

# Pwning Adobe Reader

Abusing the Reader's embedded XFA  
engine for reliable Exploitation

Sebastian Apelt  
sebastian.apelt@siberas.de

# ■■■ Agenda

- whoami
- Motivation
- (Short!) Introduction to XFA
- XFA Internals
  - XFA Objects
  - jfCacheManager
- Exploiting the Reader
- Demo
- Conclusion
- Q&A



- Sebastian Apelt (@bitshifter123)
- Co-Founder of siberas in 2009
  - IT-Security Consulting (Pentests, Code Audits, etc.)
  - Research
- Low-level addict
  - Reverse Engineering, Bughunting, Exploitation
    - > 100 CVEs in all kinds of Products
    - Pwn2Own 2014 (IE11 on Win8.1 x64)

# ■■■ Motivation

# ■■■ Motivation

- Fuzzing at siberas
  - Let's pwn the Reader @ Pwn2Own 2016!!
    - Unfortunately, no love for Reader this time ☹️
  - In 2015: XFA fuzzing on 128 cores
  - Fuzz run yielded thousands of crashes
  - So far ~ 20 Bugs identified as unique (upcoming)
  - Analysis took ages...
  - Let's take a look at a typical Reader crash!

# ■■■ Motivation

(72fc.72ec): Access violation - code c0000005 (!!! second chance !!!)  
eax=69572c30 ebx=00000002 ecx=07b2f3cc edx=05658af8 esi=0549e538 edi=07b2f3cc  
eip=20a29654 esp=0031d8c4 ebp=00000003 iopl=0      nv up ei pl nz na  
cs=0023  ss=002b  ds=002b  es=002b  fs=0053  gs=002b      efl=00210206

AcroForm!DllUnregisterServer+0x2f73ce:

20a29654  mov edx,dword ptr [eax] ds:002b:69572c30=????????

**Awesome, we have a crash!**

**But no useful function name  
(DllUnregisterServer??)**

0:000> !heap -p -a ecx  
address 07b2f3cc found in  
\_HEAP @ 11a0000  
HEAP\_ENTRY Size Prev Flags    UserPtr    UserSize - state  
07b24eb0 199c 0000 [00] 07b24eb8 0ccd8 - (busy)

Offset 0xa514 !?

**The object holding the bad  
reference is located in the  
middle of a huge buffer  
=> Page Heap useless**

0:000> kc  
AcroForm!DllUnregisterServer+0x2f73ce  
AcroForm!DllUnregisterServer+0x2f7212  
AcroForm!DllUnregisterServer+0x2f7504  
AcroForm!DllUnregisterServer+0x35f3ae  
AcroForm!DllUnregisterServer+0x358f50

**Stacktrace also not helpful**

## ■■■ Motivation

- Adobe Reader => No symbols / RTTI infos!
  - No function names
  - No object / vtable information
  - No meaningful stacktraces
  - Page Heap useless
- Root cause analysis is very hard without context
- Complicates crash triaging during fuzz runs

## ■■■ Motivation

- How do we ANALYZE crashes in XFA?
- How do we EXPLOIT these crashes?
- Obvious: We need context! We need symbols!
- No *in-depth* research about XFA internals so far:
  - Most useful: Writeups about XFA exploit from 2013 (David and Enrique of Immunity Inc, Matthieu Bonetti of Portcullis Labs)
  - Good technical analysis, but only scratching the surface



## ■■■ Motivation

- Write tools to recover contextual information
  - Lower the bar for other researchers!
  - Check <https://github.com/siberas> in the next days
- Facilitate:
  - Vulnerability discovery and root cause analysis
  - Crash triaging during fuzz runs
- Deliver XFA-specific background for exploitation

# ■■■ (Short!) Introduction to XFA

# ■■■ (Short!) Introduction to XFA

- XFA: „XML Forms Architecture“
  - Specification developed by JetForm, later Accelio (acquired by Adobe in 2002) – not a standard
  - Latest version: 3.3 (01/2012): Easy read of 1584 pages.
  - Brings *dynamic* behavior to the *static* PDF world: Forms that can dynamically change their layout!
  - Dynamic nature of XFA is powered by Javascript (Spidermonkey 24 since AR DC)
  - XFA not supported by many PDF Readers, yet (Chrome/Chromium, Firefox, Windows,...)

## ■■■ (Short!) Introduction to XFA

- XFA form data itself is an XML-structure embedded in the PDF, a so-called *XDP*-Packet
- Javascript embedded in this XDP
  - Executed upon events (e.g. document is fully loaded, user clicks on button, etc.)
- A practical example...



# (Short!) Introduction to XFA

```
<xdp:xdp xmlns:xdp="http://ns.adobe.com/xdp/">
  <config xmlns:xfa="http://www.xfa.org/schema/xci/3.0/">
    [...]
  </config>
  <template xmlns:xfa="http://www.xfa.org/schema/xf-template/3.0/">
    <subform layout="tb" name="form1">
      <pageSet>
        <pageArea id="PageArea1" name="PageArea1">
          <contentArea w="612pt" h="792pt" x="20pt" y="20pt"/>
        </pageArea>
      </pageSet>
      <field name="button1" w="41.275mm" h="9.525mm">
        <ui>
          <button highlight="inverted"/>
        </ui>
        [...]
        <event activity="click" name="event__click">
          <script contentType="application/x-javascript">
            app.alert(1337);
          </script>
        </event>
        [...]
      </field>
    </subform>
  </template>
</xdp:xdp>
```

XDP Packet is XML embedded in the PDF  
The root tag is always „xdp“

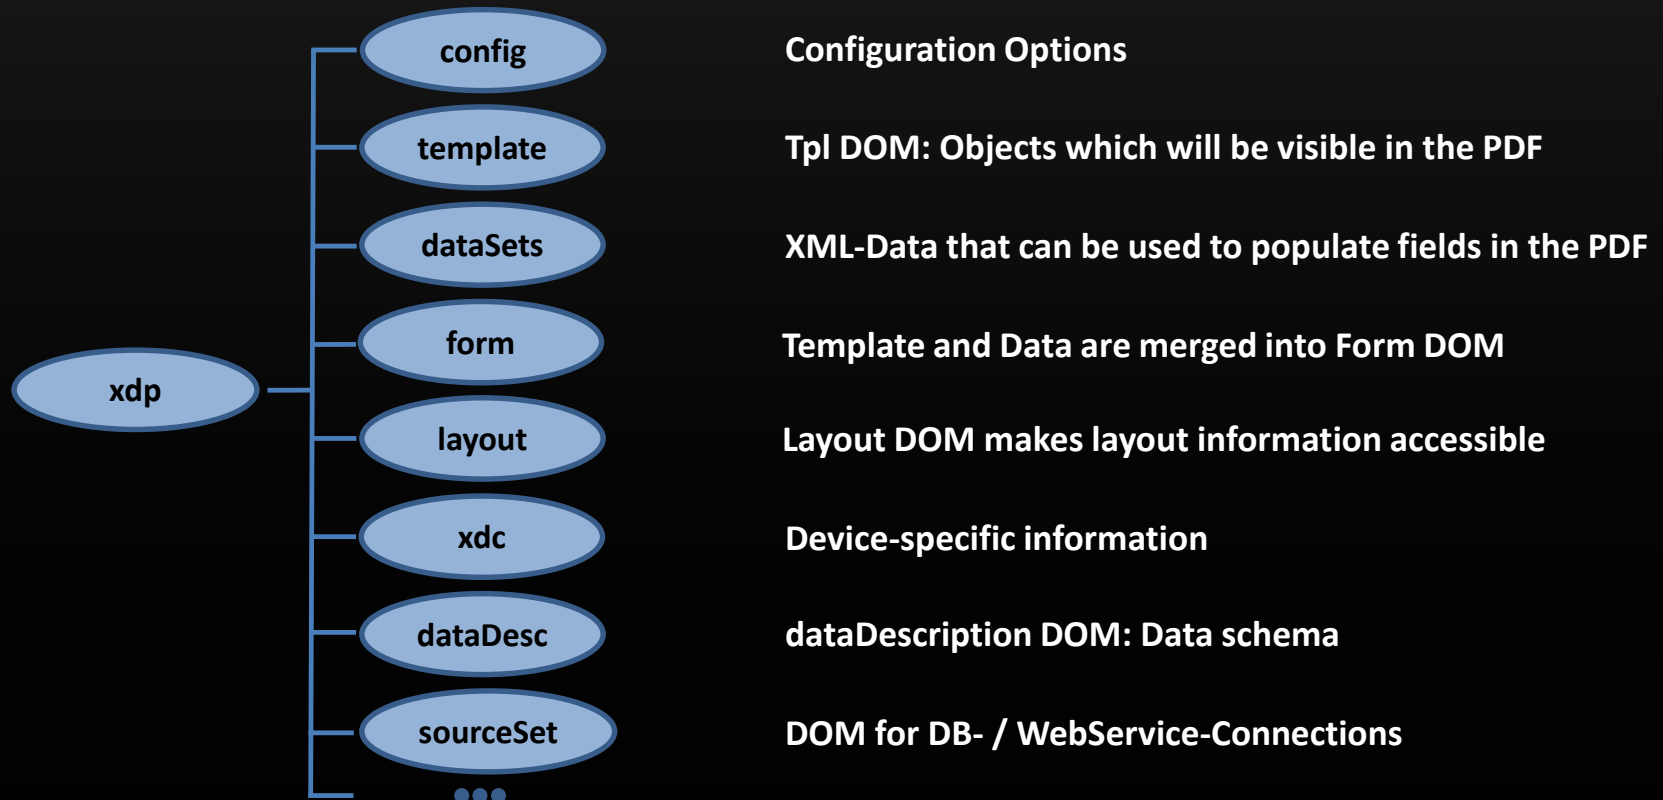
Config DOM contains configuration  
options for XFA processing

Template DOM is structured in subforms,  
containing objects like „field“, „text“, etc.

Objects can contain event objects that fire  
on certain actions (e.g. „click“)

# ■■■ (Short!) Introduction to XFA

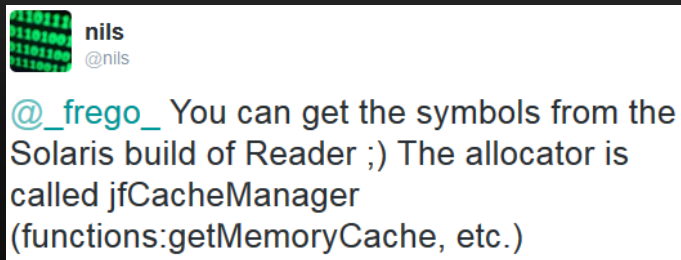
- XFA spec defines multiple DOMs
  - HUGE attack surface (> 200 objects accessible via JS)



# ■■■ XFA Internals

# ■■■ XFA Internals - General Approach

## ■ Tweet by @nils



- Nice! Some Solaris build seems to have symbols!
  - Newest version which still has symbols: Solaris v9.4.1
- We need a *reliable* heuristic to port symbols in AcroForm.api (module which implements XFA functionality) to newer AR versions



# ■■■ XFA Internals - General Approach

## ■ Problems:

- Code is rather old (2012) -> Many Code changes from v9.X to AR DC...
  - Function count: Solaris ~48 K, AR DC ~ 95 K
- Functions differ even if code stays the same (compiler optimizations like heavy inlining in v9.4.1 screw it up)
  - Tried diffing with Diaphora – Too many false positives
- Structures, objects and vtable sizes differ (slightly, but enough to make it very hard to create reliable heuristics)
- etc.

## ■■■ XFA Internals - General Approach

- Approach: Trying to *understand* Reader v9.4.1 as much as possible with the help of symbols
- Find bulletproof ways to recover the *most important* symbols, i.e.
  - Heap Mgmt functions for the custom allocator
  - Object information

# ■■■ XFA Internals - Objects

- What do we need to know about objects?
  - How to identify an object in memory
  - Vtable offsets
  - Methods and properties exposed to JavaScript
  - Offsets of the entrypoints for methods / property-getters and -setters
  - Function names of vtable entries

# ■ ■ ■ XFA Internals - Objects: Identification

## ■ First attempt: XFANode::getClassTag

```
; _DWORD __cdecl XFANode::getClassTag(XFANode *this)
public _ZNK7XFANode11getClassTagEv
_ZNK7XFANode11getClassTagEv proc near

this= dword ptr 4

mov     eax, [esp+this]
mov     eax, [eax]
mov     eax, [eax+10h]
retn

_ZNK7XFANode11getClassTagEv endp
```

classTag attribute can be found @ <XFAobj> + 0x10

```
mov     eax, ds:(_ZTV12XFAFieldImpl_ptr - 0D91324h)[ebx]
add     eax, 8
mov     [esi], eax
mov     eax, ds:(aXFA_FIELD_ptr - 0D91324h)[ebx]
mov     [eax],
mov     [esi+0Ch], eax
mov     dword ptr [esi+10h], 86h
add     esp, 10h
pop     ebx
pop     esi
leave
retn
```

From Field constructor method:  
classTag for Field-Object in  
Adobe Reader 9.4.1: 0x86

```
0:011> dd 0x6883d74
06883d74  6822bd94 00000001 00000000 683412fc
06883d84  0000008e 00000002 00011952 00000000
06883d94  00000000 00000000 068765ec 00000002
06883da4  04c39bcc 04d950a0 00000000 00000000
```

classTag for Field-Object in  
Acrobat Reader DC: 0x8e

## ■ Fail! classTags not constant across versions! 🤔

# ■ ■ ■ XFA Internals - Objects: Identification

- <XFAObj>::Type method to the rescue
- Located @ vtable+8 of each XFA-Object

**Adobe Reader 9.4.1**

```
; _DWORD XFAFieldImpl::Type(XFAFieldImpl * __hidden this)
public _ZNK12XFAFieldImpl4TypeEv ; weak
_ZNK12XFAFieldImpl4TypeEv proc near
mov     eax, 7C46h
retn
_ZNK12XFAFieldImpl4TypeEv endp
```

**Acrobat Reader DC**

```
0:011> uf poi(poi(0x6883d74)+8)
AcroForm!DllUnregisterServer+0x34020a:
67fc2490 b8467c0000 mov     eax, 7C46h
67fc2495 c3          ret
```

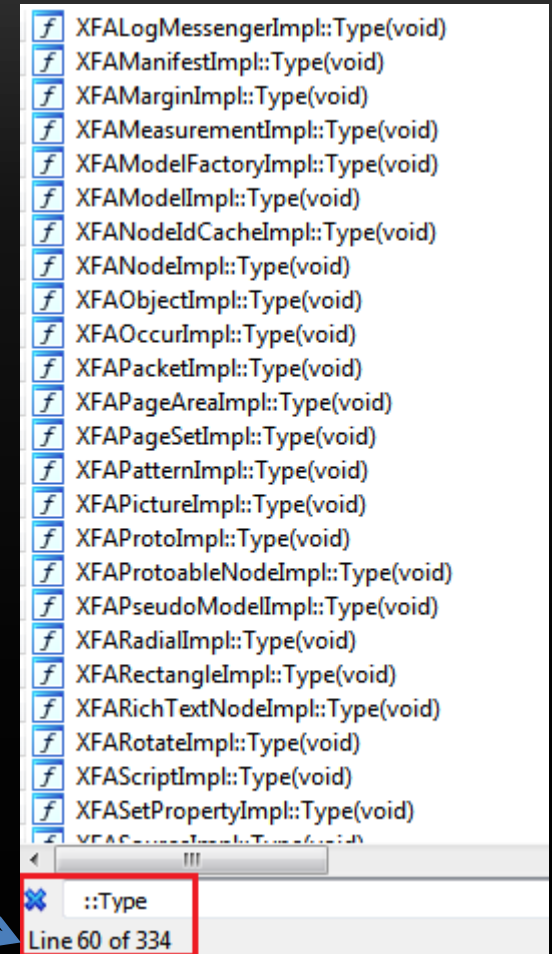
Type is 0x7C46 for both v9.4.1  
AND Acrobat Reader DC! ☺

- Type-IDs are static across versions!

# ■■■ XFA Internals - Objects: Identification

- Possible to identify every object by a binary pattern in newer versions of AcroForm.api
  - `mov eax, 7C46h`  
`retn`  
 $\Leftrightarrow$  `B8 46 7C 00 00 C3`
- Xref to the Type method gives us the vtable offset (RVA) to each object!

We can safely identify 334 objects! Not too bad!



# ■■■ XFA Internals - Objects

- What do we need to know about objects?
  - How to identify an object in memory ✓
  - Vtable offsets ✓
  - Methods and properties exposed to JavaScript
  - Offsets of the entrypoints for methods / property-getters and -setters
  - Function names of vtable entries

## ■ ■ ■ XFA Internals - Objects

- How about methods and properties?
- <XFAObj>::getScriptTable() @ vtable offset 0x34

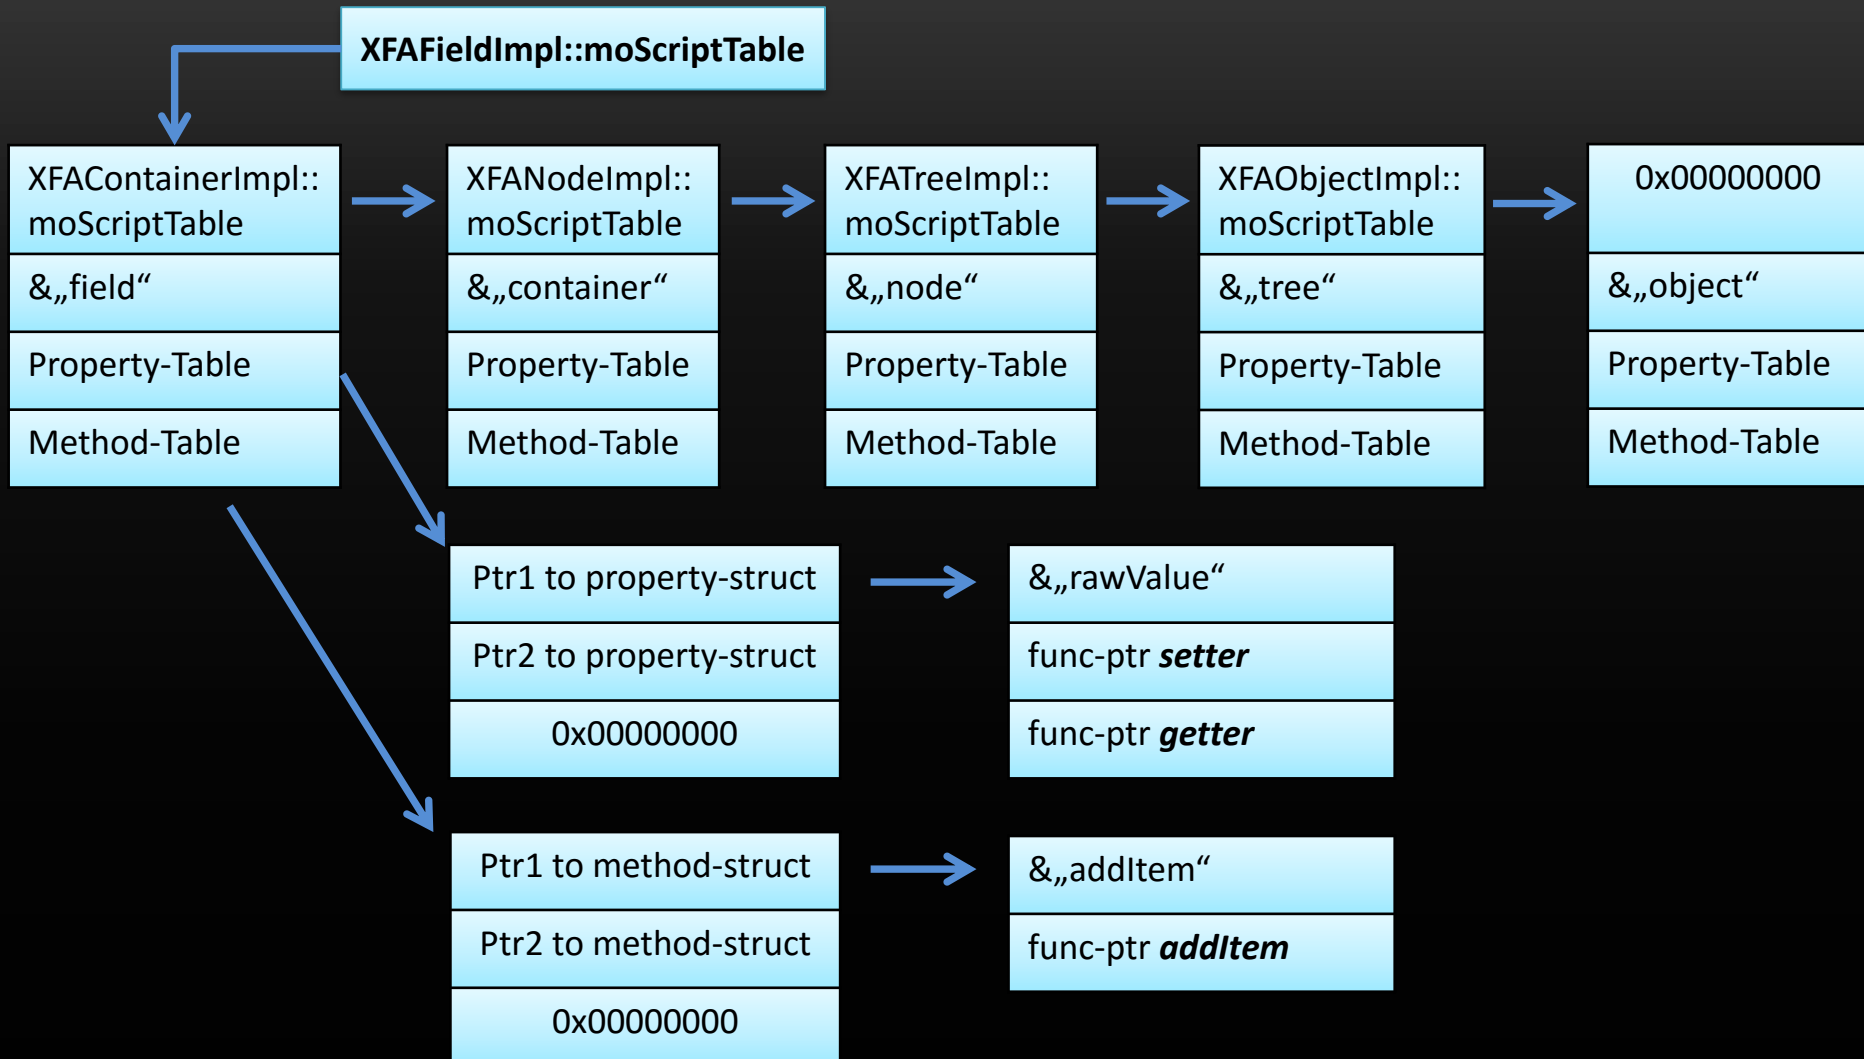
```
; _DWORD XFAFieldImpl::getScriptTable(XFAFieldImpl *__hidden this)
public _ZNK12XFAFieldImpl14getScriptTableEv
_ZNK12XFAFieldImpl14getScriptTableEv proc near
call    $+5
pop     ecx
add     ecx, 65F419h
mov     eax, ds:(_ZN12XFAFieldImpl13moScriptTableE_ptr - 0D91324h)[ecx]
retn
_ZNK12XFAFieldImpl14getScriptTableEv endp
```

XFAFieldImpl::moScriptTable

- References *moScriptTable* structure
  - Structure contains information about method and property names, function pointers, etc.



# ■■■ XFA Internals - Objects



# ■■■ XFA Internals - Objects

- What do we need to know about objects?
  - How to identify an object in memory ✓
  - Vtable offsets ✓
  - Methods and properties exposed to JavaScript ✓
  - Offsets of the entrypoints for methods / property-getters and -setters ✓
  - Function names of vtable entries

←  
TODO...  
Not trivial... ;-(

## ■■■ XFA Internals - jfCacheManager

- Most allocations in AcroForm.api are managed by a custom allocator called *jfCacheManager*
- LIFO-style heap manager
- Data buffers („blocks“) stored in big heap „chunks“
- Introduced most likely for performance reasons
- No security features...
  - No Heap Isolation (see IE, Flash, etc.)
  - No Anti-UAF like MemProtect/MemGC
  - ...

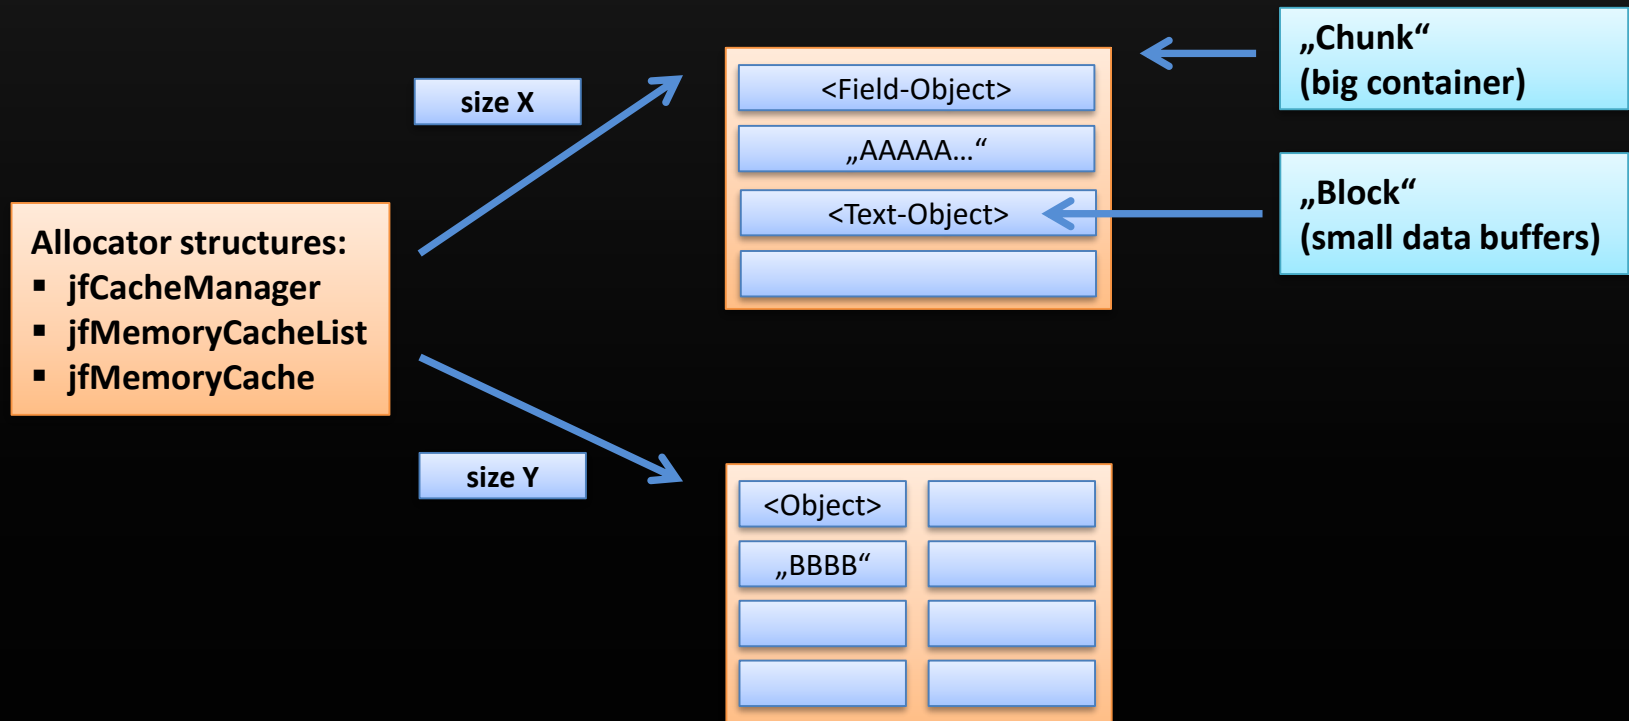
## ■■■ XFA Internals - jfCacheManager

Disclaimer: Next slides will only cover the *relevant* details of the memory manager in terms of *exploitation*!

(More in-depth analysis will be covered by a paper which will be released soon)

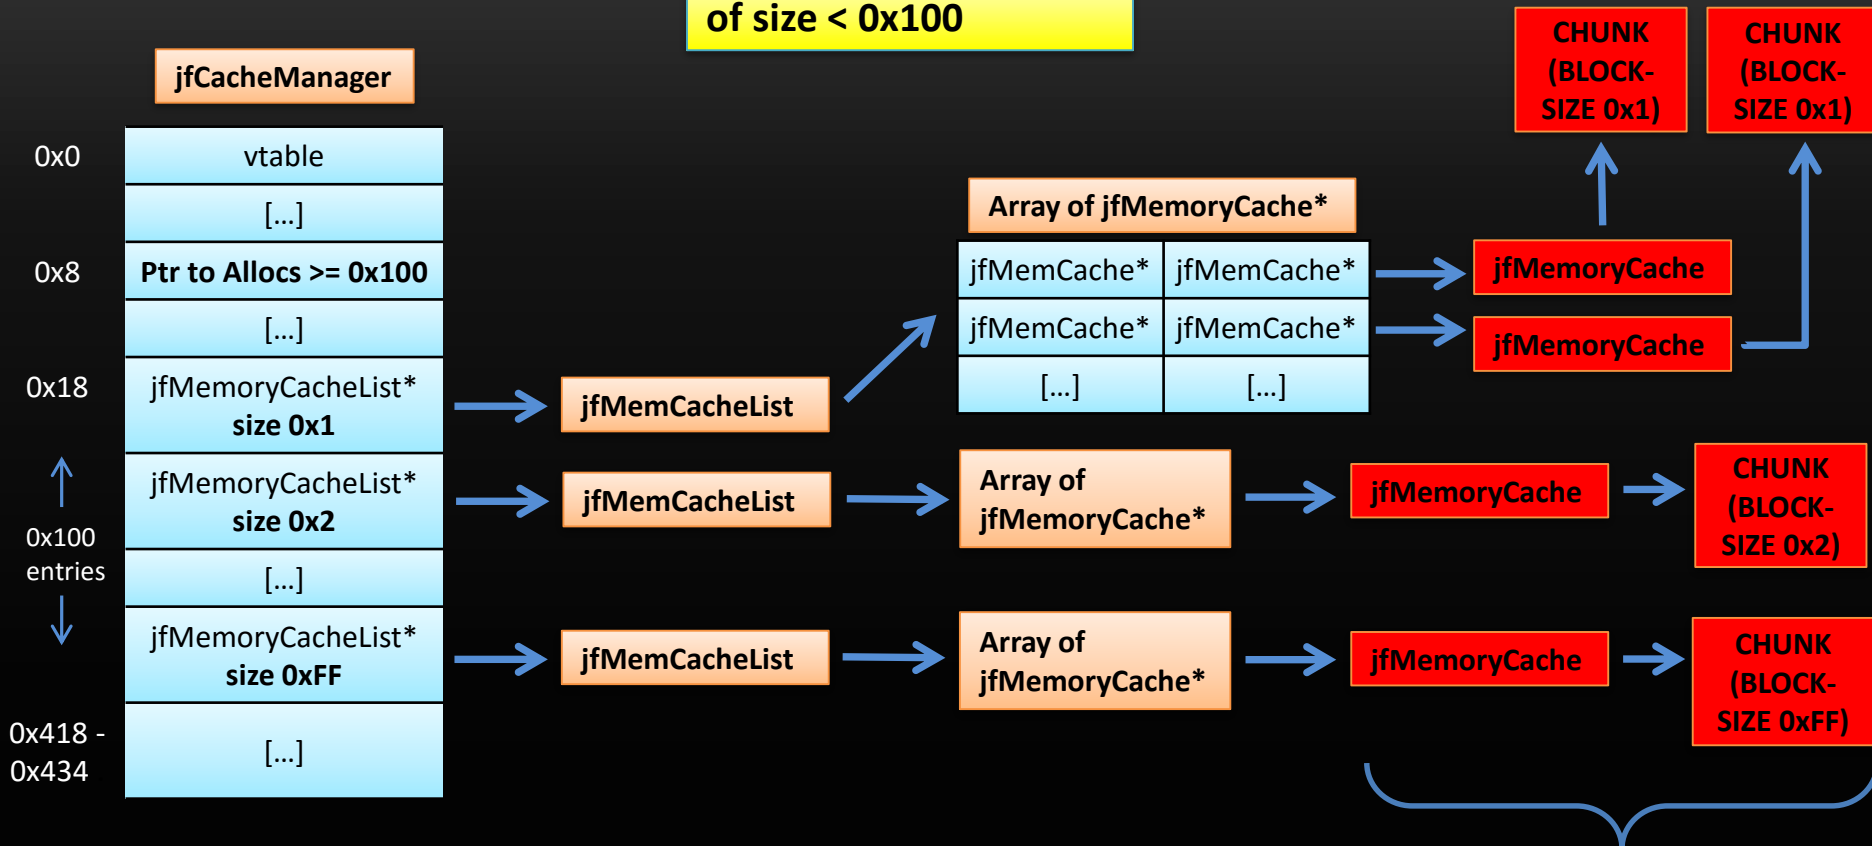
# ■■■ XFA Internals - jfCacheManager

- Very simplified version of the jfCacheManager:



# XFA Internals - jfCacheManager

Storage of allocations  
of size < 0x100



*jfMemoryCache* and the *chunks* will be relevant for exploitation!

## ■■■ XFA Internals - jfCacheManager

- *sizeof(chunk)* derived from block size:

```
base_size = 0xc350 // 50.000  
chunksize = (((size + 3) / 4) + 1) * ((base_size + size - 1) / size) * 4
```

Example: allocation size = 0x64

=> chunksize = 26 \* (0xc3b3 / 0x64) \* 4 = 0xcb20

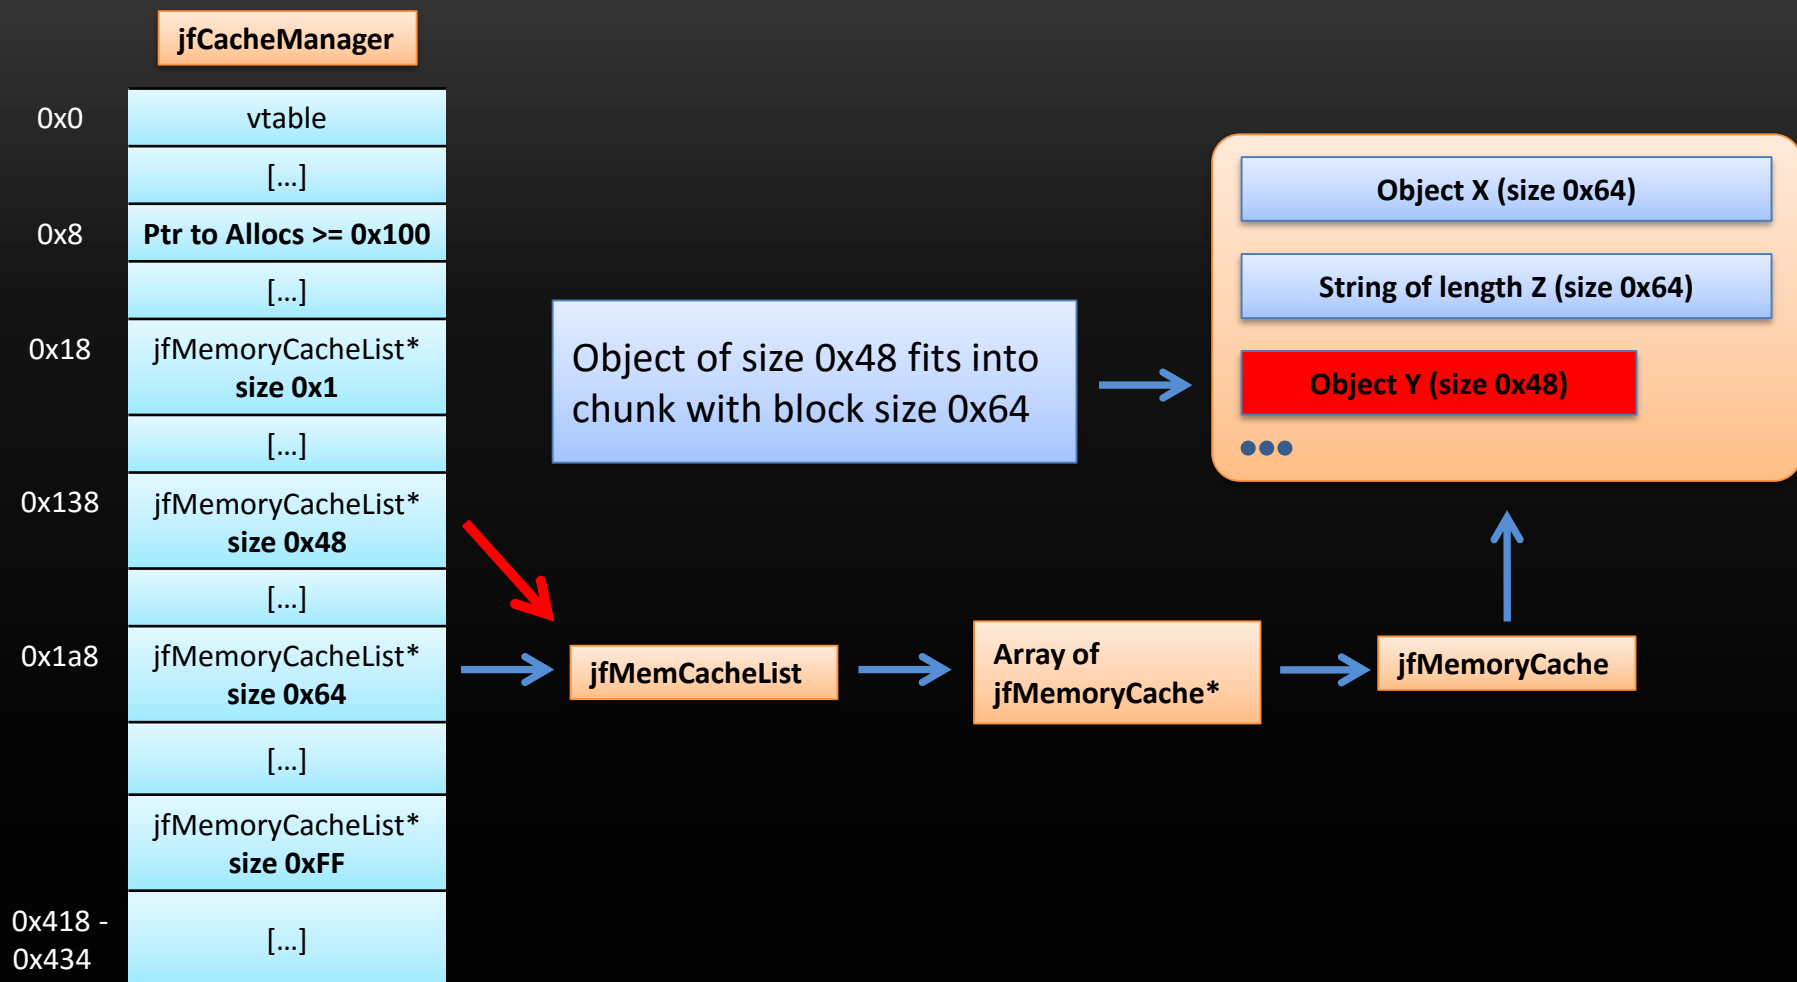
- „So, if I get a crash and I see my object located in a chunk of size 0xcb20, then sizeof(obj) == 0x64?“
  - Unfortunately not...

## ■■■ XFA Internals - jfCacheManager

- jfMemoryCacheLists can manage blocks of *multiple* sizes  
=> blocks of sizes X and Y can both end up in chunk Z!
- alloc(X) will be placed in same chunk as alloc(Y) if
  - an allocation for a size  $Y > X$  has occurred before and
  - size X is in the same „range“ as size Y
    - Ranges reach from  $2^n$  to  $(2^{n+1}-1)$  (e.g. 0x20 - 0x3f, 0x40 - 0x7f)
- In short:
  - Does the new block fit into some chunk that we already have?
  - If yes, use that chunk instead of allocating a new one!



# XFA Internals - jfCacheManager




## ■■■ XFA Internals - jfCacheManager

- Let's take a look at the structures within the chunks and what happens during alloc / free operations...

# XFA Internals - jfCacheManager

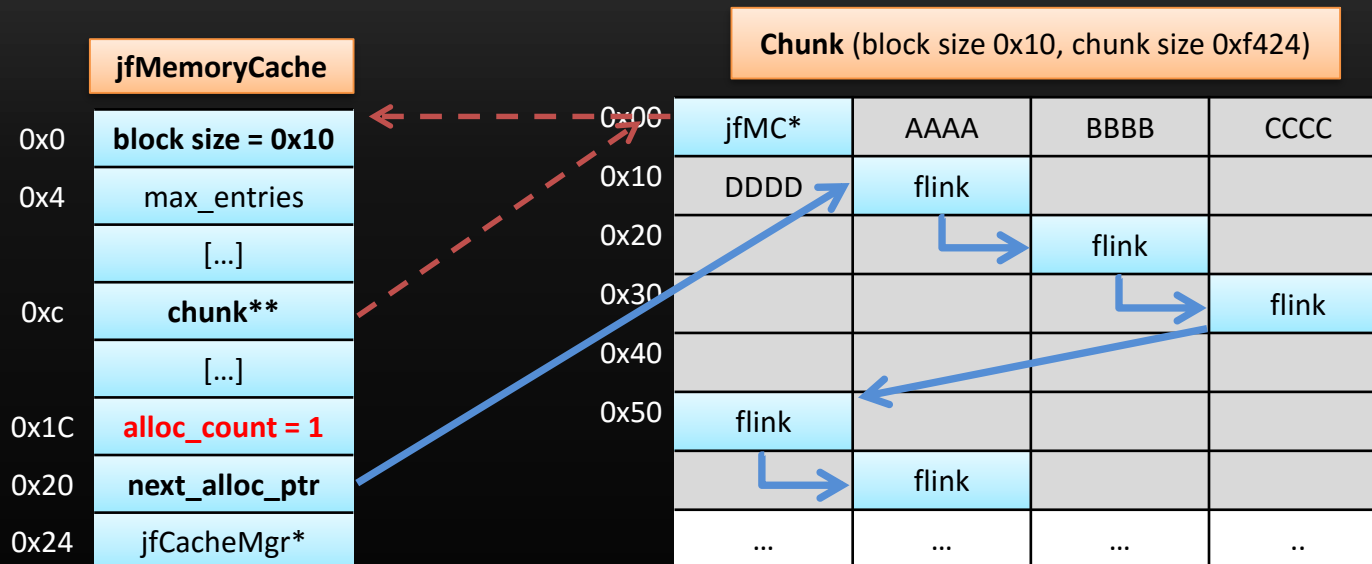
Initial state – All blocks are free



- **next\_alloc\_ptr** points to the block which will be returned with the next allocation
- **flinks** form a single linked list separating the data blocks

# XFA Internals - jfCacheManager

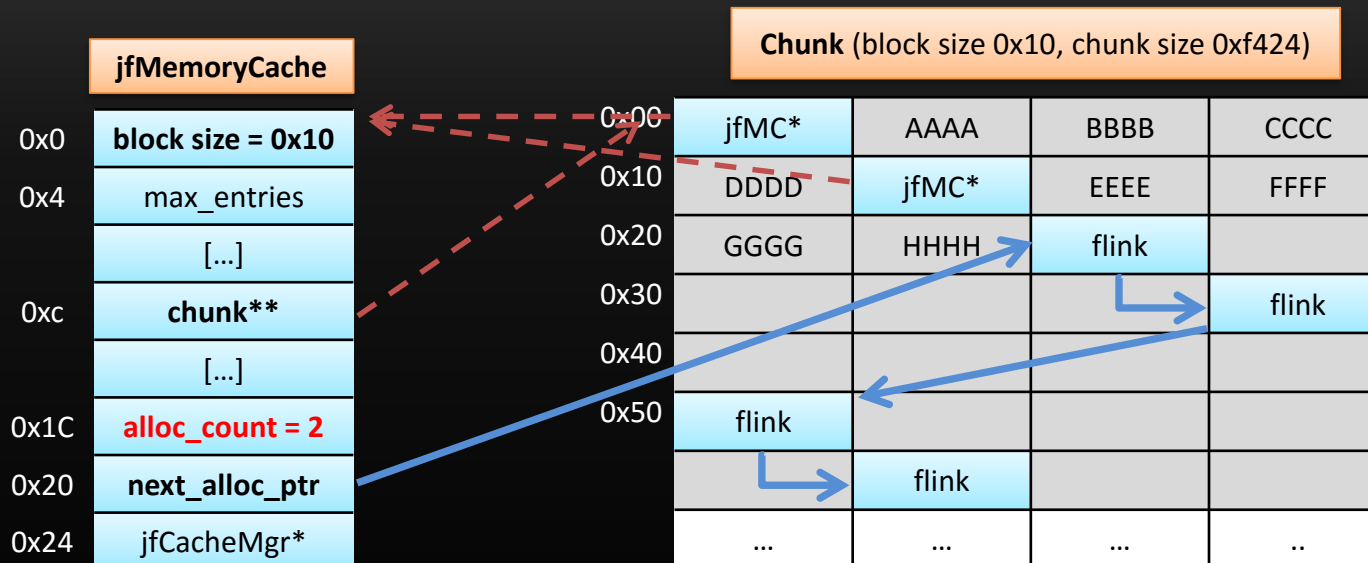
After first allocation



- *next\_alloc\_ptr* is overwritten with flink
- *flink* is overwritten with pointer back to jfMemoryCache
- *allocs\_counter* is incremented to 1

# XFA Internals - jfCacheManager

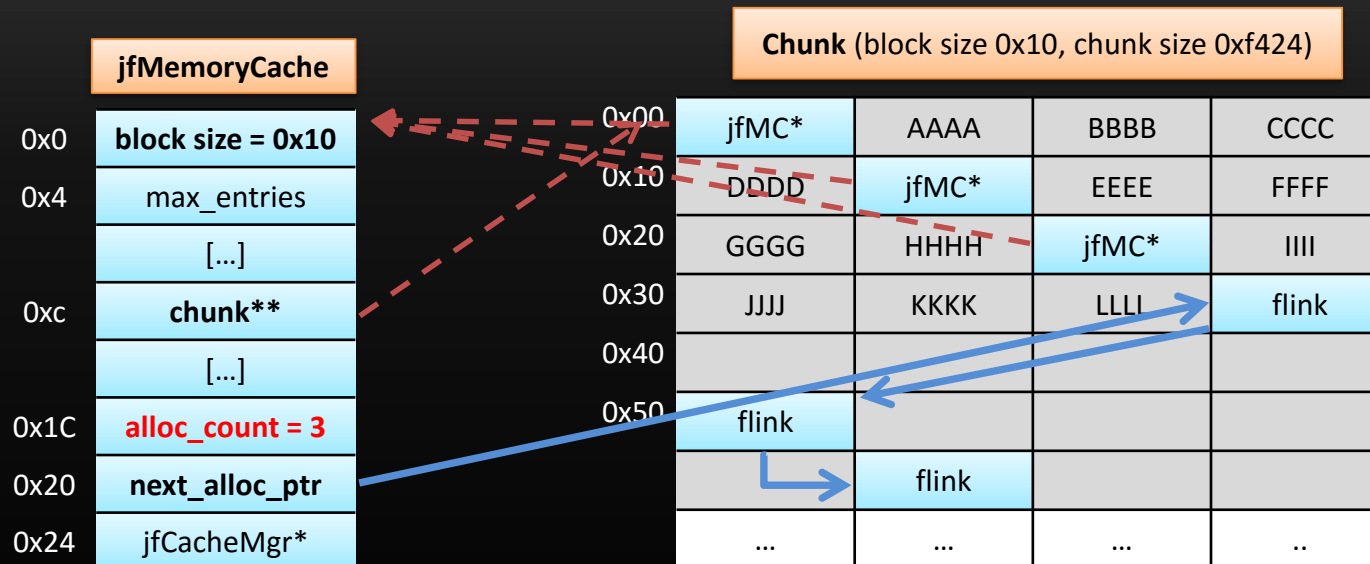
After second allocation



- *next\_alloc\_ptr* is overwritten with flink
- *flink* is overwritten with pointer back to jfMemoryCache
- *allocs\_counter* is incremented to 2

# XFA Internals - jfCacheManager

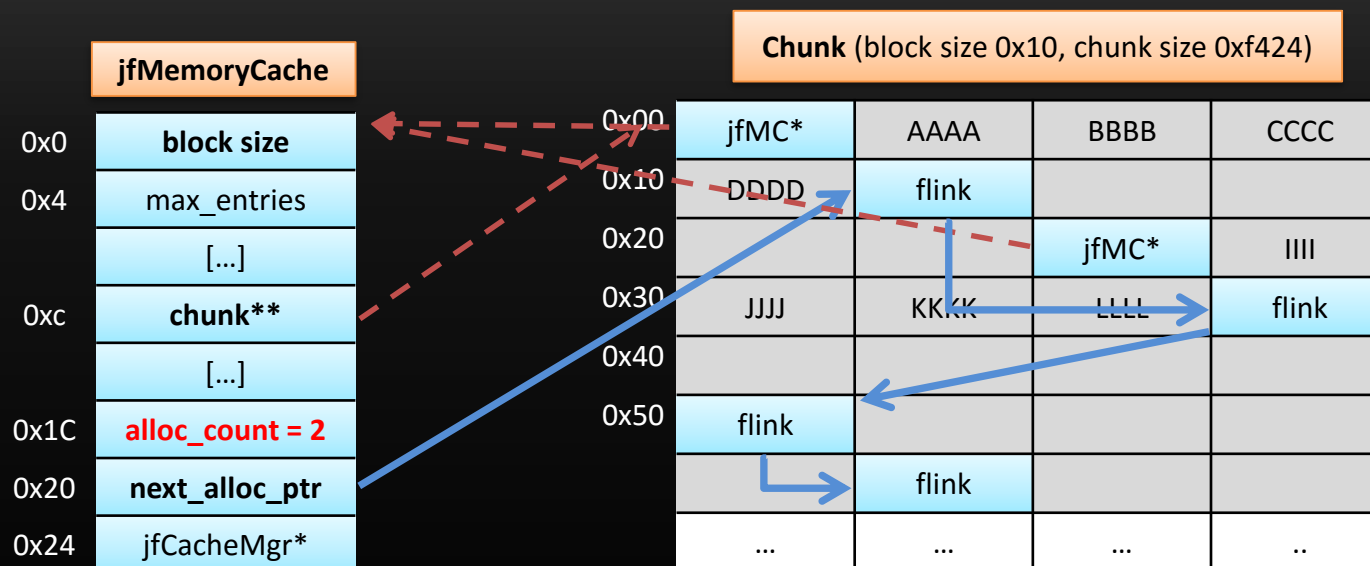
After third allocation



- *next\_alloc\_ptr* is overwritten with flink
- *flink* is overwritten with pointer back to jfMemoryCache
- *allocs\_counter* is incremented to 3

# XFA Internals - jfCacheManager

Free second block



- *next\_alloc\_ptr* is overwritten with pointer to free block - 4
- *jfMC\** is overwritten with *next\_alloc\_ptr* (becomes flink again)
- *allocs\_counter* is decremented to 2


## ■■■ XFA Internals - jfCacheManager

- Still don't like the jfCacheManager?
  - Still missing Page Heap?
- 
- Get offset „jfCacheManager\_active“ with XFAAnalyze\_funcs.py
  - Change byte from 1 to 0 in binary
  - Replace original AcroForm.api
  - You just switched off the jfCacheManager :P



[illegible]

# ■■■ Exploiting the Reader



## ■ Goals

- Bypass ASLR by corrupting specific byte(s) to cause a memory leak
- Find „flexible“ overwrite target
  - No need for a write-what-where (e.g. 0-DWORD write or a partial overwrite to a controlled address should suffice!)
- Find technique which is fast, reliable and most importantly independant from OS and AR version

# Exploiting the Reader

- Let's target the metadata contained within the chunks!
- Two possibilities:

**Chunk**

0x00	jfMC*	61616161	61616161	61616161
0x10	61616161	<b>flink</b>		
0x20			<b>jfMC*</b>	63636363
0x30	63636363	63636363	63636363	<b>flink</b>
0x40				
0x50	<b>flink</b>			
		<b>flink</b>		
	...	...	...	...

## Hit a flink

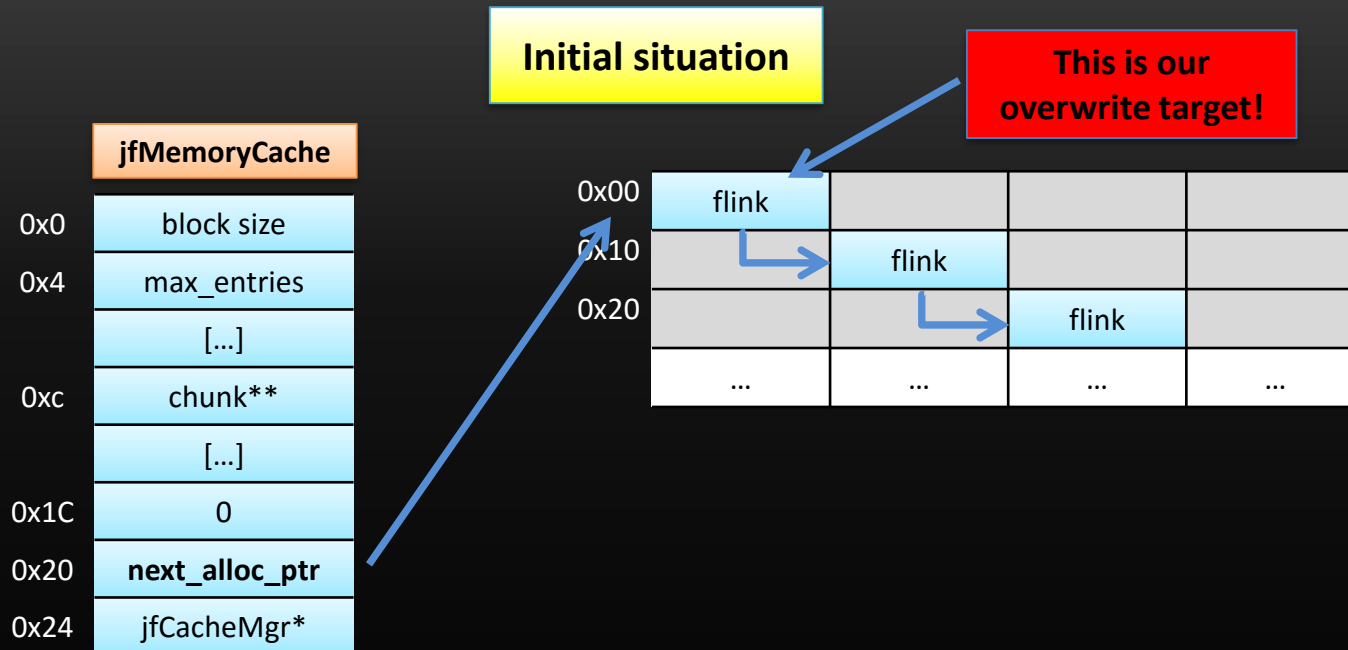
- ⇒ Block is *free*
- ⇒ Triggers when block is allocated

## Hit the jfMemoryCache\*

- ⇒ Block is *allocated*
- ⇒ Triggers when block is freed

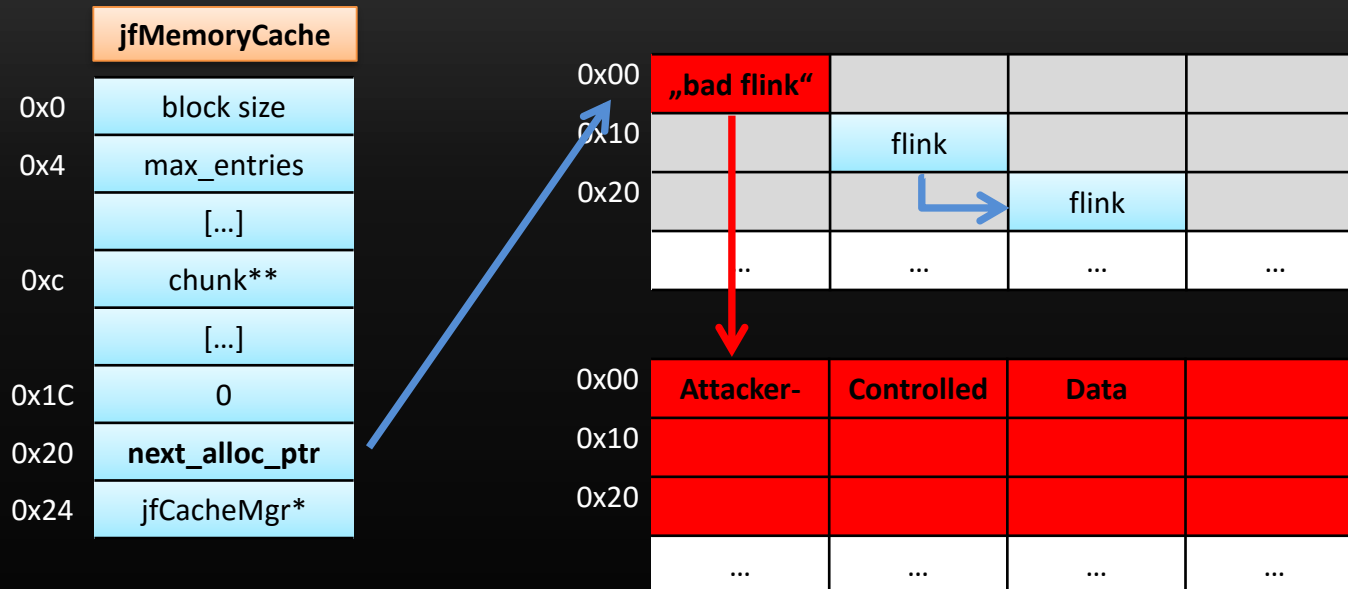
- Both methods can be abused create a memory leak!  
But hitting the *flink* is the easiest way to go

# Exploiting the Reader - Hit the flink!



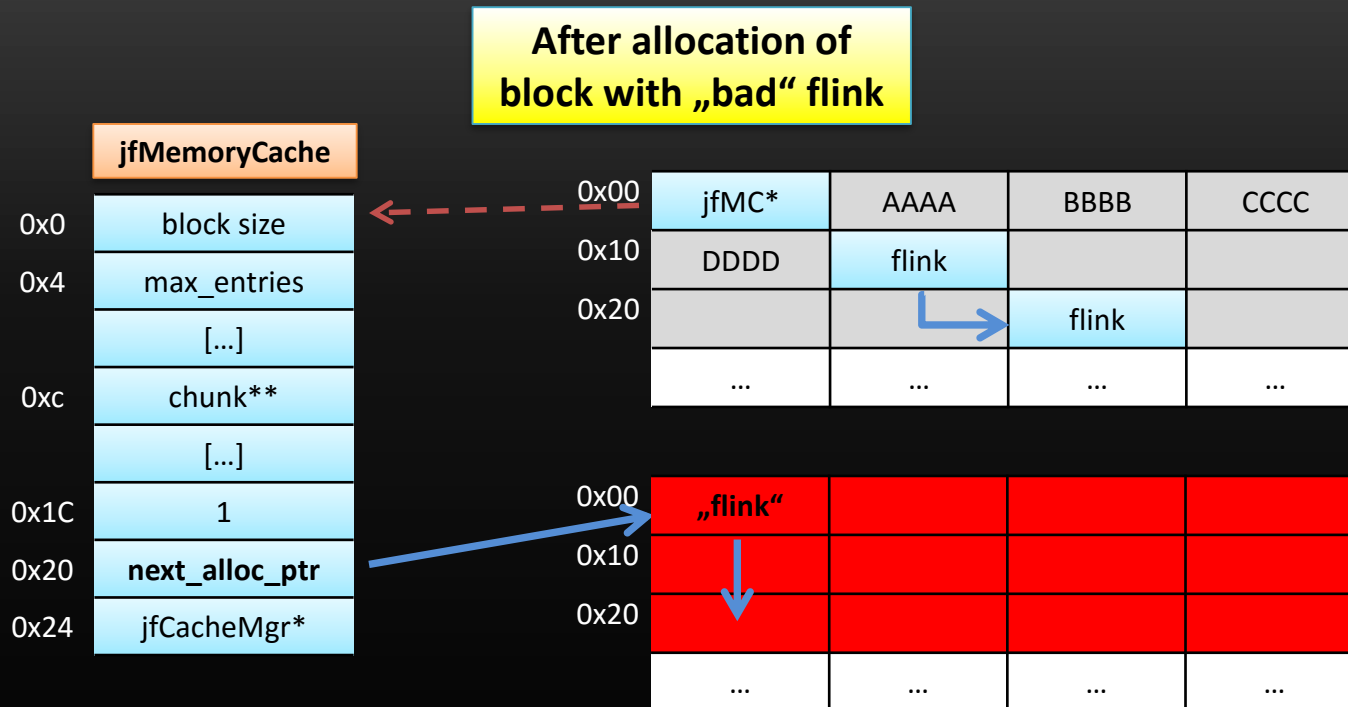
# Exploiting the Reader - Hit the flink!

After flink overwrite



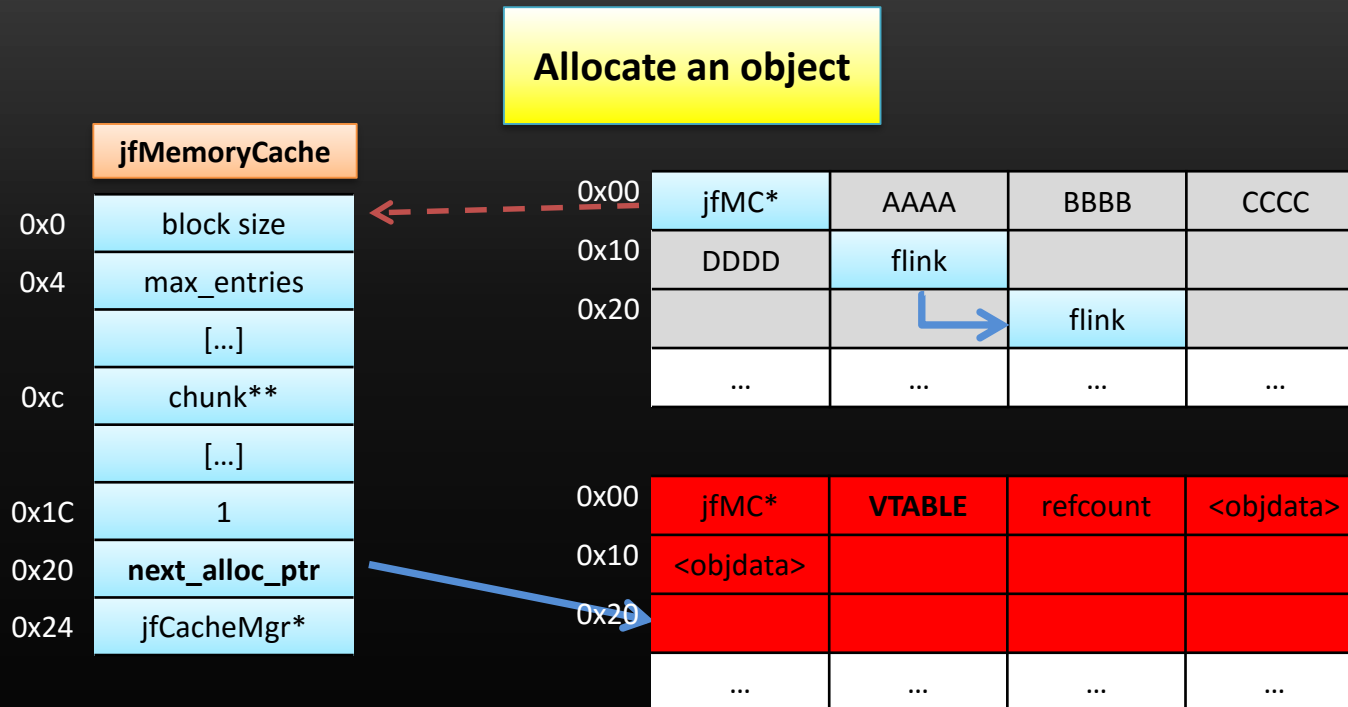
- Requirement: flink must point to controlled data after overwrite
- Still very flexible: Doable with nearly any kind of mem corruption!
- Let's see what happens when we allocate the „bad“ block

# Exploiting the Reader - Hit the flink!



- *next\_alloc\_ptr* is overwritten with the „bad“ flink
- *flink* is overwritten with pointer back to jfMemoryCache
- Now what happens when we allocate an object of size 0x10...?

# Exploiting the Reader - Hit the flink!



- Next allocation will return the data buffer after the „flink“
- The object will be placed in the middle of our controlled data  
=> We get a vtable in controlled data!!

## ■■■ Exploiting the Reader - Hit the flink!

- As soon as the vtable is in a controlled area you can just read it out
- The controlled data area can be sprayed with strings or even float arrays as „landing zone“
- Set the overwritten float or replace the string with data which will point to your ROP pivot gadget
- For floats: You can compute their binary representation after spec IEEE754:
  - 4.18356164518379836860971488084E-216 will be 0x13371337deadc0de on the heap
- GAME OVER!



# ■■■ Exploiting the Reader

Let's have a look at a practical example...

Exploitation of a 0-DWORD write has been presented @ SyScan360  
Check out my slides if you're interested ;)

Setting:



~~A 0 DWORD write primitive to an arbitrary address~~

## ■■■ Exploiting the Reader

- Let's make it harder than 0-DWORD overwrite
- For Infiltrate: Let's exploit ZDI-CAN-3507
- Originally planned for Pwn2Own 2016...
- Obvious: I can't reveal any information about the bug
- But I can describe the exploit methodology ☺
  - At least the basic steps
- WARNING:
  - The bug is ugly...
  - But: That makes it a great example to showcase the flexibility of the described flink overwrite technique!

# ■■■ Exploiting the Reader - ZDI-CAN-3507

Setting:

Write primitive of an *object-pointer* (non-XFA)  
to an arbitrary address

We can only write to an address  
where we have a 0-DWORD

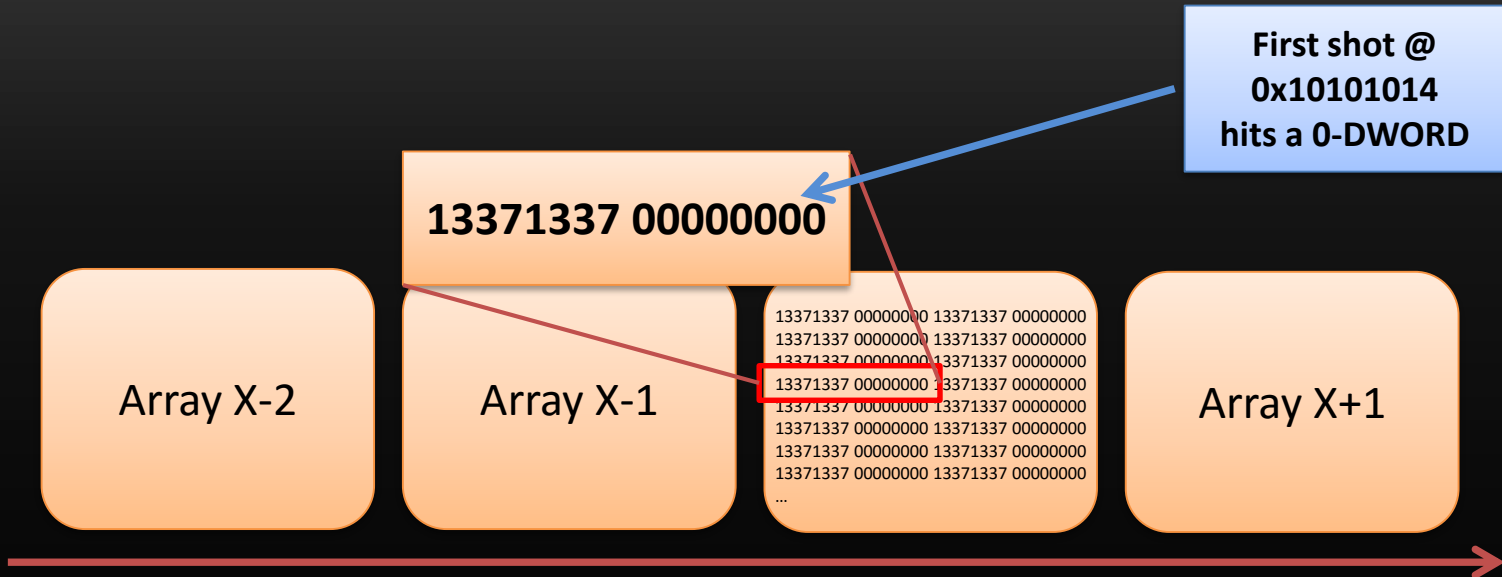


```
cmp [ecx], 0 // ecx is under control!  
jnz <no_write>  
*ecx = alloc_some_nonXFA_object()
```

## ■■■ Exploiting the Reader - ZDI-CAN-3507

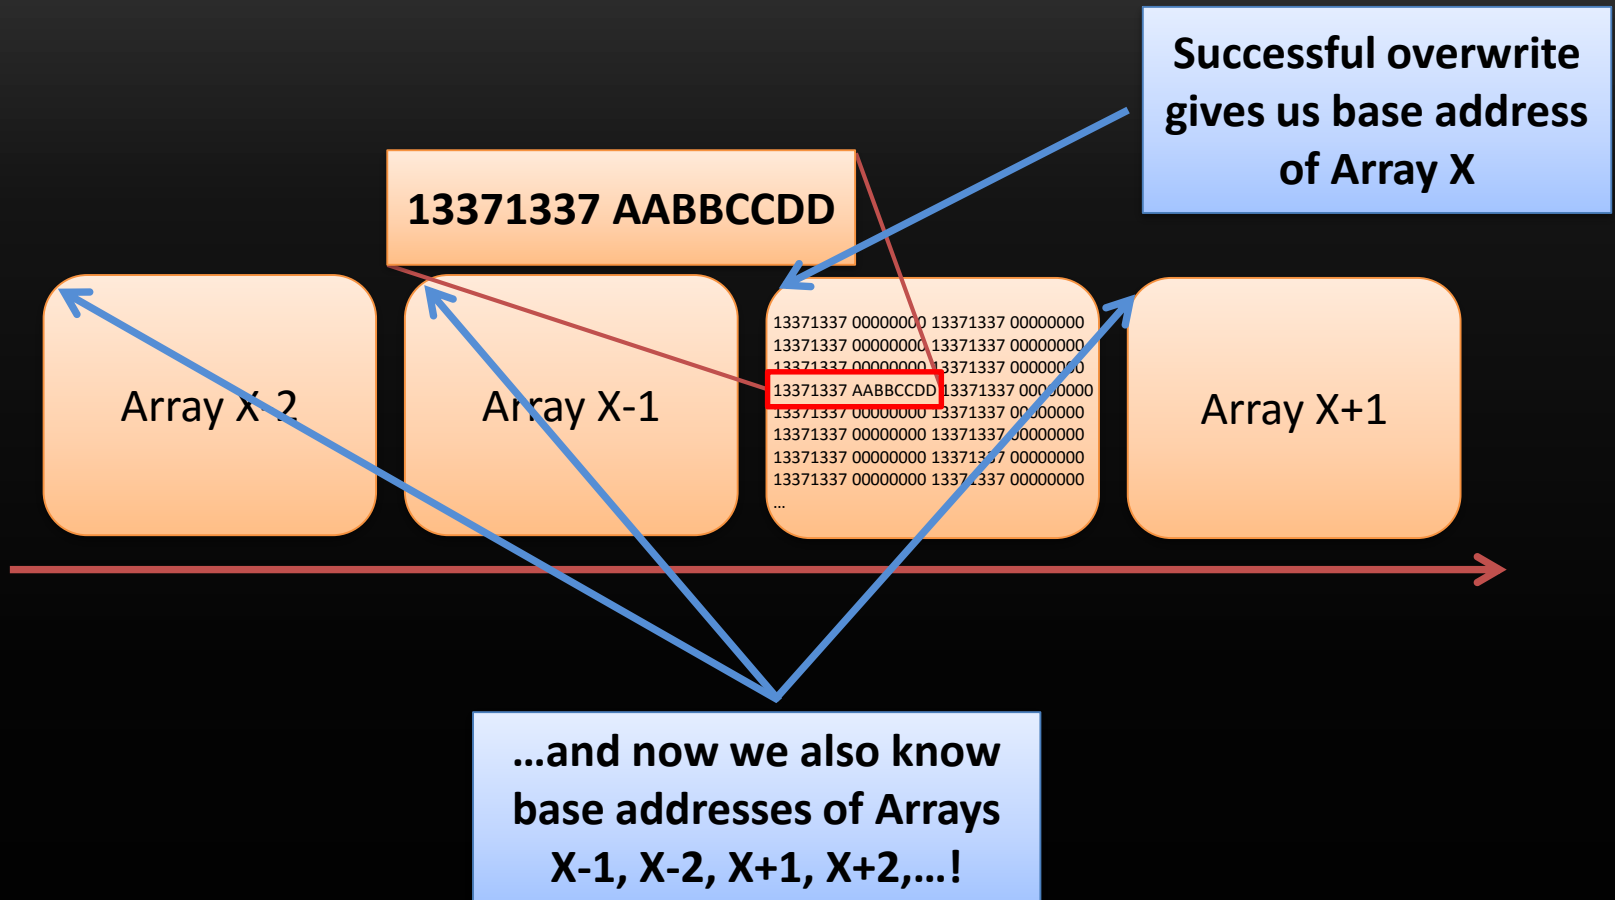
- Plan: Bypass ASLR by only triggering the vuln *twice*
  - First shot to derive information about the heap layout
  - Second shot to attack the flink
- First part is easy: Hit floating point arrays!
  - We can't shoot into heap spray of strings: No 0-DWORD...
  - Push value 1.59275155158737554072477261984e-315 into arrays => Results in binary pattern (after spec IEEE754)

```
13371337 00000000 13371337 00000000  
13371337 00000000 13371337 00000000 ...
```



- First shot will go to 0x10101014, this will be mapped by the array heap spray

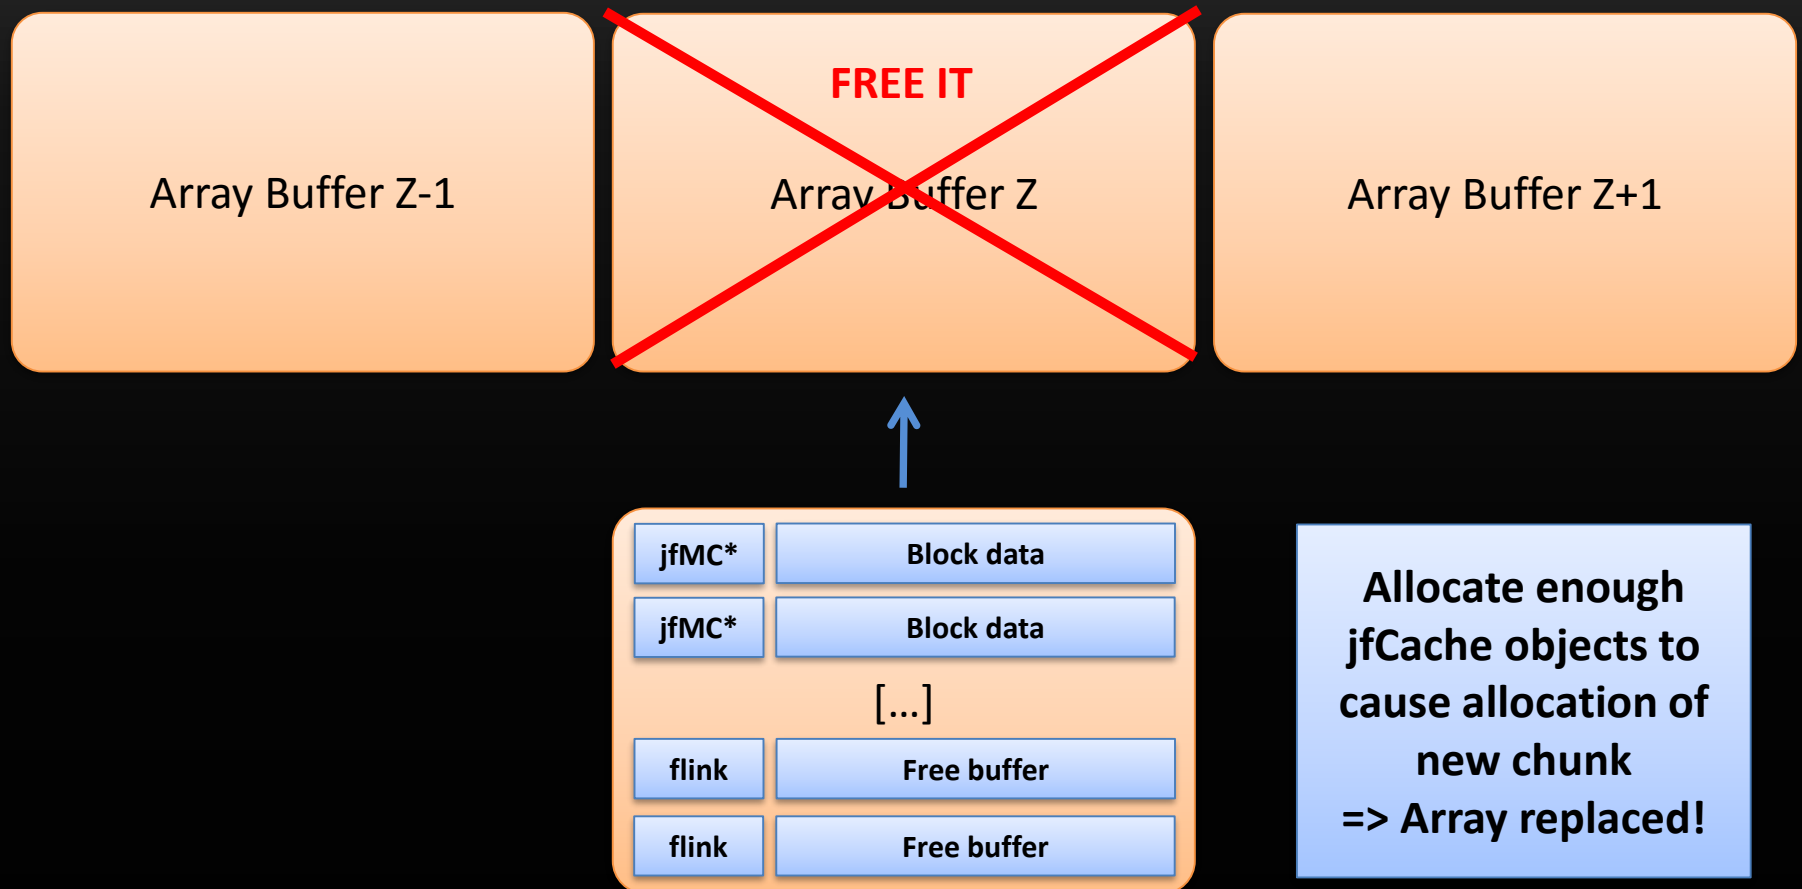
# Exploiting the Reader - ZDI-CAN-3507



# ■■■ Exploiting the Reader - ZDI-CAN-3507

- Now we need to overwrite a flink
  - A flink is an address, obviously  $\neq 0$ , but we can only write to an address where we have a 0-DW...
- Solution: Partial overwrite a flink which ends on 00's!
  - Let's manipulate the flink so that it is shifted into a neighboring float array!
  - When an object allocation with the „bad flink“ occurs, the object (and hence the vtable) is placed into the float array
- So how do I know where my flinks are in memory?
- And how do I know in where I can find the chunk that contains the flink ending on 00's (our target flink)??

# Exploiting the Reader - ZDI-CAN-3507






# Exploiting the Reader - ZDI-CAN-3507

We know the array  
base address


=> We know the  
flink addresses if  
we replace Array Z!

=> We know the  
flink addresses if we  
replace Array Z+n!

Array buffer Z-1



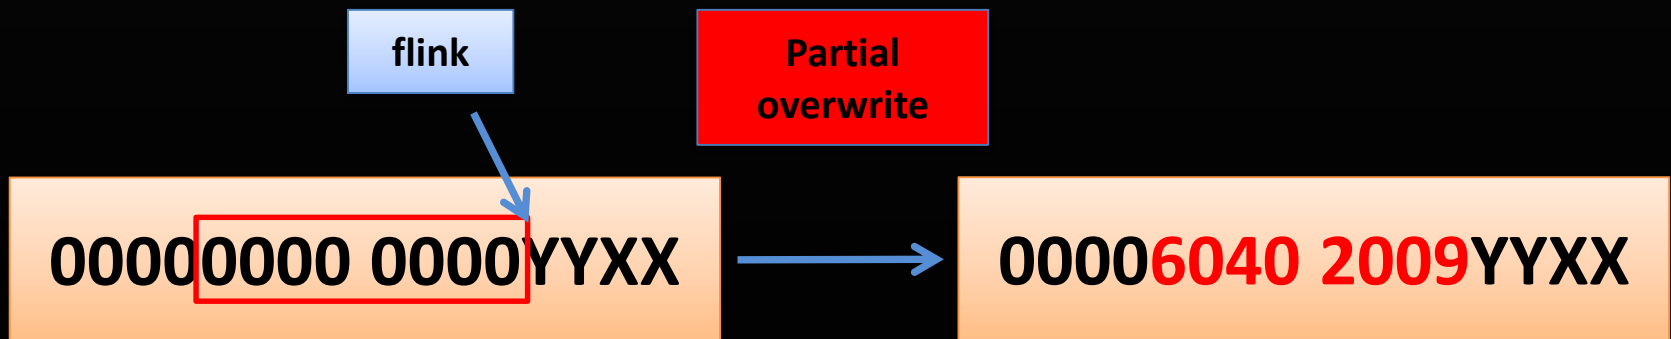
...



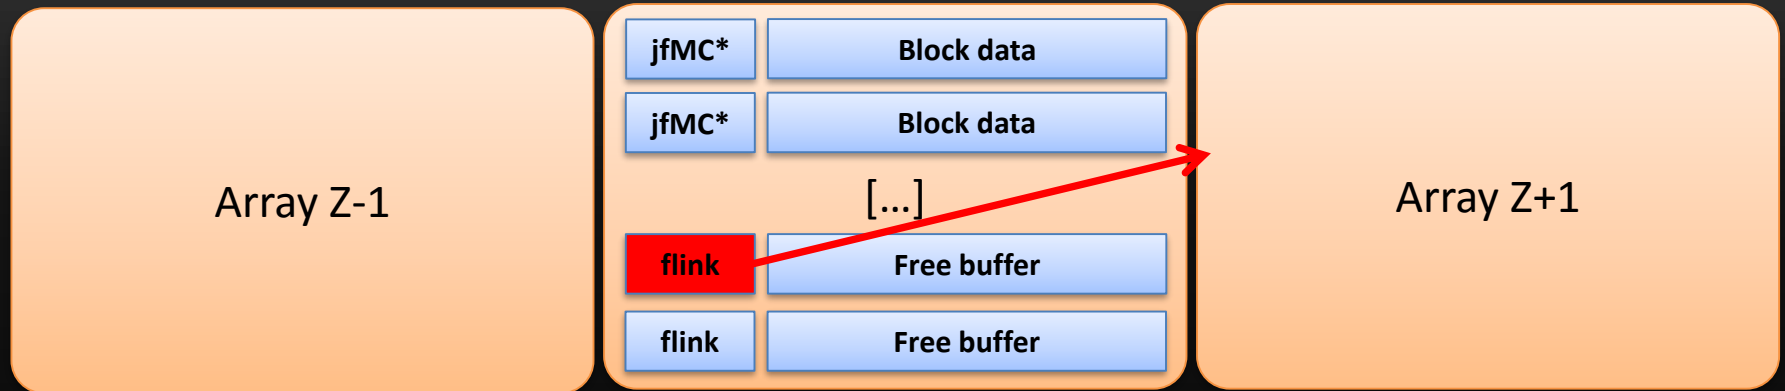
Now we can find a suitable flink  
ending on 00's  
=> This will be the overwrite target!

## Exploiting the Reader - ZDI-CAN-3507

- Knowing the flink addresses we need to search a flink of form 0xXXYY0000
  - Why not 00? You won't shift the flink into the next array!
  - Why not 000000? Very unlikely to find such a flink!
- Lower 16 bits of the flink will be overwritten with upper 16 bits of the object pointer
- Let's assume write of object pointer == 0x09204060

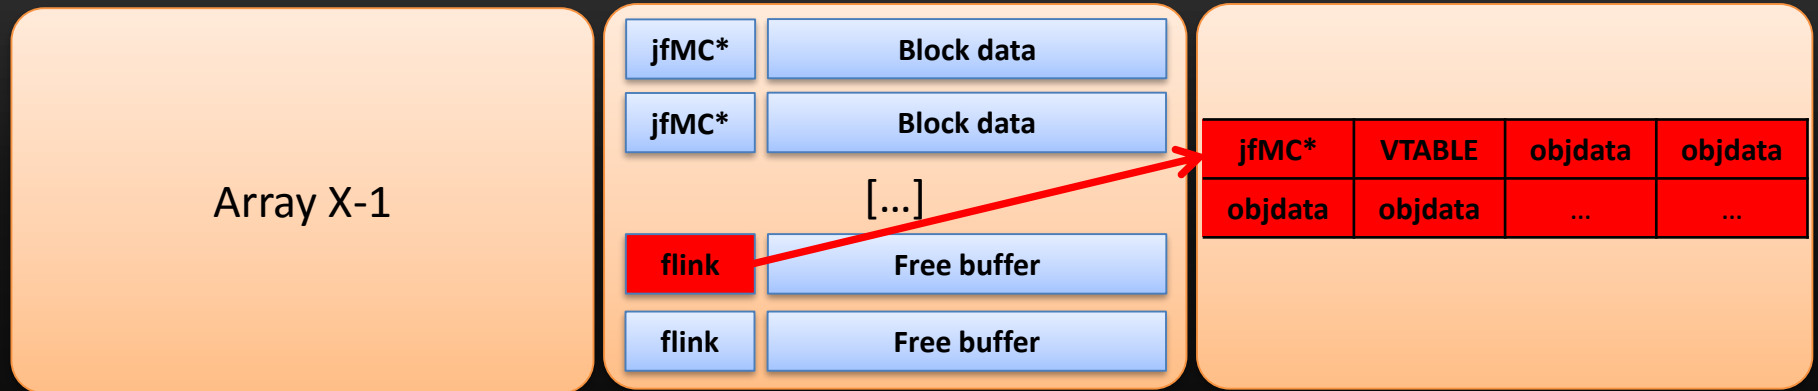


# Exploiting the Reader - ZDI-CAN-3507



- Partial overwrite: 0xXXYY0000 => 0xXXYY0920
- Flink will be shifted 0x920 bytes in this case
  - Flink should be located near to the end of the chunk so that after the overwrite it points to the next Array Z+1!

# Exploiting the Reader - ZDI-CAN-3507



- When the block with the overwritten flink is allocated the data is placed in Array Z+1
- If an object is allocated the vtable will be placed there ready to be read => ASLR bypassed! =)

# ■■■ Exploiting the Reader - ZDI-CAN-3507

- And RCE??
- Super easy!
  - Locate the vtable pointer by finding the overwritten float value in Array Z+1
  - Overwrite this float value so that we hit our stack pivot with the next vtable call
  - Reference the object with the overwritten vtable pointer to cause a vtable call and jump into your ROP
- GAME OVER.

# ■ ■ ■ Demo

# ■■■ Conclusion

## ■■■ Conclusion

- Very easy, but highly effective technique to leak data
- No global RW primitive, but enough to pwn AR
- Version-independant
- OS-independant
- Very fast: From start to pwn in ~ 1 sec possible
  - ZDI-CAN-3507 slow because vuln needs time to trigger
- Flexible technique which can be used with almost every kind of overwrite (as we have just seen)
- Custom allocator proves once again to be a perfect target in memory corruption scenarios



# Thank you for your attention! 😊

■■■ Q&A