

dict[Alpha] = 0x41414141;
dict[true] = 1.5;

var k;
for (k in dict)
{
  f(dict[k]);
}
```

Internally, Flash Player's Dictionary class is implemented using a hashtable. The hashtable derives the hash from the key atom and stores the key and value atom together in the table. When iterating over the Dictionary, the hashtable is walked from lowest to highest hash value. The hash table is always a power of two in size; this is maintained for two reasons: the hash function now becomes a fast masking operation and the constants used by the quadratic probe rely on a power of two sized table. The hash table grows to the next power of two when the number of empty cells in the table drops below 20% of the total size. To grow the hashtable, a new table is allocated and all entries are rehashed and inserted into the new table. This hashtable implementation discloses an ordering between integer (value) atoms and object (reference) atoms — the object atoms are compared directly to the integer atoms. The hash function will remove some of the most significant bits of the atoms but a large hashtable will use most of the bits. This ordering is used to disclose memory addresses of reference atoms (Objects, Strings).

2.3 Integer Sieve

Since integers are placed into the hashtable using their value as the key (of course, the any top bits will be masked off), we can determine the atom value of some ActionScript object by measuring where the new object is found when iterating over the hashtable. By recording the integers that fall before and after the newly inserted object, we can derive a bound on the atom of the new object. Since Object atoms are just pointers (with the first 3 bits modified), we can disclose as many bits of a pointer as we can grow the hashtable.

To avoid the problem of a hash collision, we perform the test twice: once with all even integers and once with

31   3 2 1 0
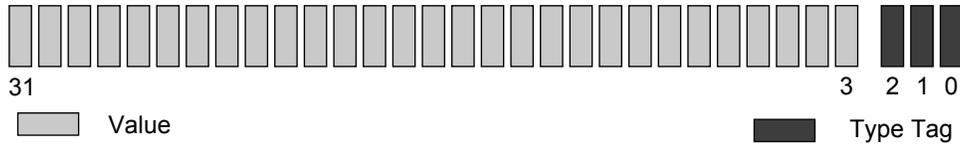
◻ Value   ◼ Type Tag

Figure 1: An ActionScript Atom

Type	Tag		Value
Untagged	000	0	
Object	001	1	8 byte aligned pointer
String	010	2	8 byte aligned pointer
Namespace	011	3	8 byte aligned pointer
"undefined"	100	4	0
Boolean	101	5	1 for True and 0 for False
Integer	110	6	2's complement, $[-2^{28}, 2^{28} - 1]$
Double	111	7	8 byte aligned pointer

Table 1: ActionScript atom tag constants

all odd integers (up to some power of two — the larger, the more bits we discover). After creating the Dictionaries, we insert the victim object into both Dictionaries. The values associated with the keys stored in the Dictionary do not impact any of this — only the keys and their ordering are used. Next, search each Dictionary using the for-in construct, recording the last key visited and breaking when the current key is the victim object. We now have two integer values (the last integers before the victim object for both the even and the odd Dictionaries).

The two integers should differ by 17. This is due to the linear probe; when a hashtable insert collides, it uses a quadratic probe to find an empty slot. The first try is at $(hash(victim) + 8)$ which always collides — $2n+8 = 2(n+4)$ or $(2n+1)+8 = 2(n+4)+1$. The next try is $(hash(victim) + 17)$ which always succeeds — $2n+17 = 2(n+8)+1$ or $(2n+1)+17 = 2(n+9)$. The only way for the two integers to differ by anything other than 17 is if the probe wrapped. Otherwise, the smaller integer is the one from the Dictionary that didn't have the collision. When the difference isn't 17 (wrapped around), the larger value is from the Dictionary that didn't have the collision. We now have the integer that, when turned into an atom is 8 smaller than the victim atom. Finally, to get the victim atom from the integer, x: $(x << 3 + 8)$ or more to the point $((x + 1) << 3)$.

That is a handful; let's walk through the execution:

```
// First, create the Dictionaries
var even = new Dictionary();
var odd = new Dictionary();

// Now, fill the Dictionary objects
// with the integer atoms
var index;
for (index = 0; index < 8; index += 1) {
```

```
    even[index * 2] = true;
    odd[index * 2 + 1] = true;
}
```

Figure 2 shows what the heap will look like after inserting the integer atoms. Next, insert the reference object to leak the address of:

```
var victim = AAAAAAAA;

even[victim] = true;
odd[victim] = true;
```

In Figure 3, you can see how the insertion of the victim atom wraps around once (the table is 16 entries long, so on a collision using these integer sieves, 17 will be added to the index modulo 16 — just adding 1 to the index where the collision took place). The last step is to iterate over both Dictionaries recording the integer just prior to finding the victim object:

```
var curr, evenPrev, oddPrev;

for (curr in even)
{
    if (curr == victim) { break; }
    evenPrev = curr;
}

for (curr in odd)
{
    if (curr == obj) { break; }
    oddPrev = curr;
}
```

After executing this snippet, `evenPrev` will contain the value 0 (atom: `0x00000006`) and `oddPrev` will

index
| 0x00000026 |
0x00800000

even
| 0x00810001 |
0x00800008

odd
| 0x00820001 |
0x00800010

0x00800008

Even sieve (0x00810000):

0x00000006	0x0000000D
0x00000016	0x0000000D
0x00000026	0x0000000D
0x00000036	0x0000000D
0x00000046	0x0000000D
0x00000056	0x0000000D
0x00000066	0x0000000D
0x00000076	0x0000000D

0x00810000

Odd sieve (0x00820000):

0x0000000E	0x0000000D
0x0000001E	0x0000000D
0x0000002E	0x0000000D
0x0000003E	0x0000000D
0x0000004E	0x0000000D
0x0000005E	0x0000000D
0x0000006E	0x0000000D
0x0000007E	0x0000000D

0x00820000

Figure 2: ActionScript heap after Integer atoms are inserted

Legend
name
address

index
| 0x00000026 |
0x00800000

victim
| 0x0081000A |
0x00800018

| AAAAAAAA... |
0x00810008

even
| 0x00810001 |
0x00800008

odd
| 0x00820001 |
0x00800010

Even sieve (0x00810000):

0x00000006	0x0000000D
0x0081000A	0x0000000D
0x00000016	0x0000000D
0x00000026	0x0000000D
0x00000036	0x0000000D
0x00000046	0x0000000D
0x00000056	0x0000000D
0x00000066	0x0000000D
0x00000076	0x0000000D

0x00810000

Odd sieve (0x00820000):

0x0000000E	0x0000000D
0x0081000A	0x0000000D
0x0000001E	0x0000000D
0x0000002E	0x0000000D
0x0000003E	0x0000000D
0x0000004E	0x0000000D
0x0000005E	0x0000000D
0x0000006E	0x0000000D
0x0000007E	0x0000000D

0x00820000

Figure 3: ActionScript heap after adding the victim atom to the sieves

contain the value 1 (atom: `0x0000000E`). Let's pretend that `evenPrev` and `oddPrev` differ by 17 (and the tables were a more reasonable size — say, 4 million). `evenPrev` is smaller; this means we can derive the lowest bits of the victim address by adding 1 to `evenPrev` (the value, 0, which is what is exposed to us in ActionScript) and shifting left 3. This results in `0x00000008`. That matches the lowest bits of the victim address.

This technique is by no means the only vectors for information leakage. If we restrict ourselves to data structure leaks, internal ids (see Python) and hash functions (as opposed to the ActionScript technique, where the hash is not available directly to the script; see .NET and Java). These functions need a unique value per heap object and the address is one such value that could be used.

3 JIT Spray

Data Execution Prevention (DEP) makes executing delivered shellcode quite difficult — the stack and default heaps are marked non-executable. Currently, the best public method to bypass DEP is to find a loaded DLL that is not ASLR protected or leak the load address of the module and use the available code to manufacture return-oriented shellcode [11, 7] that turns off DEP for the process or allocates executable memory and copies the next stage of shellcode there.

Most modern interpreters implement a just-in-time (JIT) compiler to transform the parsed input or bytecode into machine code for faster execution. JIT spraying is the process of coercing the JIT engine to write many executable pages with embedded shellcode. This shellcode will entered through the middle of a normal JIT instruction. For example, a Javascript statement such as `var x = 0x41414141 + 0x42424242;` might be compiled to contain two 4 byte constants in the executable image (for example, `mov eax, 0x41414141; mov ecx, 0x42424242; add eax, ecx`). By starting execution in the middle of these constants, a completely different instruction stream is revealed.

The rest of this section explains one implementation of this for the Adobe Flash Player ActionScript bytecode JIT engine. The end result is a Python script and an ActionScript script. The Python script generates the ActionScript. The ActionScript, when loaded and after the bytecode engine performs a JIT compile, lays out an executable page on the heap. When this executable page is entered at a known offset (`0x6A`, for the given example code below) will execute stage-0 shellcode that marks the rest of the page RWX and copies the next stage of shellcode from an ActionScript string that can be modified before compilation.

The key insight is that the JIT is predictable and must copy some constants to the executable page. Given a uniform statement (such as a long sum or any repeating pattern), those constants can encode small instructions and then control flow to the next constant's location.

3.1 Development

By placing a breakpoint on the VirtualProtect calls in the Flash Player, we can witness the JITed code that is generated from the ABC bytecode (which is compiled ActionScript). By experimentation, we are able to determine that a long XOR expression ($a \otimes b \otimes c \otimes d \otimes \ldots$) would be compiled down to a very compact set of XOR instructions. For example, after JIT-ing:

```
var y = (
    0x3c54d0d9 ^
    0x3c909058 ^
    0x3c59f46a ^
    0x3c90c801 ^
    0x3c9030d9 ^
    0x3c53535b ^
    ...
```

is turned into:

```
03470069  B8 D9D0543C   MOV EAX,3C54D0D9
0347006E  35 5890903C   XOR EAX,3C909058
03470073  35 6AF4593C   XOR EAX,3C59F46A
03470078  35 01C8903C   XOR EAX,3C90C801
0347007D  35 D930903C   XOR EAX,3C9030D9
03470082  35 5B53533C   XOR EAX,3C53535B
```

Now, note that if execution begins at 0x0347006A:

```
0347006A  D9D0          FNOP
0347006C  54            PUSH ESP
0347006D  3C 35         CMP AL,35
0347006F  58            POP EAX
03470070  90            NOP
03470071  90            NOP
03470072  3C 35         CMP AL,35
03470074  6A F4         PUSH -0C
03470076  59            POP ECX
03470077  3C 35         CMP AL,35
03470079  01C8          ADD EAX,ECX
0347007B  90            NOP
0347007C  3C 35         CMP AL,35
0347007E  D930          FSTENV DS:[EAX]
```

This is a popular GetPC method — using the floating point state save opcode [5]. In above code, the control flow from one attacker controlled DWORD operand to the next only takes up a single byte (the byte of the `CMP` opcode). By taking advantage of the single byte `XOR EAX` opcode, the `CMP AL` is a semantic NOP and does not take up the two bytes a `JMP +n` would in the attacker controlled operand bytes.

ActionScript allows the dynamic loading (and runtime generation) of bytecode. Using this type of reflection,

we can force the repeated loading of a given bytecode file spraying the constructed JIT code on the heap. This is the JIT spray. With a reasonable guess and luck, the attacker can execute shellcode despite ASLR and DEP. In the last section, we will take a look at using a combination of the two techniques to achieve reliable exploitation.

4 Putting It Together

One thing that was glossed over in the pointer inference discussion was the value that is pointed to by the leaked pointer. What does the address leak gain an attacker? What attacker controllable values are available via the leaked pointer? For Object atoms, the pointer points to a C++ instance that doesn't have any easily controllable fields. For String atoms, the pointer points to a C++ instance containing a length field and a pointer to the string buffer. This will certainly be useful in some instances, but it isn't a direct pointer to the string buffer. In the end, it's really just a pointer into the heap. To try and find a use for this pointer, we must understand the Flash heap and some of the details of the Windows memory subsystem.

Flash objects are allocated using a custom allocator which boils down to `VirtualAlloc`. When expanding the heap, the allocator tries to allocate the next chunk contiguously in 16MB increments. If that first allocation fails, it tries again without the start address hint (i.e. getting rid of the contiguous constraint Flash asks for any block of memory to fit the request). The pages for the JIT engine are allocated directly using `VirtualAlloc` by estimating the space needed. This estimate is made by counting opcodes when loading the bytecode. The allocation for all methods/functions found in a bytecode file occurs up front (and will allocate this space on-demand when the first methods are called and compiled). `VirtualAlloc` will map pages at a 64KB granularity and does so with a linear scan finding the first hole that matches the size requested. With these details, we can come up with a use for the heap address primitive.

The plan is to put the two techniques together for a reliable bypass of DEP and ASLR. To do this, we must use the heap address primitive to determine the address of a JIT block. Here is the current method we've developed in the browser environment:

1. Open a SWF file that contains enough bytecode to force an allocation of 0x01010000 bytes.
2. Open a SWF file that sprays the heap with many small ActionScript objects.
3. Open a second SWF forcing an allocation of `0x00FF0000` bytes.
4. Remove the first large SWF file; this will deallocate/unmap the 0x1010000 bytes.

5. Spray more than 16MB of small ActionScript objects, keeping them in a linked list structure.
6. Load the SWF with the JIT Spray.
7. Iterate over the linked list structure at large intervals (`0x00100000`) recording the address of the object.
8. When an unsmooth edge is found in the recorded addresses, this marks the start of the newly mapped space. This new space should be the start of the memory used for the first SWF.
9. A tight bound can be obtained for the start address by refining the above search using successively smaller intervals and updating the bounds.
10. Finally, we know the JIT spray should have reached this address we've found in the previous step + `0x01000000`.

This whole process is quite roundabout, but the Adobe Flash player doesn't seem to ever unmap memory. By not unmapping memory used by the object heap, the JIT engine will never allocate a page that was once used by the Object heap.

Once the address of a JIT sprayed block is found, the target EIP can be calculated and the exploit can be triggered. This process is quite reliable, but with the current speed of the address leak it takes between 5 and 10 minutes to trigger. If a different address leak were found this would be quicker. While testing this on a Windows Vista target with IE8, a straightforward JIT spray can be used and a predictable address is not hard to find on a lightly used browser. In a browser with a lot of memory activity, the longer method is more reliable.

This attack is not at all straightforward. It involves multiple flaws to work and requires more time than most current exploits. This is evidence of the impact current mitigations have on exploit development; exploits are increasing in complexity and prerequisites.

5 Related Work

We are not the first to take a shot at the current class of exploit mitigation techniques. Early in the lifetime of these mitigations, Tyler Durden wrote a Phrack article [4] describing an attack using printf-based info leaks. He also describes the details of the PaX ASLR implementation. In [12], Shacham et al show that the current class of ASLR implementations are weak to a brute force remote attack. Fortunately, for a modern client-side exploit, it is rare to get a large number of trials against the same target reducing the probability of guessing the base address. Most recently, Mark Dowd and Alex Sotirov teamed up to present their research on browser based exploitation techniques for bypassing DEP and ASLR [13]. Besides a useful guide to the implementation of DEP and ASLR on Windows, they propose a few methods for bypassing the mitigations. The DEP bypasses involve the use of

Java applets or .NET assemblies to introduce large sections of executable memory into the address space.

The currently most effective exploitation technique against DEP is known as return-oriented-programming (ROP). Introduced in the litereature by Shacham [11] and further developed in industry by Immunity [7], return-oriented-programming reuses existing code segments to cobble together shellcode. The found pieces of code, "gadgets", are short segments of code ending in RET-urn opcodes. The ROP program is entered by pivoting the stack to an attacker controlled buffer (see [11, 3]).

Lastly, a few researchers have focused on interpreter based exploitation. Erik Cabetas presented some of his work at SOURCE Boston [2]. Two talks at BA-Con in 2008 involved interpreter exploitation; Aaron Portnoy and Ali Rizvi-Santiago presented their research on Python [9] and Justin Ferguson spoke about his JavaScript research [6].

6   Future Research and Conclusions

Through the examination of pointer inference and our JIT spray techniques, we have shown how the design and implementation of an interpreter may have surprising impacts with regard to security. As they become more valuable, further research is needed to bound information leaks which divulge information indirectly (such as the ActionScript Dictionary ordering described in Section 2). Formal methods for bounding this leakage would be an interesting research problem.

Interpreters provide a deceptively large amount of power to the attacker. Developing the JIT spray technique into a full attack relies on the implementation details of the memory subsystem (both the OS system and the application layer existing above that). Modifying this subsystem is one way to stop the spray technique, but further research should be devoted to understanding how much an attacker can influence the memory allocation subsystem. What attacks are possible given various allocation strategies and how do these application level algorithms impact performance?

Despite modern exploit mitigations such as address space layout randomization (ASLR) and data execution prevention (DEP), browser-based exploitation is still within reach of an attacker. Despite the lack of absolute security, current mitigations have a noticable effect on the reliability and complexity of modern exploits.

7   Acknowledgments

Many people helped out when I was developing these techniques. Thanks for all the discussion in formal. I'd also like to thank Dino Dai Zovi for his initial feedback and Manish Karir and Eric Wustrow for providing quality discussions on the dictionary attack. Matt Miller at MSEC and Peleus Uhley at Adobe provided useful discussions from the vendor point of view. Lastly, thanks to my reviewers for providing clear and honest feedback.

References

[1] Tamarin project. http://www.mozilla.org/projects/tamarin/.

[2] Erik Cabetas. Vulnerabilities in application interpreters and runtimes. In *SOURCE Boston*, 2009.

[3] Dino Dai Zovi. Practical return-oriented programming. In *SOURCE Boston*, 2010.

[4] Tyler Durden. Bypassing PaX ASLR protection. *Phrack*, 0x0b(0x3b):0x09, 2002.

[5] Sinan "noir" Eran. http://www.securityfocus.com/archive/82/327100/2009-02-24/1.

[6] Justin Ferguson. Advances in attacking interpreted languages: Javascript. In *BA-Con*, 2008.

[7] Immunity. DEPLIB.

[8] Microsoft. EMET. http://blogs.technet.com/srd/archive/2009/10/27.

[9] Aaron Portnoy and Ali Rizvi-Santiago. Reverse engineering dynamic languages, a focus on python. In *BA-Con*, 2008.

[10] Paruj Ratanaworabhan, Benjamin Livshits, and Benjamin G. Zorn. Nozzle: A defense against heap-spraying code injection attacks. Technical Report MSR-TR-2008-176, Microsoft Research, 2008.

[11] Hovav Shacham. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In Sabrina De Capitani di Vimercati and Paul Syverson, editors, *Proceedings of CCS 2007*, pages 552–61. ACM Press, October 2007.

[12] Hovav Shacham, Matthew Page, Ben Pfaff, Eu-Jin Goh, Nagendra Modadugu, and Dan Boneh. On the effectiveness of address-space randomization. In Birgit Pfitzmann and Peng Liu, editors, *Proceedings of CCS 2004*, pages 298–307. ACM Press, October 2004.

[13] Alex Sotirov and Mark Dowd. Bypassing browser memory protections in Windows Vista. In *Blackhat USA*, 2008.

A   JIT Spray Exploit Demo

A video of the full JIT spray attack exploiting an Adobe Reader flaw on Windows Vista running IE8 is available at http://www.youtube.com/watch?v=HJuBpciJ3Ao.