

# Windows Metafiles

An Analysis of the EMF Attack Surface & Recent Vulnerabilities

Mateusz “j00ru” Jurczyk

Ruxcon, Melbourne 2016

# PS> whoami

- Project Zero @ Google
- Low-level security researcher with interest in all sorts of vulnerability research and software exploitation
- <http://j00ru.vexillium.org/>
- [@j00ru](#)

# Agenda

- Windows Metafile primer, GDI design, attack vectors.
- Hacking:
  - Internet Explorer (GDI)
  - Windows Kernel (ATMFD.DLL)
  - Microsoft Office (GDI+)
  - VMware virtualization (Print Spooling)
- Final thoughts.

# Windows GDI & Metafile primer

# Windows GDI

- GDI stands for *Graphics Device Interface*.
- Enables user-mode applications to use graphics and formatted text on video displays and printers.
- Major part of the system API (nearly 300 documented functions).
- Present in the OS since the very beginning (Windows 1.0 released in 1985).
  - One of the oldest subsystems, with most of its original code still running 31 years later.
  - Concidentally (?) also one of the most buggy components.

# How to draw

1. Grab a handle to a Device Context (**HDC**).
  - Identifies a persistent container of various graphical settings (pens, brushes, palettes etc.).
  - Can be used to draw to a screen (most typically), a printer, or a metafile.
  - Most trivial example:


```
HDC hdc = GetDC(NULL);
```

*(obtains a HDC for the entire screen)*


# How to draw

2. Use a drawing function.


```
Ellipse(hdc, 100, 100, 500, 300);
```



```
RoundRect(hdc, 100, 100, 500, 500, 100, 100);
```



# Windows GDI – simplified architecture





# User to kernel API mappings

Most user-mode GDI functions have their direct counterparts in the kernel:

GDI32.DLL	win32k.sys
AbortDoc	NtGdiAbortDoc
AbortPath	NtGdiAbortPath
AddFontMemResourceEx	NtGdiAddFontMemResourceEx
AddFontResourceW	NtGdiAddFontResourceW
AlphaBlend	NtGdiAlphaBlend
...	...

# Windows Metafiles

Core idea:

**Store images as lists of records directly describing GDI calls.**

# Windows Metafiles

- **Pros:**

- requires little computation work from the rasterizer itself, as it only has to call GDI functions with the supplied parameters.
- provides an official way of serializing sets of GDI operations into reproducible images.
- can work as a vector format, raster, or both.

- **Cons:**

- only works on Windows, unless full implementation of the supported graphical GDI operations is implemented externally.

# First version: WMF

- The original metafiles (WMF = **W**indows **M**eta**F**iles).
- Introduced with Windows 3.0 in 1990.
  - Not as ancient as GDI itself, but almost so.
- Initially documented in Windows 3.1 SDK (1994, volume 4).
  - A revised, more complete specification was released in 2006, and has been maintained since then.
  - A description of all records and structures can be found in the MS-WMF document.

# WMF files – 60 supported API functions

---

AnimatePaletteArc	LineToMoveToEx	SelectPaletteSetBkColor
BitBlt	OffsetClipRgn	SetBkMode
Chord	OffsetViewportOrgEx	SetDIBitsToDevice
CreateBrushIndirect	OffsetWindowOrgEx	SetMapMode
CreateDIBPatternBrush	PaintRgn	SetMapperFlags
CreateFontIndirect	PatBlt	SetPaletteEntries
CreatePalette	Pie	SetPixel
CreatePatternBrush	Polygon	SetPolyFillMode
CreatePenIndirect	Polyline	SetROP2
DeleteObject	PolyPolygon	SetStretchBltMode
Ellipse	RealizePalette	SetTextAlign
Escape	Rectangle	SetTextCharacterExtra
ExcludeClipRect	ResizePalette	SetTextColor
ExtFloodFill	RestoreDC	SetTextJustification
ExtTextOut	RoundRect	SetViewportOrgEx
FillRgn	SaveDC	SetWindowExtEx
FloodFill	ScaleViewportExtEx	SetWindowOrgEx
FrameRgn	ScaleWindowExtEx	StretchBlt
IntersectClipRect	SelectClipRgn	StretchDIBits
InvertRgn	SelectObject	TextOut

---

# Some seemingly interesting ones

---

AnimatePaletteArc	LineToMoveToEx	SelectPaletteSetBkColor
BitBlt	OffsetClipRgn	SetBkMode
Chord	OffsetViewportOrgEx	SetDIBitsToDevice
CreateBrushIndirect	OffsetWindowOrgEx	SetMapMode
CreateDIBPatternBrush	PaintRgn	SetMapperFlags
CreateFontIndirect	PatBlt	SetPaletteEntries
CreatePalette	Pie	SetPixel
CreatePatternBrush	Polygon	SetPolyFillMode
CreatePenIndirect	Polyline	SetROP2
<b>DeleteObject</b>	PolyPolygon	SetStretchBltMode
Ellipse	RealizePalette	SetTextAlign
<b>Escape</b>	Rectangle	SetTextCharacterExtra
ExcludeClipRect	ResizePalette	SetTextColor
ExtFloodFill	<b>RestoreDC</b>	SetTextJustification
ExtTextOut	RoundRect	SetViewportOrgEx
FillRgn	<b>SaveDC</b>	SetWindowExtEx
FloodFill	ScaleViewportExtEx	SetWindowOrgEx
FrameRgn	ScaleWindowExtEx	StretchBlt
IntersectClipRect	SelectClipRgn	StretchDIBits
InvertRgn	<b>SelectObject</b>	TextOut

---

# WMF: there's more!

- The format also supports a number of records which do not directly correspond to GDI functions.
  - Header with metadata.
  - Embedded EMF.
  - Records directly interacting with the printer driver / output device.
  - End-of-file marker.
  - ...

# WMF: there's more!

- Generally, the most interesting records can be found in two sections:

Name	Section	Description
<a href="#">Bitmap</a> record types	<a href="#">2.3.1</a>	Manage and output bitmaps.
Control record types	<a href="#">2.3.2</a>	Define the start and end of a WMF <a href="#">metafile</a> .
Drawing record types	<a href="#">2.3.3</a>	Perform graphics drawing orders.
Object record types	<a href="#">2.3.4</a>	Create and manage graphics objects.
State record types	<a href="#">2.3.5</a>	Specify and manage the graphics configuration.
Escape record types	<a href="#">2.3.6</a>	Specify extensions to functionality that are not directly available thrc



# Windows Metafile – example

```
...
R0003: [017] META_SETMAPMODE          (s=12) {iMode(8=MM_ANISOTROPIC)}
R0004: [011] META_SETVIEWPORTTEXT     (s=16) {szlExtent(1920,1200)}
R0005: [009] META_SETWINDOWEXT     (s=16) {szlExtent(1920,1200)}
R0006: [010] META_SETWINDOWORGE      (s=16) {ptlOrigin(-3972,4230)}
R0007: [009] META_SETWINDOWEXT     (s=16) {szlExtent(7921,-8462)}
R0008: [049] META_CREATEPALETTE       (s=960) {ihPal(1) LOGPAL[ver:768, entries:236]}
R0009: [048] META_SELECTPALETTE      (s=12) {ihPal(Table object: 1)}
R0010: [052] META_REALIZEPALETTE     (s=8)
R0011: [039] META_CREATEBRUSHINDIRECT (s=24) {ihBrush(2), style(0=BS_SOLID, color:0x00FFFFFF)}
R0012: [037] META_SELECTOBJECT       (s=12) {Table object: 2=OBJ_BRUSH.(BS_SOLID)}
R0013: [037] META_SELECTOBJECT       (s=12) {Stock object: 8=OBJ_PEN.(PS_NULL)}
R0014: [019] META_SETPOLYFILLMODE    (s=12) {iMode(1=ALTERNATE)}
R0015: [086] META_POLYGON16          (s=320) {rclBounds(89,443,237,548), nbPoints:73, P1(-2993,398) - Pn(-2993,398)}
R0016: [038] META_CREATEPEN          (s=28) {ihPen(3), style(0=PS_SOLID | COSMETIC), width(0), color(0x00000000)}
...
```

# WMF: still very obsolete

- Even though already quite complex, the format didn't turn out to be very well thought-out for modern usage.
- It's still supported by GDI, and therefore some of its clients (e.g. Microsoft Office, Paint, some default Windows apps).
- Has been basically forgotten in any real-world use-cases for the last decade or more.

# WMF: discouraged from use

- Even Microsoft gives a lot of reasons not to use it anymore:

The following are the limitations of this format:

- A Windows-format metafile is application and device dependent. Changes in the application's mapping modes or the device resolution affect the appearance of metafiles created in this format.
- A Windows-format metafile does not contain a comprehensive header that describes the original picture dimensions, the resolution of the device on which the picture was created, an optional text description, or an optional palette.
- A Windows-format metafile does not support the new curve, path, and transformation functions. See the list of supported functions in the table that follows.
- Some Windows-format metafile records cannot be scaled.
- The metafile device context associated with a Windows-format metafile cannot be queried (that is, an application cannot retrieve device-resolution data, font metrics, and so on).

# Next up: EMF (Enhanced MetaFiles)

- Already in 1993, Microsoft released an improved revision of the image format, called EMF.
- Documented in the official MS-EMF specification.
- Surpasses WMF in a multitude of ways:
  - uses 32-bit data/offset width, as opposed to just 16 bits.
  - device independent.
  - supports a number of new GDI calls, while maintaining backward compatibility with old records.

# Enhanced Metafile – example

```
...
R0121: [039] EMR_CREATEBRUSHINDIRECT (s=24) {ihBrush(2), style(1=BS_NULL)}
R0122: [037] EMR_SELECTOBJECT (s=12) {Table object: 2=OBJ_BRUSH.(BS_NULL)}
R0123: [040] EMR_DELETEOBJECT (s=12) {ihObject(1)}
R0124: [090] EMR_POLYPOLYLINE16 (s=44) {rclBounds(128,-256,130,-254), nPolys:1, nbPoints:2, P1(386,-765) - Pn(386,-765)}
R0125: [019] EMR_SETPOLYFILLMODE (s=12) {iMode(1=ALTERNATE)}
R0126: [039] EMR_CREATEBRUSHINDIRECT (s=24) {ihBrush(1), style(0=BS_SOLID, color:0x00A86508)}
R0127: [037] EMR_SELECTOBJECT (s=12) {Table object: 1=OBJ_BRUSH.(BS_SOLID)}
R0128: [040] EMR_DELETEOBJECT (s=12) {ihObject(2)}
R0129: [058] EMR_SETMITERLIMIT (s=12) {Limit:0.000}
R0130: [091] EMR_POLYPOLYGON16 (s=60) {rclBounds(127,-259,138,-251), nPolys:1, nbPoints:6, P1(384,-765) - Pn(384,-765)}
R0131: [040] EMR_DELETEOBJECT (s=12) {ihObject(1)}
R0132: [040] EMR_DELETEOBJECT (s=12) {ihObject(3)}
R0133: [014] EMR_EOF (s=20) {nPalEntries:0, offPalEntries:16, nSizeLast:20}
...
```

# EMF: interesting records at first glance

2.3.3	Comment Record Types.....	105
2.3.3.1	EMR_COMMENT Record.....	106
2.3.3.2	EMR_COMMENT_EMFPLUS Record.....	107
2.3.3.3	EMR_COMMENT_EMFSPOOL Record.....	107
2.3.3.4	EMR_COMMENT_PUBLIC Record Types .....	108
2.3.3.4.1	EMR_COMMENT_BEGINGROUP Record .....	109
2.3.3.4.2	EMR_COMMENT_ENDGROUP Record .....	110
2.3.3.4.3	EMR_COMMENT_MULTIFORMATS Record.....	111
2.3.3.4.4	EMR_COMMENT_WINDOWS_METAFILE Record.....	112

# EMF: interesting records at first glance

2.3.6	Escape Record Types .....	159
2.3.6.1	EMR_DRAWESCAPE Record .....	161
2.3.6.2	EMR_EXTESCAPE Record.....	161
2.3.6.3	EMR_NAMEDESCAPE Record.....	162

# EMF: interesting records at first glance

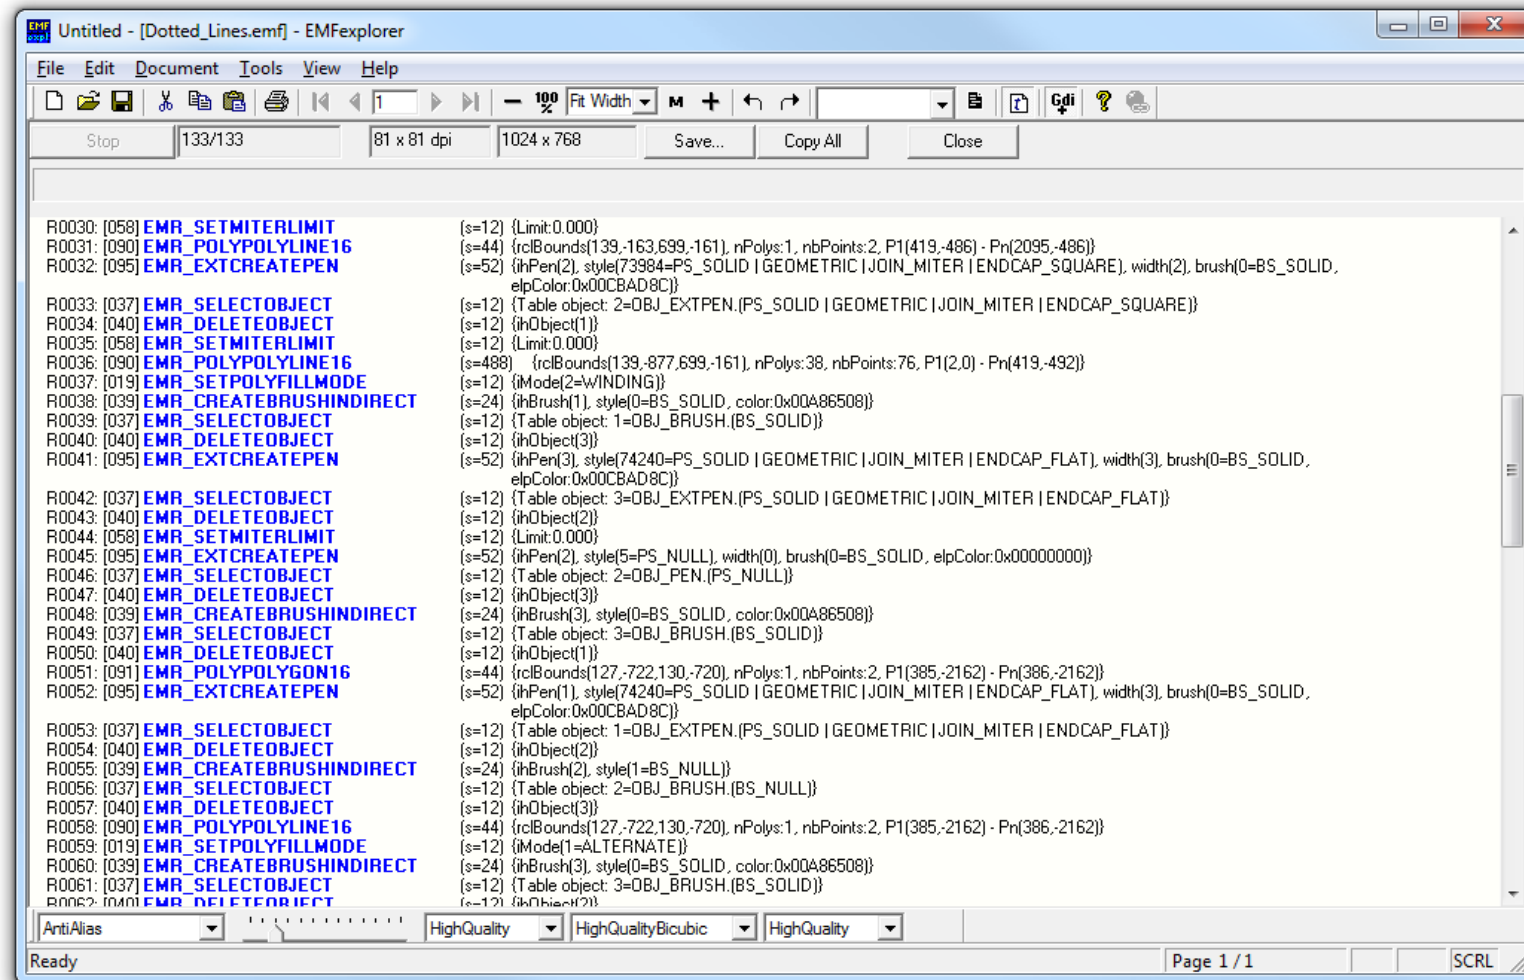
2.3.9	OpenGL Record Types .....	180
2.3.9.1	EMR_GLSBOUNDEDRECORD Record .....	182
2.3.9.2	EMR_GLSRECORD Record .....	182



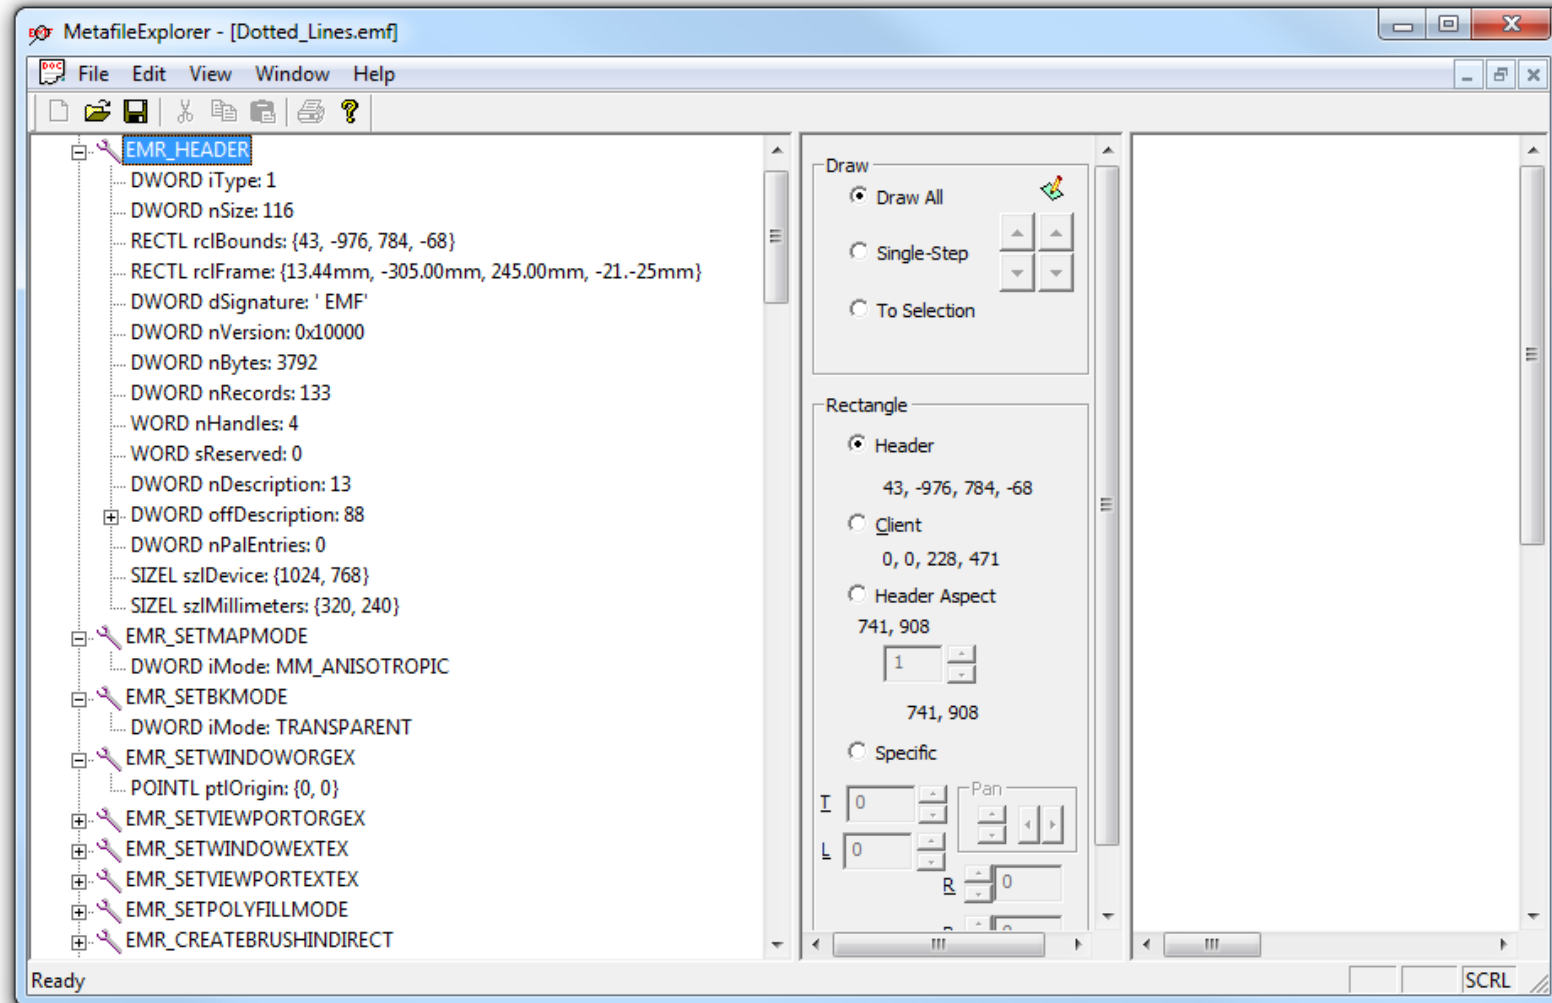
# EMF: current support

- Despite being only 3 years younger than WMF, EMF has remained in current usage until today.
  - Not as a mainstream image format, but still a valid attack vector.
- A variety of attack vectors:
  - Win32 GDI clients – most notably Internet Explorer.
  - GDI+ clients – most notably Microsoft Office.
  - Printer drivers, including those used in virtualization technology.

# Toolset – examination (EMFexplorer)



# Toolset – examination (MetafileExplorer)



# Toolset – reading & writing (pyemf)

```
#!/usr/bin/env python
import os
import pyemf
import sys

def main(argv):
    if len(argv) != 2:
        print "Usage: %s /path/to/poc.emf" % argv[0]
        sys.exit(1)

    emf = pyemf.EMF(width = 100, height = 100, density = 1)
    emf.CreateSolidBrush(0x00ff00)
    emf.SelectObject(1)
    emf.Polygon([(0, 0), (0, 100), (100, 100), (100, 0)])

    emf.save(argv[1])

if __name__ == "__main__":
    main(sys.argv)
```

# The latest: EMF+

- GDI had all the fundamental primitives, but lacked many complex features (anti-aliasing, floating point coords, support for JPEG/PNG etc.).
- Windows XP introduced a more advanced library called GDI+ in 2001.
  - Built as a user-mode gdiplus.dll library, mostly on top of regular GDI (gdi32.dll).
  - Provides high-level interfaces for C++ and .NET, therefore is much easier to use.
  - GDI+ itself is written in C++, so all the typical memory corruption bugs still apply.

# The latest: EMF+

- Since there is a new interface, there must also be a new image format with its serialized calls.
- Say hi to EMF+!
- Basically same as EMF, but representing GDI+ calls.
- Come in two flavours: **EMF+ Only** and **EMF+ Dual**.
  - „Only” contains exclusively GDI+ records, and can only be displayed with GDI+.
  - „Dual” stores the picture with two sets of records, compatible with both GDI/GDI+ clients.

## 2.3 EMF+ Records

This section specifies the Records, which are grouped into the following categories:

Name	Section	Description
Clipping record types	<a href="#">2.3.1</a>	Specify clipping regions and operations.
Comment record types	<a href="#">2.3.2</a>	Specify arbitrary private data in the <a href="#">EMF+ metafile</a> .
Control record types	<a href="#">2.3.3</a>	Specify global parameters for EMF+ metafile processing.
Drawing record types	<a href="#">2.3.4</a>	Specify graphics output.
Object record types	<a href="#">2.3.5</a>	Define reusable graphics objects.
Property record types	<a href="#">2.3.6</a>	Specify properties of the <a href="#">playback device context</a> .
State record types	<a href="#">2.3.7</a>	Specify operations on the state of the playback device context.
Terminal Server record types	<a href="#">2.3.8</a>	Specify graphics processing on a <a href="#">terminal server</a> .
Transform record types	<a href="#">2.3.9</a>	Specify properties and transforms on <a href="#">coordinate spaces</a> .

# Formats and implementations in Windows

- Three formats in total to consider: WMF, EMF, EMF+.
- Three libraries: GDI, GDI+ and **MF3216**.
  - MF3216.DLL is a system library with just one meaningful exported function:  
**ConvertEmfToWmf**.
  - Used for the automatic conversion between WMF/EMF formats in the Windows clipboard.
    - „Synthesized” formats **CF\_METAFILEPICT** and **CF\_ENHMETAFILE**.
- No bugs found there. ☹



# Formats and implementations in Windows


Library	Supported formats
GDI	WMF, EMF
GDI+	WMF, EMF, EMF+
MF3216	EMF

In this talk, we'll focus on auditing and exploiting the EMF parts, as this is where the most (interesting) issues were discovered.

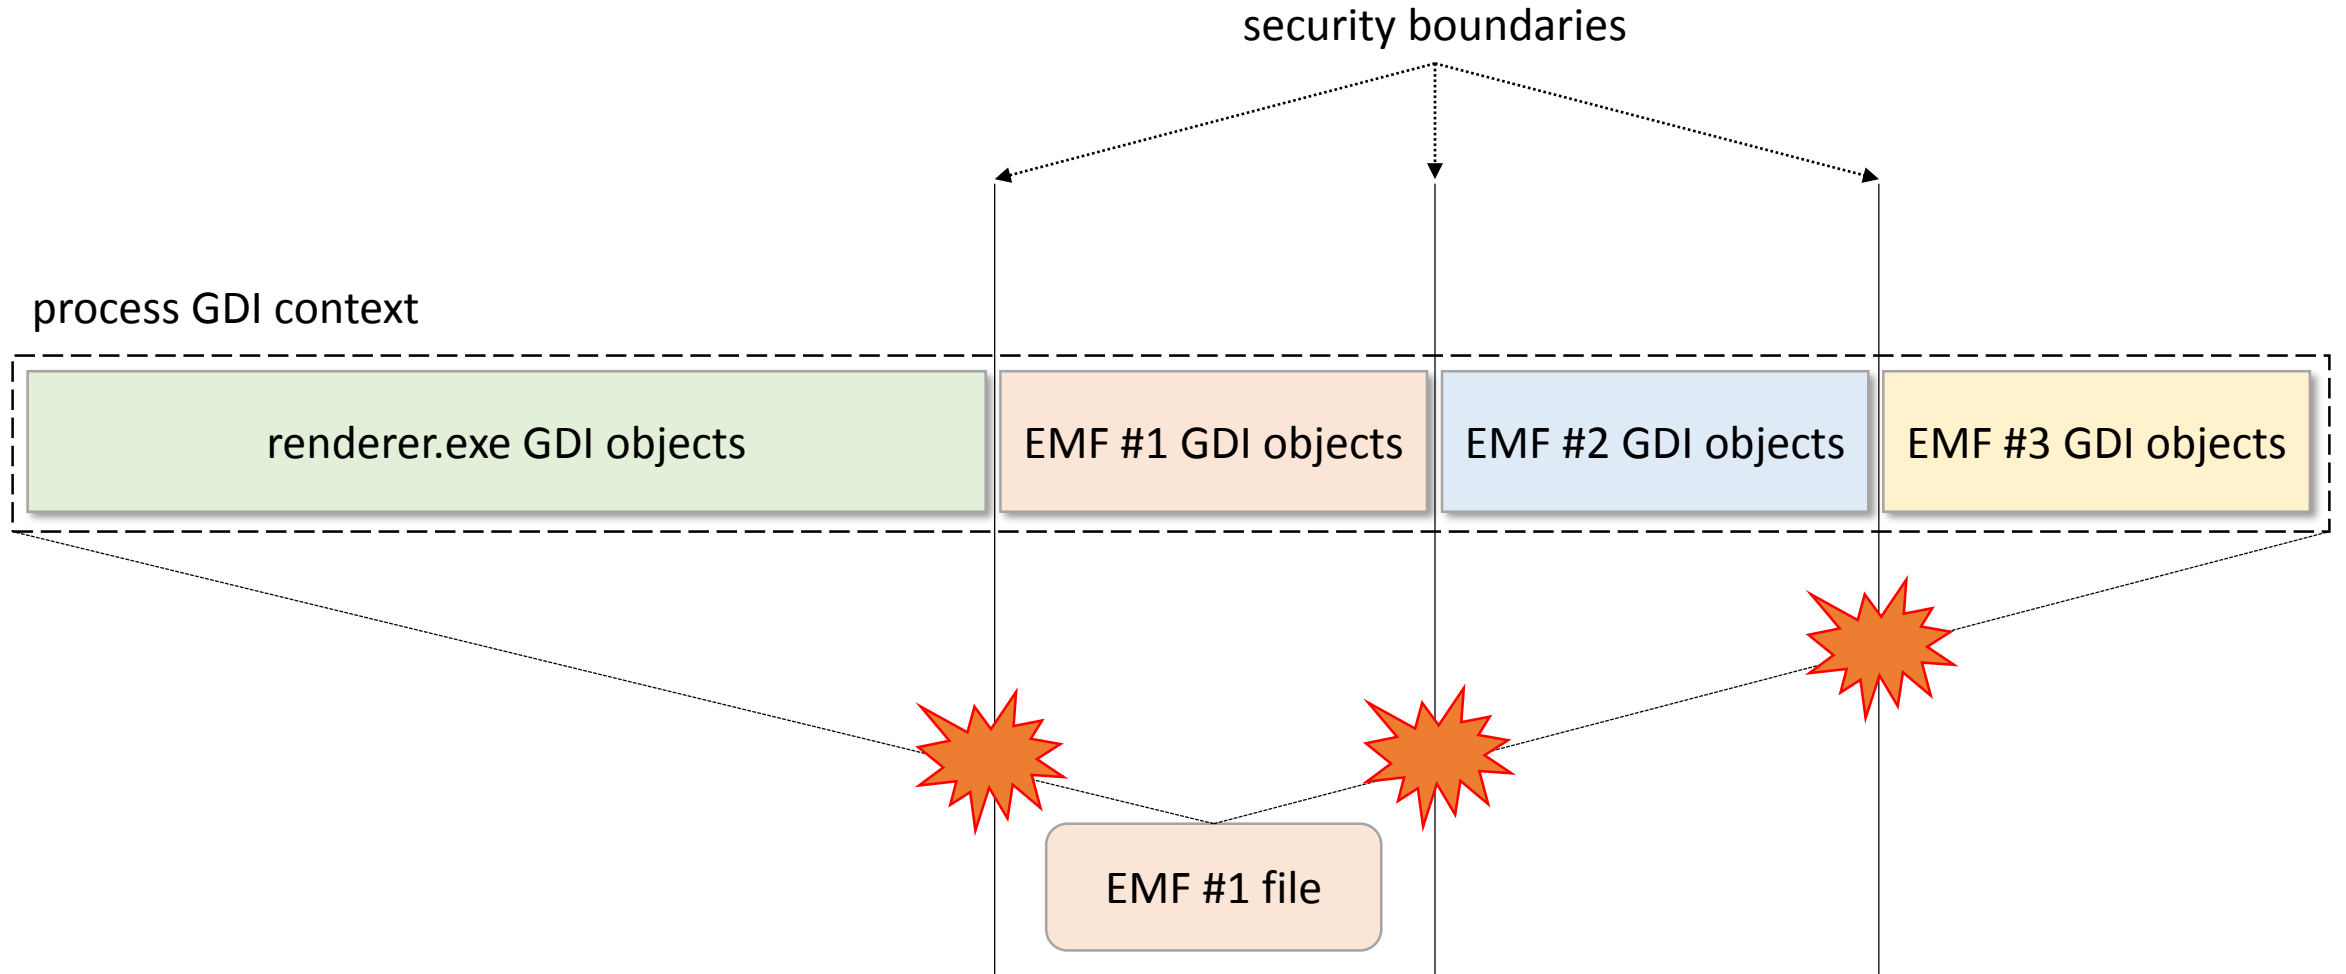
# Attack scenario

- In all cases, Metafiles are processed in the user-mode context of the renderer process, in the corresponding DLL.
  - GDI, GDI+ and MF3216 iterate through all input records and translate them into GDI/GDI+ calls.
- Memory corruption bugs will result in arbitrary code execution in that context.
- **Important:** Metafiles directly operate on the GDI context of the renderer.
  - Can create, delete, change and use various GDI objects on behalf of the process.
  - In theory, it should only have access to its own objects and be self-contained.
  - However, any bugs in the implementation could enable access to *external* graphics objects used by the program.
  - A peculiar case of „privilege escalation“.

# Attack scenario: GDI context priv. escal.



# Attack scenario: GDI context priv. escal.



# Types of Metafile bugs

## 1. **Memory corruption bugs**

- Buffer overflows etc. due to mishandling specific records.
- Potentially exploitable in any type of renderer.
- Impact: typically RCE.

## 2. **Memory disclosure bugs**

- Rendering uninitialized or out-of-bounds heap memory as image pixels.
- Exploitable only in contexts where displayed images can be read back (web browsers, remote renderers).
- Impact: information disclosure (stealing secret information, defeating ASLR etc.).

## 3. **Invalid interaction with the OS and GDI object mismanagement.**

- Impact, exploitability = ???, depending on the specific nature of the bug.

# Let's get started!

- Earlier this year, I started manually auditing the available EMF implementations.
- This has resulted in 10 CVEs from Microsoft and 3 CVEs from VMware (covering several dozen of actual bugs).
- Let's look into the root causes and exploitation of the most interesting ones.
  - Examples are shown based on Windows 7 32-bit, but most of the research applies to both bitnesses and versions up to Windows 10.

Auditing GDI

# Getting started

- To get some general idea of where the functionality in question is implemented and what types of bugs were found in the past, it makes sense to check prior art.
- A „wmf vulnerability” query yields just one result:

**the SetAbortProc bug!**



# SetAbortProc WMF bug (CVE-2005-4560)

- Discovered on **December 27, 2005**. Fixed on **January 5, 2006**.
- Critical bug, allowed 100% reliable RCE while using GDI to display the exploit (e.g. in Internet Explorer).
- Called „**Windows Metafile vulnerability**“, won Pwnie Award 2007.
- No memory corruption involved, only documented features of WMF.
- So what was the bug?

# The GDI API...

## SetAbortProc function

The **SetAbortProc** function sets the application-defined abort function that allows a print job to be canceled during spooling.

### Syntax

C++

```
int SetAbortProc(  
    In HDC      hdc,  
    _In_ ABORTPROC lpAbortProc  
);
```

function pointer



... and the WMF counterpart

#### **2.1.1.17      MetafileEscapes Enumeration**

The MetafileEscapes Enumeration specifies **printer driver** functionality that might not be directly accessible through WMF records defined in the [RecordType Enumeration \(section 2.1.1.1\)](#).

**SETABORTPROC:** Sets the application-defined function that allows a **print job** to be canceled during printing.

# In essence...

... the format itself supported calling:


```
SetAbortProc(hdc, (ABORTPROC)"controlled data");
```

and having the function pointer called afterwards.

Code execution by design.

# Lessons learned

1. The format may (un)officially *proxy* calls to interesting / dangerous API calls, so the semantics of each function and its parameters should be checked for unsafe behavior.
2. The handling of WMF takes place in a giant switch/case in `gdi32!PlayMetaFileRecord`.



# What about EMF bugs?

- Searching for „emf vulnerability” yields more diverse results.
- Most recent one: „[Yet Another Windows GDI Story](#)” by Hossein Lotfi (@hosselot).
  - Fixed in April 2015 as part of MS15-035, assigned CVE-2015-1645.
  - A heap-based buffer overflow due to an unchecked assumption about an input „size” field in one of the records ([SETDIBITSTODEVICE](#)).
  - In large part an inspiration to start looking into EMF security myself.

# Lessons learned

- Main function for playing EMF records is `gdi32!PlayEnhMetaFileRecord`.
- Each record type has its own class with two methods:
  - `::bCheckRecord()` – checks the internal integrity and correctness of the record.
  - `::bPlay()` – performs the actions indicated in the record.



# GDI32 ::bCheckRecord array

```
.text:7DAFD874 dword_7DAFD874 dd 90909090h ; DATA XREF: IsValidEnhMetaRecord(x,x)+25↑r
.text:7DAFD878 int (__thiscall MR::*const * const afnbMRCheck)(struct tagHANDLETABLE *) dd offset
.text:7DAFD87C dd offset MRBP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD880 dd offset MRBP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD884 dd offset MRBP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD888 dd offset MRBP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD88C dd offset MRBP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD890 dd offset MRBPP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD894 dd offset MRBPP::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD898 dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD89C dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8A0 dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8A4 dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8A8 dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8AC dd offset MREOF::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8B0 dd offset MRSETPIXELV::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8B4 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8B8 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8BC dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8C0 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8C4 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8C8 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8CC dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8D0 dd offset MRSETCOLORADJUSTMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8D4 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8D8 dd offset MRGDICOMMENT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8DC dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8E0 dd offset MRDD::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8E4 dd offset MR::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8E8 dd offset MRCREATEBRUSHINDIRECT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8EC dd offset MRCREATEBRUSHINDIRECT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8F0 dd offset MRCREATEBRUSHINDIRECT::bCheckRecord(tagHANDLETABLE *)
.text:7DAFD8F4 dd offset MRCREATEBRUSHINDIRECT::bCheckRecord(tagHANDLETABLE *)
```

# GDI32 ::bPlay array

```
.text:7DAD4E2C dword_7DAD4E2C dd 90909090h ; DATA XREF: PlayEnhMetaFileRecord(x,x,x,x)+32↑r
.text:7DAD4E30 int (__thiscall MR::*const * const afnbMRPlay)(void *, struct tagHANDLETABLE *, unsigned int)
.text:7DAD4E34 dd offset MRPOLYBEZIER::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E38 dd offset MRPOLYGON::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E3C dd offset MRPOLYLINE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E40 dd offset MRPOLYBEZIERTO::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E44 dd offset MRPOLYLINETO::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E48 dd offset MRPOLYPOLYLINE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E4C dd offset MRPOLYPOLYGON::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E50 dd offset MRSETWINDOWEXTTEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E54 dd offset MRSETVIEWPORTEXTTEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E58 dd offset MRSETVIEWPORTORGEEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E5C dd offset MRSETBRUSHORGEEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E60 dd offset MREOF::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E64 dd offset MRSETPIXELV::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E68 dd offset MRSETMAPPERFLAGS::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E6C dd offset MRSETMAPMODE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E70 dd offset MRSETBKMODE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E74 dd offset MRSETPOLYFILLMODE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E78 dd offset MRSETROP2::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E7C dd offset MRSETSTRETCHBLTMODE::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E80 dd offset MRSETTEXTALIGN::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E84 dd offset MRSETCOLORADJUSTMENT::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E88 dd offset MRSETTEXTCOLOR::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E90 dd offset MRSETBKCOLOR::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E94 dd offset MROFFSETCLIPRGN::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E98 dd offset MRMOVEETOEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4E9C dd offset MRSETMETARGN::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4EA0 dd offset MREXCLUDECLIPRECT::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4EA4 dd offset MRINTERSECTCLIPRECT::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4EA8 dd offset MRSCALEVIEWPORTEXTTEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4EAC dd offset MRSCALEWINDOWEXTTEX::bPlay(void *,tagHANDLETABLE *,uint)
.text:7DAD4EB0 dd offset MRSAVEDC::bPlay(void *,tagHANDLETABLE *,uint)
```

That's a starting point.

# CVE-2016-0168

<b>Impact:</b>	File Existence Information Disclosure
<b>Record:</b>	EMR_CREATECOLORSPACE, EMR_CREATECOLORSPACEW
<b>Exploitable in:</b>	Internet Explorer
<b>CVE:</b>	CVE-2016-0168
<b><i>google-security-research</i> entry:</b>	722
<b>Fixed:</b>	MS16-055, 10 May 2016

# Minor bug #1 in EMR\_CREATECOLORSPACEW

- The quality of the code can be immediately recognized by observing many small, but obvious bugs.
- `MRCREATECOLORSPACEW::bCheckRecord()` checks that the size of the record is  $\geq 0x50$  bytes long:

```
.text:7DB01AEF    mov     eax, [esi+4]
.text:7DB01AF2    cmp     eax, 50h
.text:7DB01AF5    jb      short loc_7DB01B1E
```

- Then immediately proceeds to read a `.cbData` field at offset `0x25C`:

```
.text:7DB01AF7    mov     ecx, [esi+25Ch]
```

- Result: out-of-bounds read by `0x20C` bytes.

# Minor bug #2 in EMR\_CREATECOLORSPACEW

- Then, the .cbData from invalid offset 0x25C is used to verify the record length:

```
.text:7DB01AF7    mov     ecx, [esi+25Ch]
.text:7DB01AFD    add     ecx, 263h
.text:7DB01B03    and     ecx, 0FFFFFFFCh
.text:7DB01B06    cmp     eax, ecx
.text:7DB01B08    ja      short loc_7DB01B1E
```

- The above translates to:

```
if (... && record.length <= ((record->cbData + 0x263) & ~3) && ...) {
    // Record valid.
}
```

# Minor bug #2 in EMR\_CREATECOLORSPACEW

- Two issues here:
  1. Obvious integer overflow making a large .cbData pass the check.
  2. Why would the record length be **smaller** then the data declared within? It should be **larger**!
- It all doesn't matter anyway, since the data is not used in any further processing.

# Minor bug #3 in EMR\_CREATECOLORSPACEW

- The `.lcsFilename` buffer of the user-defined `LOGCOLORSPACEW` structure is not verified to be nul-terminated.
  - May lead to out-of-bound reads while accessing the string.
- As clearly visible, there are lots of unchecked assumptions in the implementation, even though only minor so far.
  - Keeps our hopes up for something more severe.



# The file existence disclosure

- Back to the functionality of `EMR_CREATECOLORSPACE[W]` records: all they do is call `CreateColorSpace[W]` with a fully controlled `LOGCOLORSPACE` structure:

```
typedef struct tagLOGCOLORSPACE {  
    DWORD        lcsSignature;  
    DWORD        lcsVersion;  
    DWORD        lcsSize;  
    LCSCSTYPE     lcsCSType;  
    LCSGAMUTMATCH lcsIntent;  
    CIEXYZTRIPLE  lcsEndpoints;  
    DWORD        lcsGammaRed;  
    DWORD        lcsGammaGreen;  
    DWORD        lcsGammaBlue;  
    TCHAR        lcsFilename[MAX_PATH];  
} LOGCOLORSPACE, *LPLOGCOLORSPACE;
```

# Inside CreateColorSpaceW

- The function builds a color profile file path using internal `gdi32!BuildIcmProfilePath`.
  - if the provided filename is relative, it is appended to a system directory path.
  - otherwise, absolute paths are left as-is.
- All paths are accepted, except for those starting with two "/" or "\" characters:

```
if ((pszSrc[0] == '\\\' || pszSrc[0] == '/') &&  
    (pszSrc[1] == '\\\' || pszSrc[1] == '/')) {  
    // Path denied.  
}
```

# Inside CreateColorSpaceW

- This is supposedly to prevent specifying remote UNC paths starting with the "\\\" prefix, e.g. \\192.168.1.13\C\Users\test\profile.icc.
- However, James Forshaw noted that this check is not effective, as the prefix can be also represented as "\\?\UNC\".
- The check is easily bypassable with:

\\?\UNC\192.168.1.13\C\Users\test\profile.icc

# CreateColorSpaceInternalW: last step

- After the path is formed, but before invoking the `NtGdiCreateColorSpace` system call, the function opens the file and immediately closes it to see if it exists:

```
HANDLE hFile = CreateFileW(&FileName, GENERIC_READ, FILE_SHARE_READ, 0,
                           OPEN_EXISTING, FILE_ATTRIBUTE_NORMAL, 0);
if (hFile == INVALID_HANDLE_VALUE) {
    GdiSetLastError(2016);
    return 0;
}
CloseHandle(hFile);
```

# Consequences

- In result, we can have `CreateFileW()` called over any chosen path.
  - If it succeeds, the color space object is created and the function returns success.
  - If it fails, the GDI object is not created and the handler returns failure.
- Sounds like information disclosure potential.
  - How do we approach exploitation e.g. in Internet Explorer?

# Intuitive way: leaking the return value

- Since the return value of `CreateFileW()` determines the success of the record processing, we could maybe leak this bit?
  - Initial idea: use `EMR_CREATECOLORSPACE` as the first record, followed by a drawing operation.
  - If the drawing is never executed (which can be determined with the `<canvas>` tag), the call failed.

# Intuitive way: leaking the return value

- Unfortunately impossible.
- The `gdi32!_bInternalPlayEMF` function (called by `PlayEnhMetaFile` itself) doesn't abort image processing when one record fails.
  - A „success“ flag is set to `FALSE`, and the function proceeds to further operations.
- All records are always executed, and the return value is a flag indicating if at least one of the records failed during the process.

# Can't we leak the final return value?

- No, not really.
- The return value of `PlayEnhMetaFile` is discarded by Internet Explorer in `mshtml!CImgTaskEmf::Decode`:

```
.text:64162B49      call     ds:__imp__PlayEnhMetaFile@12
.text:64162B4F      or      dword ptr [ebx+7Ch], 0FFFFFFFFh
.text:64162B53      lea     eax, [esp+4C8h+var_49C]
```



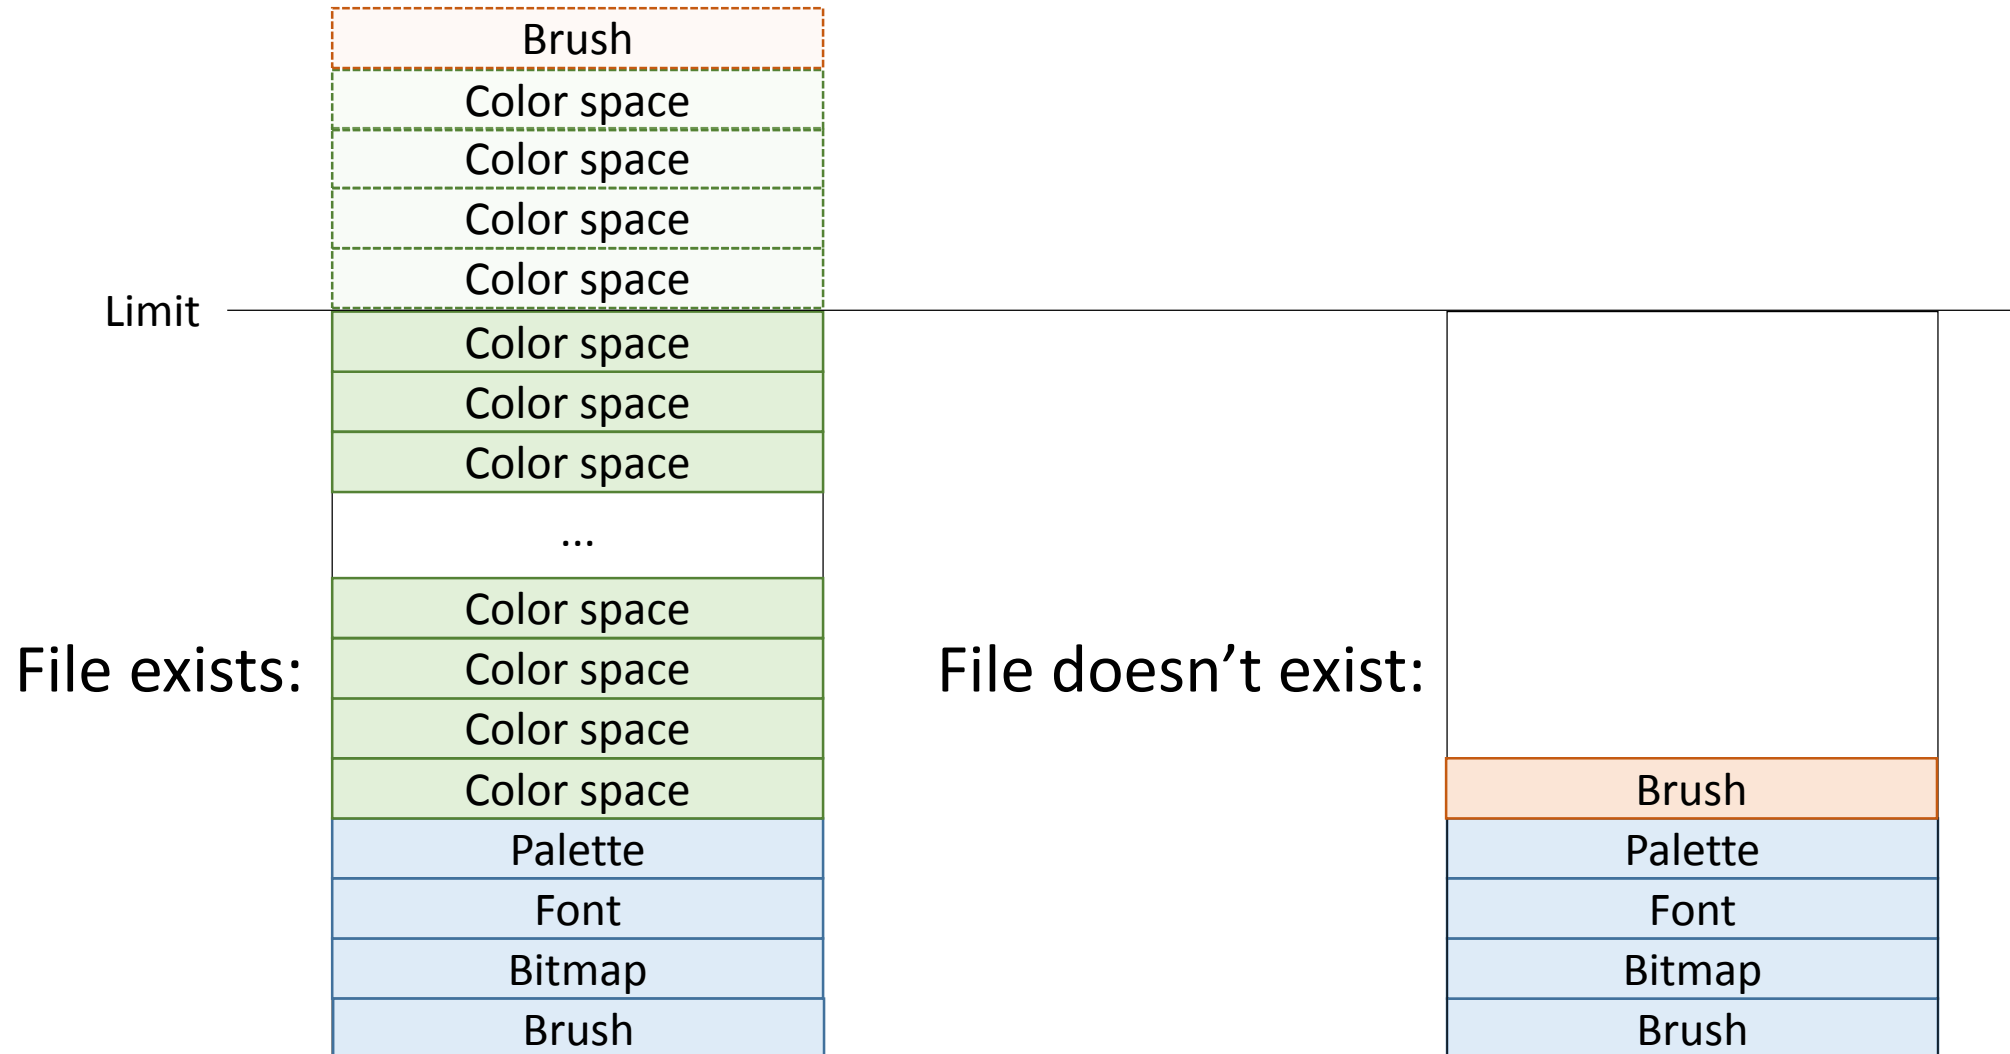
# Other disclosure options

- The other indicator could be the creation of a color space object via `NtGdiCreateColorSpace`.
- Leaking it directly is not easy (if at all possible), but maybe there is some side channel?

# Using the GDI object limit

- Every process in Windows is limited to max. 10,000 GDI objects by default.
  - The number can be adjusted in the registry, but isn't for IE.
- If we use 10,000 EMR\_CREATECOLORSPACEW records with the file path we want to check, then:
  - If the file exists, we'll have 10,000 color space objects, reaching the per-process limit.
  - If it doesn't, we won't have any color spaces at all.
- We're now either at the limit, or not. If we then create a brush (one more object) and try to paint, then:
  - If the file exists, the brush creation will fail and the default brush will be used.
  - If it doesn't, the brush will be created and used for painting.

# GDI object limit as oracle illustrated



**DEMO**

# Vulnerability impact

- Arbitrary file existence disclosure, useful for many purposes:
  - Recognizing specific software (and versions) that the user has installed, for targetted attacks.
  - Tracking users (by creating profiles based on existing files).
  - Tracking the opening times of offline documents (e.g. each opening in Microsoft Office could trigger a ping to remote server via SMB).
  - Blindly scanning network shares available to the user.

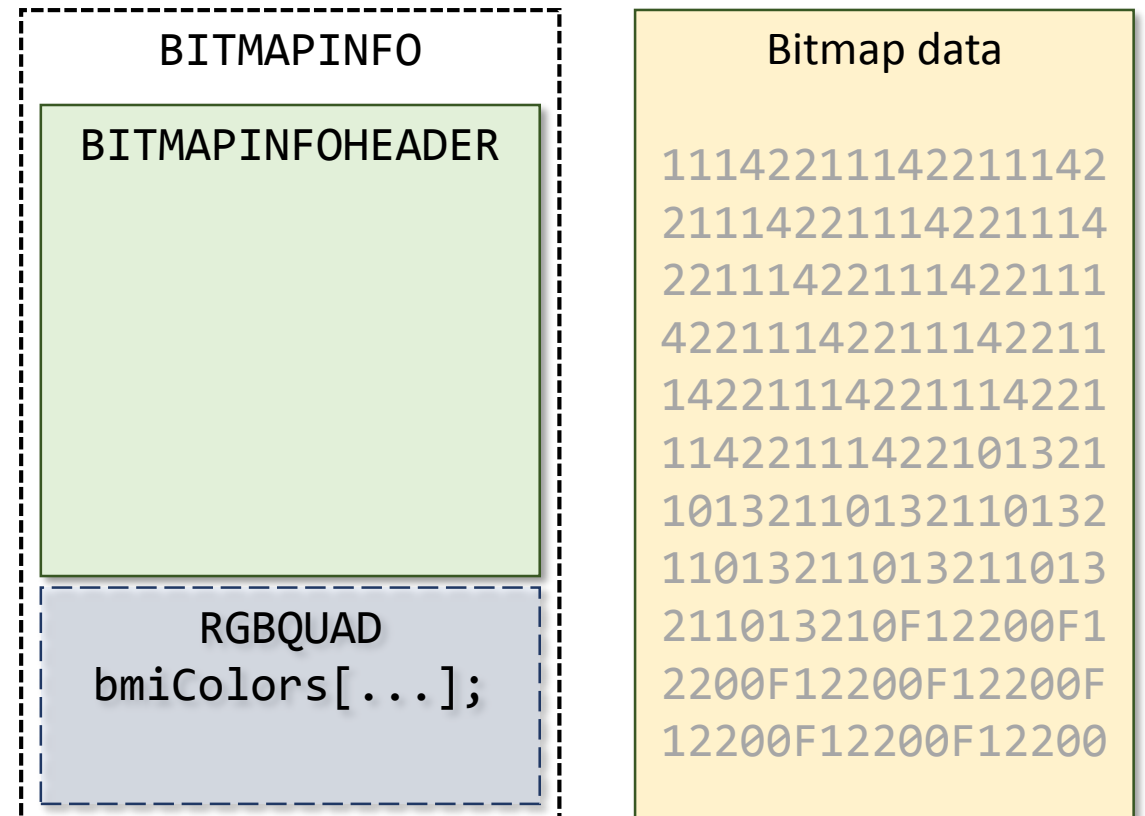
# CVE-2016-3216

<b>Impact:</b>	Memory disclosure
<b>Record:</b>	Multiple records (10)
<b>Exploitable in:</b>	Internet Explorer
<b>CVE:</b>	CVE-2016-3216
<b><i>google-security-research</i> entry:</b>	757
<b>Fixed:</b>	MS16-074, 14 June 2016

# Device Independent Bitmaps (DIBs)

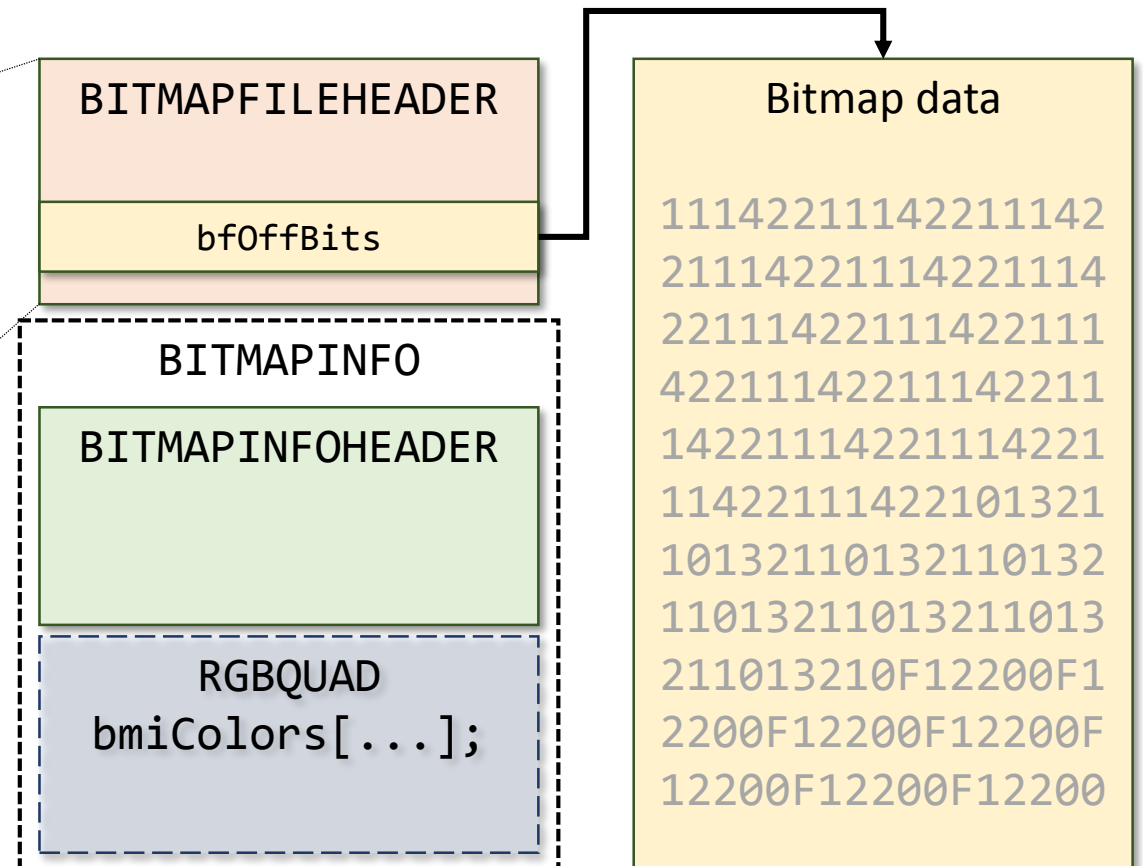
In Windows GDI, raster bitmaps are usually stored in memory in the form of DIBs:

- Short header containing basic metadata about the image, followed by optional palette.
- The image data itself.



# .BMP files are just DIBs, too.

```
typedef struct tagBITMAPFILEHEADER {  
    WORD    bfType;  
    DWORD   bfSize;  
    WORD    bfReserved1;  
    WORD    bfReserved2;  
    DWORD   bfOffBits;  
} BITMAPFILEHEADER;
```





# BITMAPINFOHEADER, the trivial header

```
typedef struct tagBITMAPINFOHEADER {  
    DWORD biSize;  
    LONG biWidth;  
    LONG biHeight;  
    WORD biPlanes;  
    WORD biBitCount;  
    DWORD biCompression;  
    DWORD biSizeImage;  
    LONG biXPelsPerMeter;  
    LONG biYPelsPerMeter;  
    DWORD biClrUsed;  
    DWORD biClrImportant;  
} BITMAPINFOHEADER;
```

- Short and simple structure.
- 40 bytes in length (in typical form).
- Only 8 meaningful fields.

# Is it really so trivial to handle?

- **biSize** needs to be sanitized (can only be a few valid values).
- **biWidth**, **biHeight**, **biPlanes**, **biBitCount** can cause integer overflows (often multiplied with each other).
- **biHeight** can be negative to indicate bottom-up bitmap.
- **biPlanes** must be 1.
- **biBitCount** must be one of {1, 2, 4, 8, 16, 24, 32}.
  - For **biBitCount** < 16, a color palette can be used.
  - The size of the color palette is also influenced by **biClrUsed**.

# Is it really so trivial to handle?

- **biCompression** can be BI\_RGB, BI\_RLE8, BI\_RLE4, BI\_BITFIELDS, ...
  - Each compression scheme must be handled correctly.
- **biSizeImage** must correspond to the actual image size.
- The palette must be sufficiently large to contain all entries.
- The pixel data buffer must be sufficiently large to describe all pixels.
- Encoded pixels must correspond to the values in header (e.g. not exceed the palette size etc.).

# Many potential problems

1. The decision tree for correctly handling a DIB based on its header is very complex.
2. Lots of corner cases to cover and implementation bugs to avoid.
3. A consistent handling across various parts of code is required.

# GDI functions operating on DIB (directly)

```
int StretchDIBits(  
    _In_      HDC      hdc,  
    _In_      int      XDest,  
    _In_      int      YDest,  
    _In_      int      nDestWidth,  
    _In_      int      nDestHeight,  
    _In_      int      XSrc,  
    _In_      int      YSrc,  
    _In_      int      nSrcWidth,  
    _In_      int      nSrcHeight,  
    _In_ const VOID    *lpBits,  
    _In_ const BITMAPINFO *lpBitsInfo,  
    _In_      UINT      iUsage,  
    _In_      DWORD     dwRop  
);
```

pointer to image data

pointer to DIB header

```
int SetDIBitsToDevice(  
    _In_      HDC      hdc,  
    _In_      int      XDest,  
    _In_      int      YDest,  
    _In_      DWORD     dwWidth,  
    _In_      DWORD     dwHeight,  
    _In_      int      XSrc,  
    _In_      int      YSrc,  
    _In_      UINT      uStartScan,  
    _In_      UINT      cScanLines,  
    _In_ const VOID    *lpvBits,  
    _In_ const BITMAPINFO *lpbmi,  
    _In_      UINT      fuColorUse  
);
```

# GDI functions operating on DIB (indirectly)

```
HBITMAP CreateBitmap(  
    _In_      int    nWidth,  
    _In_      int    nHeight,  
    _In_      UINT   cPlanes,  
    _In_      UINT   cBitsPerPel,  
    _In_ const VOID *lpvBits  
);
```

```
BOOL MaskBlt(  
    _In_ HDC      hdcDest,  
    _In_ int       nXDest,  
    _In_ int       nYDest,  
    _In_ int       nWidth,  
    _In_ int       nHeight,  
    _In_ HDC      hdcSrc,  
    _In_ int       nXSrc,  
    _In_ int       nYSrc,  
    _In_ HBITMAP  hbmMask,  
    _In_ int       xMask,  
    _In_ int       yMask,  
    _In_ DWORD     dwRop  
);
```

# Data sanitization responsibility

- In all cases, it is the API caller's responsibility to make sure the headers and data are correct and adequate.
- Passing in fully user-controlled input data is somewhat problematic, as the application code would have to „clone“ GDI's DIB handling.
- Guess what? EMF supports multiple records which contain embedded DIBs.

# EMF records containing DIBs

- EMR\_ALPHABLEND
- EMR\_BITBLT
- EMR\_MASKBLT
- EMR\_PLGBLT
- EMR\_STRETCHBLT
- EMR\_TRANSPARENTBLT
- EMR\_SETDIBITSTODEVICE
- EMR\_STRETCHDIBITS
- EMR\_CREATEMONOBRUSH
- EMR\_EXTCREATEPEN



# The common scheme

- Two pairs of (offset, size) for both the header and the bitmap:

**offBmi (4 bytes):** A 32-bit unsigned integer that specifies the offset from the start of this record to the DIB header, if the record contains a DIB.

**cbBmi (4 bytes):** A 32-bit unsigned integer that specifies the size of the DIB header, if the record contains a DIB.

**offBits (4 bytes):** A 32-bit unsigned integer that specifies the offset from the start of this record to the DIB bits, if the record contains a DIB.

**cbBits (4 bytes):** A 32-bit unsigned integer that specifies the size of the DIB bits, if the record contains a DIB.

# Necessary checks in the EMF record handlers

- In each handler dealing with DIBs, there are four necessary consistency checks:
  1. `cbBmiSrc` is adequately large for the header to fit in.
  2. `(offBmiSrc, offBmiSrc + cbBmiSrc)` resides fully within the record.
  3. `cbBitsSrc` is adequately large for the bitmap data to fit in.
  4. `(offBitsSrc, offBitsSrc + cbBitsSrc)` resides fully within the record.

# Checks were missing in many combinations

Record handlers	Missing checks
MRALPHABLEND::bPlay MRBITBLT::bPlay MRMASKBLT::bPlay MRPLGBLT::bPlay MRSTRETCHBLT::bPlay MRTRANSPARENTBLT::bPlay	#1, #2
MRSETDIBITSTODEVICE::bPlay	#3
MRSTRETCHDIBITS::bPlay	#1, #3
MRSTRETCHDIBITS::bPlay MRCREATEMONOBRUSH::bPlay MREXTCREATEPEN::bPlay	#1, #2, #3, #4


\* This was just after a cursory look; Microsoft might have fixed more.

# The consequence

- Due to missing checks, parts of the image description could be loaded from other parts of the process address space (e.g. adjacent heap allocations):
  - DIB header
  - Color palette
  - Pixel data
- Uninitialized or out-of-bound heap memory could be disclosed with the palette or pixel data.

# Proof of concept

- I hacked up a PoC file with an `EMR_STRETCHBLT` record, containing an 8-bpp DIB with palette entries going beyond the file.
- Result: garbage bytes being displayed as image pixels.
- The same picture being displayed three times in a row in IE:



- The data can be read back using HTML5, in order to leak module addresses and other sensitive data.

**DEMO**

Auditing ATMFD.DLL

# Looking further into the list of EMF records

2.3.6	Escape Record Types .....	159
2.3.6.1	EMR_DRAWESCAPE Record .....	161
2.3.6.2	EMR_EXTESCAPE Record.....	161
2.3.6.3	<u>EMR_NAMEDESCAPE Record.....</u>	162



# NamedEscape?

- **DrawEscape()** and **ExtEscape()** are both documented functions.
  - They are also pretty well explored and researched.
- What's **NamedEscape()**?
  - The function is exported from gdi32.dll.
  - However, no documentation is provided by Microsoft.
  - Internally, it is a simple wrapper for **win32k!NtGdiExtEscape**, the same syscall that **Escape()** and **ExtEscape()** use.
    - Passes along two input arguments which are otherwise set to 0.

# What do the specs say?

### 2.3.6.3 EMR\_NAMEDESCAPE Record

The MR\_NAMEESCAPE record passes arbitrary information to a specified printer driver.

**Note:** Fields that are not described in this section are specified in section [2.3.6](#).

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Type																															
Size																															
iEscape																															
cjDriver																															
cjIn																															
DriverName (variable)																															
...																															
Data (variable)																															
...																															

# Sending data to a driver by name

- In `[Ext]Escape()`, the driver is identified by the HDC.
- Here, we can directly specify the driver's name!
- The interface is similar to IOCTLs.
  - Escape code (32-bit value).
  - Input buffer of controlled size.
  - Output buffer of controlled size (missing from `EMR_NAMEDESCAPE`).
- What is the actual attack surface?
  - Let's search online for „NamedEscape“.

# First result

## atmfd NamedEscape(0x2514) buffer-underflow vulnerability


Project Member Reported by [taviso@google.com](mailto:taviso@google.com), Jul 2 2015

A buffer-underflow vulnerability exists when using NamedEscape(0x2514) in atmfd.

```
kd> kv
Child-SP      RetAddr      : Args to Child                               : Call Site
fffff880`059e8458 fffff800`02ac4e69 : 00000000`0000003b 00000000`c0000005 fffff960`00197fac fffff880`059e8d20 : nt!KeBugCheckEx
fffff880`059e8460 fffff800`02ac47bc : fffff880`059e94c8 fffff880`059e8d20 00000000`00000000 fffff800`02af1630 : nt!KiBugCheckDispatch+0x69
fffff880`059e85a0 fffff800`02af113d : fffff800`02ceb248 fffff800`02c23514 fffff800`02a51000 fffff880`059e94c8 : nt!KiSystemServiceHandler+0x7c
fffff880`059e85e0 fffff800`02aeff15 : fffff800`02c1931c fffff880`059e8658 fffff880`059e94c8 fffff800`02a51000 :
nt!RtlpExecuteHandlerForException+0xd
fffff880`059e8610 fffff800`02b00e81 : fffff880`059e94c8 fffff880`059e8d20 fffff880`00000000 00000000`00000001 : nt!RtlDispatchException+0x415
fffff880`059e8c00 fffff800`02ac4f42 : fffff880`059e94c8 00000000`00000000 fffff880`059e9570 fffff900`c3f81000 : nt!KiDispatchException+0x135
fffff880`059e9390 fffff800`02ac3aba : 00000000`00000000 00000000`00000008 00000000`00000400 00000000`00000000 : nt!KiExceptionDispatch+0xc2
fffff880`059e9570 fffff960`00197fac : 00000000`00000001 fffff900`c3f81000 00000000`00000001 42424242`41414141 : nt!KiPageFault+0x23a (TrapFrame @
fffff880`059e9570)
fffff880`059e9700 fffff960`001a0411 : 00000000`00000000 00000000`00000001 fffff900`c3f81000 fffff900`00000000 :
win32k!SURFACE::bDeleteSurface+0x264
fffff880`059e9850 fffff960`00197940 : 00000000`00000000 00000000`00000000 fffff900`c1e05330 fffff900`00000000 : win32k!NtGdiCloseProcess+0x2c9
fffff880`059e98b0 fffff960`00197087 : 00000000`00000000 00000000`00000001 fffffa80`074d7b50 00000000`00000001 : win32k!GdiProcessCallout+0x200
fffff880`059e9930 fffff800`02d990cd : 00000000`00000000 00000000`00000000 00000000`00000000 fffffa80`074d7b00 : win32k!W32pProcessCallout+0x6b
fffff880`059e9960 fffff800`02d7d2b0 : 00000000`00000000 00000000`00000001 fffffa80`06f42600 00000000`00000000 : nt!PspExitThread+0x4d1
fffff880`059e9a60 fffff800`02ac4b53 : fffffa80`06f42600 fffff880`00000000 fffffa80`074d7b50 00000000`00000000 : nt!NtTerminateProcess+0x138
fffff880`059e9ae0 00000000`76f5de7a : 00000000`00000000 00000000`00000000 00000000`00000000 00000000`00000000 : nt!KiSystemServiceCopyEnd+0x13
(TrapFrame @ fffff880`059e9ae0)
00000000`0008e318 00000000`00000000 : 00000000`00000000 00000000`00000000 00000000`00000000 00000000`00000000 : 0x76f5de7a
```

This bug is subject to a 7 day disclosure deadline, as the issue is being exploited in the wild. If 7 days elapse without a broadly available patch, then the bug report will automatically become visible to the public.

A small testcase is attached.

 **test.c**  
6.5 KB [Download](#)

 **fontdata.h**  
3.6 MB [Download](#)

# Hacking Team ATMFD.DLL 0-day

- Discovered in the leaked data dump on July 7, 2015.
- Fixed by Microsoft on July 14 (MS15-077, CVE-2015-2387).
- Local privilege escalation to ring-0 through vulnerable ATMFD.DLL.
  - Bug triggered through `NamedEscape("ATMFD.DLL", 0x2514)`.
  - Also used in the exploit: `NamedEscape("ATMFD.DLL", 0x250A)`.
  - Hey, I know this driver!

# NamedEscape + ATMFD


- ATMFD.DLL is a very special case for the NamedEscape interface.
  - It is one of a few, or perhaps the only driver using this interface for communication.
- It is even specifically checked for in the `win32k!GreNamedEscape` function:

```
.text:BF9DC326    cmp     esi, PDEV * gppdevATMFD
.text:BF9DC32C    jnz     short loc_BF9DC33F
.text:BF9DC32E    push    offset aAtmfd_dll ; "atmfd.dll"
.text:BF9DC333    push    ebx                ; wchar_t *
.text:BF9DC334    call    __wcsicmp
```

# Finding the handler function

- Locating the escape function within ATMFD.DLL is easy.
  - Just search for some magic values – e.g. 0x2514 – in (hex)decimal in IDA Pro or Hex-Rays.
  - You'll find it right away.
- In case of the latest Windows 7 32-bit, the address is 0x14654.

# A broad control flow graph





# What do we learn?

- 13 escape codes supported, each expecting a specific input length:

Escape code	Input data length
0x2502	0
0x2509	194
0x250A	12
0x250B	194
0x250C	48
0x250D	>88
0x250E	1656
0x250F	0
0x2510	6
0x2511	32
0x2512	1124
0x2513	148
0x2514	≥6

# Analysis was difficult

- Sure we know the escape codes and input data sizes, but:
  - No debug symbols are available, so no function names, structures, data types etc.
  - Unknown functionality of the codes.
  - Unknown format of input and output data.
  - Unknown internal structures.
  - No public documentation available.
- Not the most convenient target to look into (very high entry bar).
- Intended to have a deeper look in 2015, but got distracted and gave up.

# Giving it another shot

- When I noticed that the functionality was also reachable from within EMF in 2016, I decided to give it another shot.
  - Web browser → ring-0 execution potential?
- Let's see what other system modules use the `NamedEscape()` function!

☒ Find text:

namedescape

☐ Whole words only

☐ Case sensitive

☐ RegEx (2)

☐ Hex

☐ Find files NOT containing the text

☒ ANSI charset (Windows)

☐ ASCII charset (DOS)

☐ Unicode UTF-16

☐ UTF8

☐ Office xml (docx, xlsx, odt etc)+EPUB

Search results:

[2 files and 0 directories found]

c:\Windows\System32\qdi32.dll

c:\Windows\System32\atmlib.dll

# ATMLIB.DLL?

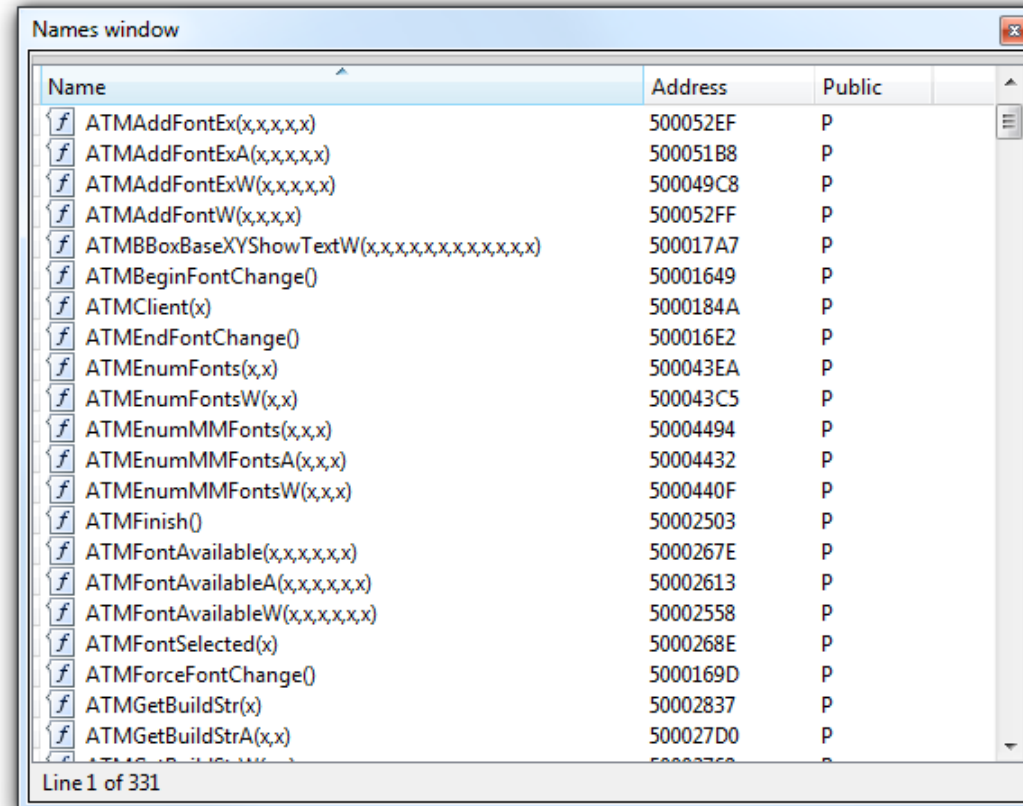
Description	
File description	Windows NT Open Type/Type 1 API Library.
Type	Application extension
File version	5.1.2.248
Product name	Adobe Type Manager
Product version	5.1 Build 248
Copyright	©1983-1990, 1993-2004 Adobe Systems Inc.
Size	33,5 KB
Date modified	2016-05-13 23:27
Language	English (United States)
Legal trademarks	Adobe, Multiple Master, ATM, Adobe Type Ma
Original filename	ATMLIB.DLL

# The missing part of ATM

- Part of the *Adobe Type Manager* suite.
  - Family of computer programs for rasterizing PostScript fonts (Type 1 and OpenType).
  - Ported to Windows (3.0, 3.1, 95, 98, Me) by patching into the OS at a very low level.
  - First officially incorporated into Windows in NT 4.0.
  - ATMFD.DLL is the kernel-mode font driver.
  - ATMLIB.DLL is the user-mode counterpart, which provides the ATM API to client applications.

# Best part about ATMLIB.DLL?

- Debug symbols available from the Microsoft servers!























The screenshot shows a debugger's 'Names window' displaying a list of exported functions from ATMLIB.DLL. The window has a title bar 'Names window' and a close button. The list is organized into three columns: 'Name', 'Address', and 'Public'. Each entry in the 'Name' column is preceded by a small icon of a function symbol. The 'Address' column shows the memory address for each function, and the 'Public' column indicates if the function is public (marked with 'P'). The list includes various functions such as ATMLAddFontEx, ATMLAddFontExA, ATMLAddFontExW, ATMLAddFontW, ATMLBBoxBaseXYShowTextW, ATMLBeginFontChange, ATMLClient, ATMLEndFontChange, ATMLEnumFonts, ATMLEnumFontsW, ATMLEnumMMFonts, ATMLEnumMMFontsA, ATMLEnumMMFontsW, ATMLFinish, ATMLFontAvailable, ATMLFontAvailableA, ATMLFontAvailableW, ATMLFontSelected, ATMLForceFontChange, ATMLGetBuildStr, and ATMLGetBuildStrA. The status bar at the bottom indicates 'Line 1 of 331'.

Name	Address	Public
ATMLAddFontEx(x,x,x,x,x)	500052EF	P
ATMLAddFontExA(x,x,x,x,x)	500051B8	P
ATMLAddFontExW(x,x,x,x,x)	500049C8	P
ATMLAddFontW(x,x,x,x)	500052FF	P
ATMLBBoxBaseXYShowTextW(x,x,x,x,x,x,x,x,x,x,x)	500017A7	P
ATMLBeginFontChange()	50001649	P
ATMLClient(x)	5000184A	P
ATMLEndFontChange()	500016E2	P
ATMLEnumFonts(x,x)	500043EA	P
ATMLEnumFontsW(x,x)	500043C5	P
ATMLEnumMMFonts(x,x,x)	50004494	P
ATMLEnumMMFontsA(x,x,x)	50004432	P
ATMLEnumMMFontsW(x,x,x)	5000440F	P
ATMLFinish()	50002503	P
ATMLFontAvailable(x,x,x,x,x,x)	5000267E	P
ATMLFontAvailableA(x,x,x,x,x,x)	50002613	P
ATMLFontAvailableW(x,x,x,x,x,x)	50002558	P
ATMLFontSelected(x)	5000268E	P
ATMLForceFontChange()	5000169D	P
ATMLGetBuildStr(x)	50002837	P
ATMLGetBuildStrA(x,x)	500027D0	P

```
signed int __fastcall CallDriver(int a1, int a2, int a3, char a4)
{
    const wchar_t *v4; // esi@1
    signed int v6; // [esp+10h] [ebp-1Ch]@3

    v4 = L"ATMFD.DLL";
    if ( !callMSDriver )
        v4 = L"ATMFDA.DLL";
    v6 = NamedEscape(0, v4, a1, a2, a3, a4 != 0 ? a2 : 0, a4 != 0 ? a3 : 0);
    if ( !v6 )
        v6 = -219;
    return v6;
}
```






xrefs to CallDriver(x,x,x,x)			
Direction	Typ	Address	Text
	p	ATMProperlyLoaded()+23	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMProperlyLoaded()+36	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMBeginFontChange()+29	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMEndFontChange()+47	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetVersion()+25	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMSetFlags(x,x)+38	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p sub_50001F93+AC	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p sub_50001F93+10C	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p sub_50001F93+14D	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMFontAvailableW(x,x,x,x,...	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetFontBBox(x,x)+57	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetBuildStrW(x,x)+17	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetBuildStrA(x,x)+17	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMMakePFMW(x,x,x,x)+D1	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMMakePFMW(x,x,x,x)+10C	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetNtmFieldsW(x,x,x)+...	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p DIIMain(x,x,x)+31	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetFontPathsW(x,x)+99	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetMenuNameW(x,x)+...	call @CallDriver@16 ; CallDriver(x,x,x,x)
	Do...	p ATMGetPostScriptNameW(x...	call @CallDriver@16 ; CallDriver(x,x,x,x)

OK
Cancel
Search
Help

Line 1 of 20

# Reverse engineering escape codes

Name	Escape code	Input data length
ATMProperlyLoaded	0x2502	0
ATMBeginFontChange	0x2503	0
ATMEndFontChange	0x2506	0
ATMFontAvailable, ATMGetPostScriptName	0x2509	194
ATMGetFontBBox	0x250A	12
ATMGetMenuName	0x250B	194
ATMGetGlyphName	0x250C	48
ATMMakePFM	0x250D	>88
ATMGetFontPaths	0x250E	1656
ATMGetVersion	0x250F	0
ATMSetFlags	0x2510	6
?	0x2511	32
?	0x2512	1124
ATMGetNtmFields	0x2513	148
ATMGetGlyphList	0x2514	≥6

-  ATMLIB & ATMFD
-  ATMLIB only
-  ATMFD only

# Googling for the symbol names...

- We can find three extremely interesting documents:
  1. **Adobe Type Manager Software API With Multiple Master Fonts: Macintosh**,  
Technical Note #5074, 14 February 1992, Adobe Systems Incorporated
  2. **Adobe Type Manager Software API: Windows**, Technical Note #5073,  
24 January 1997, Adobe Systems Incorporated
  3. **Adobe Type Manager® Software API for Windows® 95 and Windows NT® 4**,  
Technical Note #5642, 26 June 1998, Adobe Systems Incorporated

# From there...

- Function declarations.
- Structure definitions.
- Constant and enumeration names.
- Overall overview of various ATM mechanics.

# From there... (functions)

```
ATMFontAvailable extern BOOL WINAPI ATMFontAvailable (  
    LPSTR    lpFacename,  
    int      nWeight,  
    BYTE     cItalic,  
    BYTE     cUnderline,  
    BYTE     cStrikeOut,  
    int ATMFAR*lpFromOutline);
```

*Note: Version 1.0*

**ATMFontAvailable()** checks whether a Type 1 font outline is available and can be rendered. All of the parameters except *\*lpFromOutline* correspond to parameters passed to the Windows GDI function **CreateFont()**.

The *lpFacename* parameter is a long pointer to the face name of the font to be verified. **ATMFontAvailable()** returns False if ATM can not respond to a request for the given font. If the function returns True, then ATM can respond in some way to the given font. See the description of the *\*lpFromOutline* parameter, below.

# From there... (structures)

```
ATMFontSpec    typedef struct {  
                    char        faceName [LF_FACESIZE];  
                    WORD        styles;  
                } ATMFontSpec, ATMFAR *LPATMFontSpec;  
  
ATMEnumFontProc  typedef BOOL  
                    (ATMCALLBACK ATMEnumFontProc) (  
                    LPLOGFONT    lpLogFont,  
                    LPSTR        lpPostScriptName,  
                    WORD          flags,  
                    DWORD         dwUserData);  
  
                typedef ATMEnumFontProc ATMFAR *LPATMEnumFontProc;  
  
ATMFontMetricsHeader  typedef struct{  
                    WORD        mmVersion;  
                    WORD        mmFlags;  
                    char        mmCopyright[72];
```

# From there... (constants)

## 7 ATM Return Values and Flags

The following return values, flags, type bits, and flag bits are used by the functions defined in the ATM 4.01 API. They are supported in both the 16-bit and 32-bit libraries, except where noted. Additionally, they are supported for all functions manipulating single- and double-byte fonts.

### 7.1 Return values for non-Boolean functions

**ATM\_NOERR**    0  
The normal return value

**ATM\_INVALIDFONT**    -1  
An invalid font error; the font is not consistent

# Reverse engineering the escape handlers

- With all this, analysis of relevant ATMFD functions becomes much easier.
  - Operation names are roughly known, and they carry information about the escape's functionality.
  - Some structures are fully known, other can be recovered through RE of ATMLIB.DLL.
  - The semantics of ATMFD's return values and other enums are much clearer now.
- We can directly call ATMLIB.DLL functions and do run-time debugging.
- Some strings in ATMFD.DLL can be helpful, as well.



Let's manually audit all 13 escape codes implemented  
in ATMFD.

# GPZ #781

<b>Impact:</b>	Out-of-bound read
<b>Escape code</b>	0x2511
<b>CVE:</b>	None
<b><i>google-security-research</i> entry:</b>	781
<b>Fixed:</b>	WontFix

# Escape code 0x2511

- The escape code is not referenced in ATMLIB.
  - Unknown name or functionality.
- Required input buffer size is 32 bytes:

```
case 0x2511:  
    if ( cbInput == 32 ) {  
        ret = ATMUnspecifiedScramble(lpBuffer);  
        goto label_return;  
    }  
    break;
```

# ATMUnspecifiedScramble

- Doesn't operate on any font objects, only the input data.
- The input structure can be reverse engineered to the following:

```
struct ATM_2511_input {  
    DWORD dword_0;  
    DWORD dword_4;  
    DWORD dword_8;  
    WORD  word_C;  
    WORD  padding;  
    DWORD dwords_10[4];  
};
```

# Unknown logic

```
if (input->dword_0 <= 1) {  
    if (input->dword_8 == 0) {  
        DWORD value = GetUnspecifiedScrambledValue();  
        global_dword_1 = input->dword_8 = value;  
    }  
    if (global_dword_1 != 0) {  
        if (input->word_C > 32) {  
            input->word_C = 32;  
        }  
        for (WORD i = 0; i < input->word_C; i++) {  
            DWORD value = Scramble(input->dwords_10[i], global_dword_1);  
            if (global_dwords_2[i] == 0) {  
                global_dwords_2[i] = value;  
            }  
            global_bools_3[i] = global_dwords_2[i] != value;  
        }  
        global_dword_1 = 0;  
    }  
}
```

# Unknown logic

- `GetUnspecifiedScrambledValue()` transforms a static 32-bit integer with logic/arithmetic operations and returns it.
- `Scramble(x, y)` combines two 32-bit integers into one and returns it.
- The purpose of the logic is undetermined, but also irrelevant.
- Have you noticed that:
  - The `dwords_10` array at the end of the structure only has 4 elements (enforced by the required size of the structure).
  - The function makes it possible to operate on up to 32 elements of the table!

# Out-of-bounds read access

- Accesses to `input->dwords_10[4..31]` are all invalid.
- That's an overread by as much as  $28 \times 4 = 112$  bytes!
- Not particularly useful, could cause a DoS by crashing the kernel.
- But... remember that the NamedEscape surface is available through EMF?
  - The ring-0 out-of-bounds access could be triggered remotely, e.g. through Internet Explorer or Microsoft Office.

# Something's wrong...

- When trying to repro this through IE, I reached the affected code, triggered the out-of-bounds access, but never got a system crash!
  - Even with Special Pools enabled.
- What's up? Wasn't the pool allocation supposed to end up near the end of a page boundary at least once?
- It turned out that the input buffer was not on the pools, but kernel stack!



# Where does the buffer come from?

- Let's look into [win32k!NtGdiExtEscape](#), the top-level handler of the system call:

```
.text:BF822691 loc_BF822691:
.text:BF822691     lea     eax, [ebp+var_3C]
.text:BF822694     mov     [ebp+input_buffer], eax
...
.text:BFAB95AD     cmp     esi, 32
.text:BFAB95B0     jle     loc_BF822691
.text:BFAB95B6     cmp     esi, 2710000h
.text:BFAB95BC     jg      short loc_BFAB95D2
.text:BFAB95BE     push    706D7447h      ; Tag
.text:BFAB95C3     push    esi            ; NumberOfBytes
.text:BFAB95C4     push    21h           ; PoolType
.text:BFAB95C6     call    ds:__imp__ExAllocatePoolWithTag@12
.text:BFAB95CC     mov     [ebp+input_buffer], eax
```

In C:

```
if (NumberOfBytes > 32) {  
    lpBuffer = ExAllocatePoolWithTag(...);  
} else {  
    lpBuffer = &local_buffer;  
}
```

# Close, but no cigar 😞

- The input buffer size must be exactly 32 bytes.
- For all sizes  $\leq 32$  bytes, a local buffer is used for storage.
  - Performance optimization.
- There are always more than 112 bytes (being overread) of stack memory after the local buffer.
  - Higher level stack frames, `KTRAP_FRAME`, padding etc.
  - Due to this extremely unfortunate coincident, a kernel crash may never occur.

# Local information disclosure?

- As the out-of-bounds values are persistently stored (in some form) by ATMFD, it could be possible to extract them back to user-mode.
- Only in a local scenario.
- Not trivial, if at all possible:
  - The values are severely mangled before being saved.
  - There is no obvious route to reading them back through the available interfaces.
- Microsoft classified the issue as WontFix.
  - Non-exploitable by pure accident, but the bug is still there and could become exploitable if conditions change.
- A great example of some very obscure functionality included in the ATMFD escape interface.

# CVE-2016-3220

<b>Impact:</b>	Pool-based buffer overflow
<b>Escape code</b>	0x250C
<b>CVE:</b>	CVE-2016-3220
<b><i>google-security-research</i> entry:</b>	785
<b>Fixed:</b>	MS16-074, 14 June 2016

# ATMGetGlyphName()

- Not an official symbol, but a name assigned based on analysis of ATMLIB.
- Basic facts (on x86):
  - The input buffer size is enforced to be 48 bytes.
  - As the name implies, the function operates on a specific font object, and returns the name of one of its glyphs.
  - **The font is identified by its kernel-mode address, placed at offset 4.**

# Say what?

- A user-mode client identifies a font object with a kernel-mode address.
- Legacy mechanism, implemented as an optimization or to simplify the overall code logic.
- How does the client know the address?

# Obtaining font kernel address

```
PVOID address;  
GetFontData(hdc, 'ebdA', 0, &address, sizeof(PVOID));
```



# GetFontData()

- The function is used to read data from specific SFNT tables of the DC's font file.
  - cmap, head, hhea, hmtx, maxp, etc.
- „ebdA” (backwards for „Adbe”) is a magic table ID, separately handled by ATMFD.
- If the special ID is used and the size of the request is the length of the native word, the kernel-mode font address is returned instead of actual font data.
- Kernel ASLR bypass by design.

# Back on the subject

- The control flow of the escape code handler is deep and complex.
- Let's examine each stage of execution respectively.

# ATMGetGlyphName() step by step #1

1. The i/o buffer size is enforced to be 48 bytes.
2. The font object is located based on the kernel-mode address passed by the client.
3. The font file contents are mapped into memory (?).
4. The function checks if it's a Type 1 or OpenType font.
  - The Type 1 implementation is not particularly interesting, let's follow the OTF one.

# ATMGetGlyphName() step by step #2

5. A function is called with a controlled 16-bit glyph index and a pointer to offset 8 of the i/o buffer (to copy the name there).
  - Let's name it `FormatOpenTypeGlyphName()`.
6. To retrieve the actual glyph name from the .OTF file, another function is used, let's call it `GetOpenTypeGlyphName()`.
  - Here's where the interesting stuff happens.

# GetOpenTypeGlyphName()

- If the glyph ID is between 0 and 390, the name is obtained from a hard-coded list of names:

```
.data:000528E8 off_528E8      dd offset a_notdef_0      ; DATA XREF: GetOpenTypeGlyphName+24↑r
.data:000528E8              ; ".notdef"
.data:000528EC      dd offset aSpace_0      ; "space"
.data:000528F0      dd offset aExclam      ; "exclam"
.data:000528F4      dd offset aQuotedbl     ; "quotedbl"
.data:000528F8      dd offset aNumbersign   ; "numbersign"
.data:000528FC      dd offset aDollar      ; "dollar"
.data:00052900      dd offset aPercent     ; "percent"
.data:00052904      dd offset aAmpersand    ; "ampersand"
.data:00052908      dd offset aQuoteright   ; "quoteright"
.data:0005290C      dd offset aParenleft    ; "parenleft"
.data:00052910      dd offset aParenright   ; "parenright"
.data:00052914      dd offset aAsterisk     ; "asterisk"
.data:00052918      dd offset aPlus        ; "plus"
```

# GetOpenTypeGlyphName()

- Otherwise, the name is extracted from the .OTF file itself, by reading from the Name INDEX.
- String arrays are represented with a list of offsets of consecutive strings.
  - The length of each entry can be determined by subtracting the offset of N+1 and N.

# Name INDEX structure

Type	Name	Description
Card16	count	Number of objects stored in INDEX
OffSize	offSize	Offset array element size
Offset	offset [count+1]	Offset array (from byte preceding object data)
Card8	data[<varies>]	Object data

# GetOpenTypeGlyphName() pseudo-code

```
PushMarkerToStack();
```

```
int glyph_name_offset = ReadCFFEntryOffset(glyph_id);
```

```
int next_glyph_name_offset = ReadCFFEntryOffset(glyph_id + 1);
```

```
*pNameLength = next_glyph_name_offset - glyph_name_offset;
```

```
EnsureBytesAreAvailable(next_glyph_name_offset - glyph_name_offset);
```

```
PopMarkerFromStack();
```



# Internal font stack

- Each font object has an internal array of 16 elements, each 32-bit wide.
- ATMFD debug messages can help us understand their meaning:

```
"fSetPriv->HeldDataKeys[ fSetPriv->nHeldDataKeys-1] == MARK"
```

```
"fSetPriv->nHeldDataKeys >= 0"
```

```
"fSetPriv->nHeldDataKeys > 0"
```

```
"fSetPriv->nHeldDataKeys < MAXHELDATAKEYS"
```


# Internal font stack

- The stack is internally called **He1dDataKeys**.
- The element counter is **nHe1dDataKeys**.
- **MAXHELDDATAKEYS** equals 16.
- A special marker value of -1 is called **MARK**.
- It's still not very clear, what the purpose of the stack is.

# Stack management

- For memory safety, it's important that operations on the stack are balanced.
  - Otherwise, adjacent fields in the font structure, or adjacent allocations on the pools could be overwritten.
- In the code above, it all looks good: **1x PUSH** and **1x POP** afterwards.
  - As long as the functions in between don't perform any stack operations by themselves.

# GetOpenTypeGlyphName() pseudo-code

 Fully controlled

```
PushMarkerToStack();
```

```
int glyph_name_offset = ReadCFFEntryOffset(glyph_id);
```

```
int next_glyph_name_offset = ReadCFFEntryOffset(glyph_id + 1);
```

```
*pNameLength = next_glyph_name_offset - glyph_name_offset;
```

```
EnsureBytesAreAvailable(next_glyph_name_offset - glyph_name_offset);
```

```
PopMarkerFromStack();
```

# EnsureBytesAreAvailable()

- Custom name based on reverse engineering.
- Probably not only ensures bytes are available, but also retrieves them.
- By fully controlling the 32-bit parameter, we can cause it to fail.
- How does it handle failure?

# Error handling

- An exception is generated and handled internally by the function.
- As part of it, all items up to and including -1 are popped from the stack.
  - This interferes with the stack balance, as the element counter is decreased again as part of normal execution.
- After the escape's handler execution, `nHeldDataKeys` is smaller by 1 than before.
  - We can indefinitely set it to -1, -2, -3, ... and write data to those indexes.
  - The result is a pool-based buffer underflow.

# Pool-based buffer underflow

- Persistently decrementing the counter by 1 requires writing 0xffffffff to the current out-of-bounds element.
  - With some pool massaging, this primitive should be sufficient to get arbitrary code execution.
  - Other values may be written to the stack, too (mostly kernel-mode addresses), which should further facilitate exploitation.
- The core of a basic proof of concept is very simple.

```
PVOID address;
GetFontData(hdc, 'ebdA', 0, &address, sizeof(PVOID));

while (1) {
    BYTE buffer[48] = { 0 };
    *(WORD *)&buffer[2] = 391;
    *(PVOID *)&buffer[4] = address;

    NamedEscape(NULL, L"ATMFD.DLL", 0x250C,
                sizeof(buffer), buffer,
                sizeof(buffer), buffer);
}
```



## **SPECIAL\_POOL\_DETECTED\_MEMORY\_CORRUPTION** (c1)

Special pool has detected memory corruption. Typically the current thread's stack backtrace will reveal the guilty party.

Arguments:

Arg1: fe67ef50, address trying to free

Arg2: fe67ee28, address where bits are corrupted

Arg3: 006fa0b0, (reserved)

Arg4: 00000023, caller is freeing an address where nearby bytes within the same page have been corrupted

Debugging Details:

-----

...

STACK\_TEXT:

```
9f4963e4 82930dd7 00000003 c453df12 00000065 nt!RtlpBreakWithStatusInstruction
9f496434 829318d5 00000003 fe67e000 fe67ee28 nt!KiBugCheckDebugBreak+0x1c
9f4967f8 82930c74 000000c1 fe67ef50 fe67ee28 nt!KeBugCheck2+0x68b
9f496818 82938b57 000000c1 fe67ef50 fe67ee28 nt!KeBugCheckEx+0x1e
9f49683c 8293963d fe67ef50 fe67e000 fe67ef50 nt!MiCheckSpecialPoolSlop+0x6e
9f49691c 82973b90 fe67ef50 00000000 fe67ef50 nt!MmFreeSpecialPool+0x15b
9f496984 96a609cc fe67ef50 00000000 fe67ef60 nt!ExFreePoolWithTag+0xd6
9f496998 96b44ec1 fe67ef60 09fe969f 00000000 win32k!VerifierEngFreeMem+0x5b
```

...

# Vulnerability conditions and requirements

- A specially crafted OpenType font must be loaded in the system.
  - Name INDEX with two specific, 32-bit offset entries.
  - Trivial in a local scenario, but could also be possible in a remote one, for targets which support embedded fonts.
- The kernel-mode address of the font object must be specified in the i/o buffer.
  - Not a problem in a local scenario, as shown above.
  - Nearly impossible in a remote scenario due to insufficient interaction capabilities.
  - On 32-bit platforms, there is realistically ~25 unknown bits, so ~33m possible addresses.
    - Maybe could be brute-forced within somewhat realistic file sizes.

# Vulnerability conditions and requirements

- To get any benefit from the memory corruption, pool memory must be massaged, to overwrite some actually meaningful data.
  - Possible in a local scenario, very difficult or nearly impossible in a remote one.
- In summary:
  - *Elevation of Privileges* as a local user.
  - Maybe a DoS with some luck and a specific configuration (x86) in a remote scenario.

# NamedEscape attack surface summary

- Extremely old and obscure communication interface.
  - Bad coding practices, such as sharing ring-0 addresses with ring-3 code.
  - It was probably long forgotten and would likely stay that way if not for the HackingTeam 0-day.
- Unfortunately no browser→kernel exploits found.
  - It was close, and a long shot anyway.
  - Some interesting issues were uncovered, anyway.
  - I also rediscovered the HT vulnerability.
- Audited manually as a whole, but some bugs could obviously still lurk there.
- A prime example of a deep system interface that the EMF files are able to easily touch.

# Auditing GDI+

Out of time, please see the full slide deck released after the conference.


# Hacking VMware Workstation

# EMF in print spooling

- EMF files are also used heavily in print spooling.
- This opens up more format-related attacks vectors, in the form of printer drivers (and other related software).
- One such feasible target is VMware Workstation.

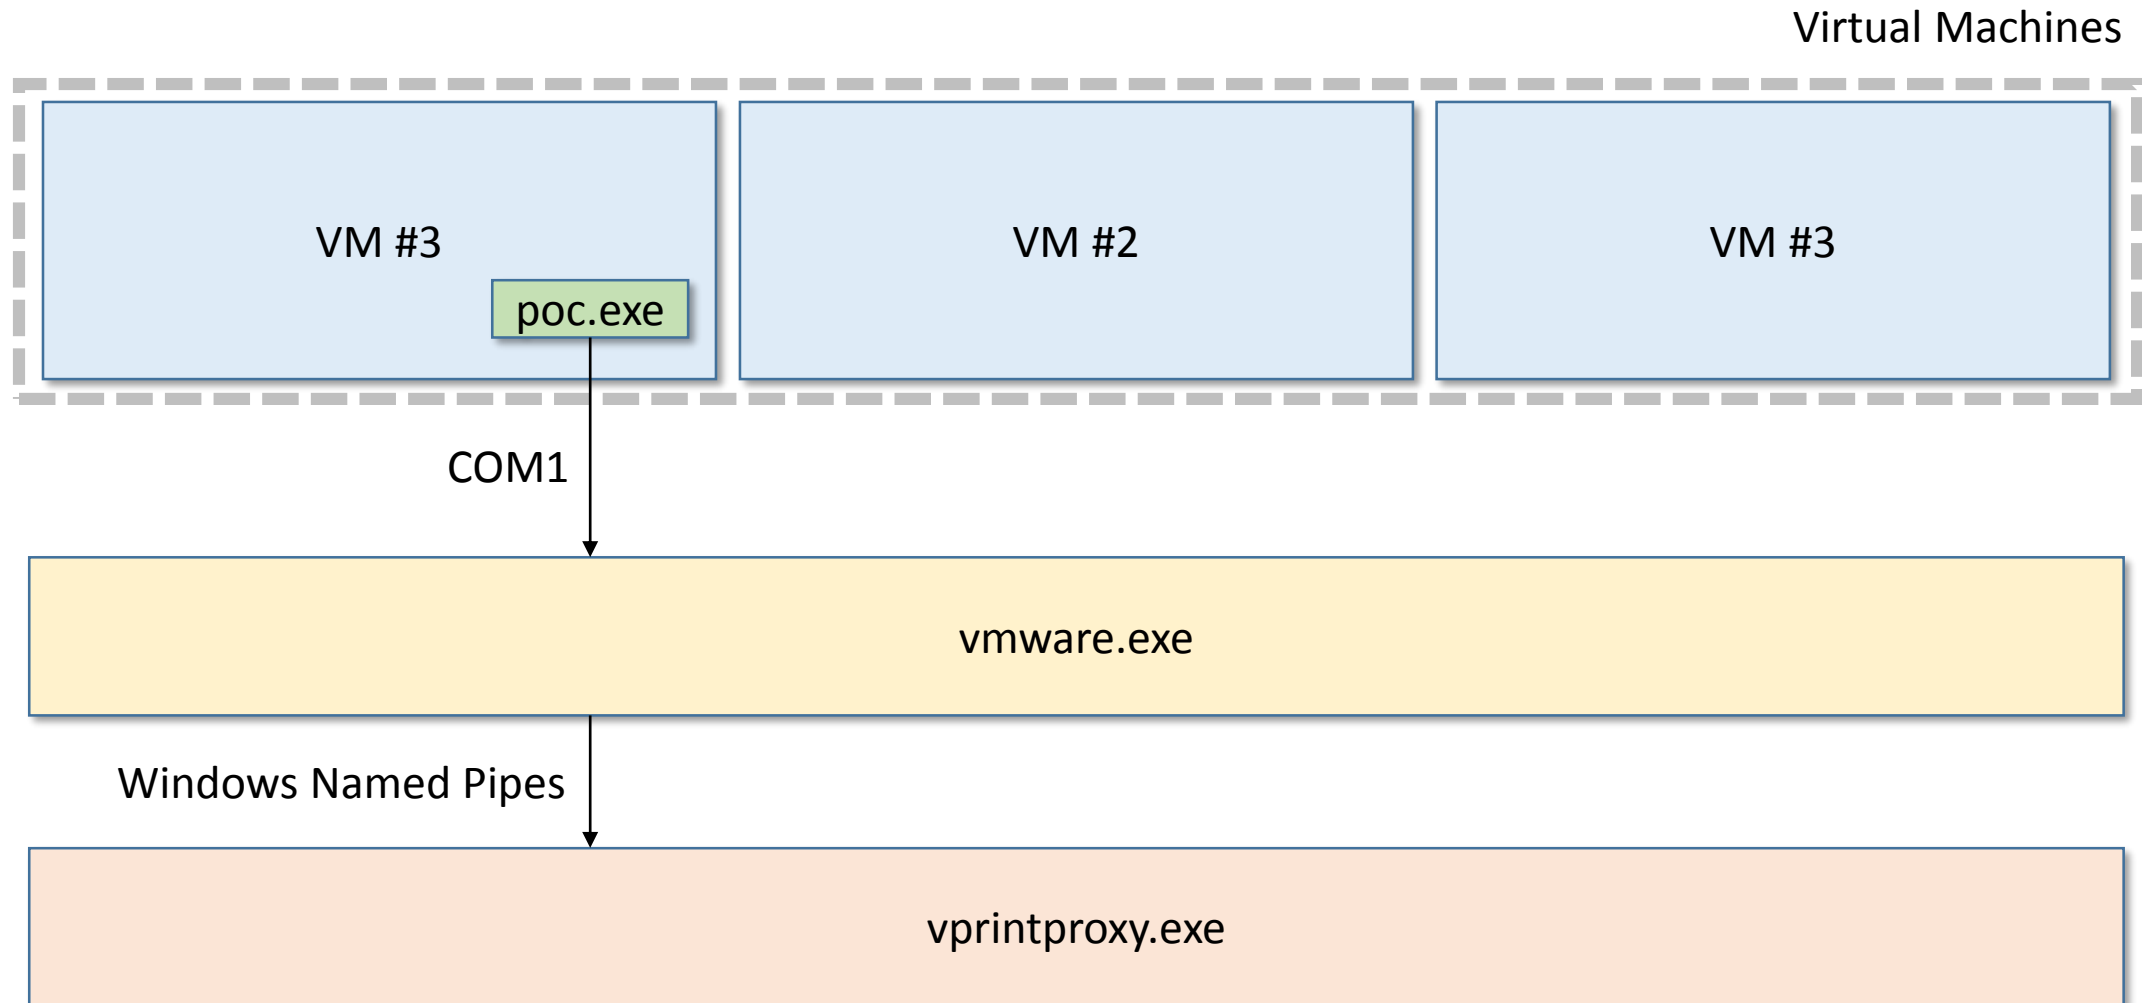
# Virtual printers

- A feature which allows a virtual machine to print documents to printers available on the host (basically printer sharing).
- A feasible VM escape attack vector.
- To my best knowledge, it was enabled by default in 2015, but it's no longer the case (likely thanks to bugs reported by Kostya Kortchinsky).
- Still a frequently used option.





# Architecture



# Architecture

- The attacked process is `vprintproxy.exe` running on the host.
  - Receives almost verbatim data sent by an unprivileged process in a guest system.
  - Quite a communication channel.
- The data is sent in the form of EMFSPOOL files.
  - Similar to EMF, with the extra option to embed fonts in various formats.

# TPView

- More specifically, the most interesting EMF handling takes place in [TPview.dll](#).
  - Together with some other printer-related libraries, they all seem to be developed by a third party, ThinPrint.
- Mostly just falls back to GDI, but also performs specialized handling of several record types.
  - Used to be full of simple bugs, but Kostya found (nearly) all of them!
  - Took another look, discovered a double-free and out-of-bounds `memset()`, but that's all (issues #848 and #849).

# JPEG2000 decoding

- There was one last custom EMF record which seemed completely unexplored.
  - ID = 0x8000.
  - Based on debug strings, it was clear that it was related to JPEG2000 decoding.
- I am no expert at JPEG2K, and the code doesn't seem to be convenient for manual auditing.
- Let's fuzz it?

# Approaching the fuzzing

- Best fuzzing: on Linux, at scale, with AddressSanitizer and coverage feedback.
- After some research, it turns out that the JPEG2000 decoder is authored by yet another vendor, LuraTech.
  - Commercial license, source code not freely available.
- So, are we stuck with TPview.dll wrapped by VMware Workstation?
  - Still feasible, but more complex, slower, and less advanced.

# More research

- After some more digging, I found out that the same vendor released a freeware JPEG2000 decoding plugin for the popular IrfanView program.
  - [JPEG2000.DLL](#).
    - Cursory analysis shows that this is the same or a very similar code base.
- The plugin interface is an extremely simple to use, and resembles the following definition.

```
HGLOBAL ReadJPG2000(IN PCHAR lpFilename,  
                    IN DWORD dwUnknown,  
                    OUT PCHAR lpStatus,  
                    OUT PCHAR lpFormat,  
                    OUT LPDWORD lpWidth,  
                    OUT LPDWORD lpHeight);
```

# Getting there...

- Thanks to this, we can already quickly fuzz-test the implementation in a single process on Windows, without running VMware at all.
  - A wrapper program for loading the DLL and calling the relevant function is <50 LOC long.
- However, I'd really prefer to have this on Linux...



# Fuzzing DLL on Linux

- Why not, really?
- The preferred base address is 0x10000000, which is available in the address space.
  - Relocations not required; sections must be mapped with respective access rights.
- Other actions:
  - Resolve necessary imports.
  - Obtain the address of the exported function.
  - Call it to execute the decoding.
- Should work!

# Resolving imports

- The Import Table may be the only troublesome part.
  - WinAPI functions not available on Linux.
- The DLL imports from **ADVAPI32**, **KERNEL32**, **MSVCRT**, **SHELL32** and **USER32**.
  - C Runtime imports can be directly redirected to libc.
  - All the other ones would have to be rewritten or at least stubbed-out.

# KERNEL32 imports

- Three WinAPI functions used in decoding: **GlobalAlloc**, **GlobalLock** and **GlobalUnlock**:

```
void *GlobalAlloc(uint32_t uFlags, uint32_t dwBytes) __attribute__((stdcall));
void *GlobalAlloc(uint32_t uFlags, uint32_t dwBytes) {
    void *ret = malloc(dwBytes);
    if (ret != NULL) {
        memset(ret, 0, dwBytes);
    }
    return ret;
}

void *GlobalLock(void *hMem) __attribute__((stdcall));
void *GlobalLock(void *hMem) {
    return hMem;
}

bool GlobalUnlock(void *hMem) __attribute__((stdcall));
bool GlobalUnlock(void *hMem) {
    return true;
}
```

# Missing libc imports

- Two MSVCRT-specific imports were found, which had to be reimplemented:

```
long long _ftol(double val) __attribute__((cdecl));  
long long _ftol(double val) {  
    return (long long)val;  
}
```

```
double _CIpow(double x, double y) __attribute__((cdecl));  
double _CIpow(double x, double y) {  
    return pow(x, y);  
}
```

It works!

```
$ ./loader JPEG2000.dll test.jp2
```

```
[+] Successfully loaded image (9b74ba8), format:  
JPEG2000 - Wavelet, width: 4, height: 4
```

# Running the fuzzing

- An internally available JPEG2000 input file corpus was used.
- The mutation strategy was adjusted to hit the 50/50 success/failure rate.
- Left the dumb fuzzer running for a few days, and...
  - ... 186 crashes with unique stack traces were found.

# Crash reproduction

- Keep in mind the crashes are still in the plugin DLL, not VMware Workstation.
- `vprintproxy.exe` is very convenient to use: creates a named pipe and reads exactly the same data that is written to COM1.
  - Once again we can check testcases without starting up any actual VMs.
- PageHeap enabled for better bug detection and deduplication.

# Final results

Instruction	Reason
add [eax+edx*4], edi	Heap buffer overflow
cmp [eax+0x440], ebx	Heap out-of-bounds read
cmp [eax+0x8], esi	Heap out-of-bounds read
cmp [edi+0x70], ebx	Heap out-of-bounds read
cmp [edi], edx	Heap out-of-bounds read
cmp dword [eax+ebx*4], 0x0	Heap out-of-bounds read
cmp dword [esi+eax*4], 0x0	Heap out-of-bounds read
div dword [ebp-0x24]	Division by zero
div dword [ebp-0x28]	Division by zero
fld dword [edi]	NULL pointer dereference
idiv ebx	Division by zero
idiv edi	Division by zero
imul ebx, [edx+eax+0x468]	Heap out-of-bounds read
mov [eax-0x4], edx	Heap buffer overflow
mov [ebx+edx*8], eax	Heap buffer overflow
mov [ecx+edx], eax	Heap buffer overflow
mov al, [esi]	Heap out-of-bounds read
mov bx, [eax]	NULL pointer dereference
mov eax, [ecx]	NULL pointer dereference
mov eax, [edi+ecx+0x7c]	Heap out-of-bounds read

Instruction	Reason
mov eax, [edx+0x7c]	Heap out-of-bounds read
movdqa [edi], xmm0	Heap buffer overflow
movq mm0, [eax]	NULL pointer dereference
movq mm1, [ebx]	NULL pointer dereference
movq mm2, [edx]	NULL pointer dereference
movzx eax, byte [ecx-0x1]	Heap out-of-bounds read
movzx eax, byte [edx-0x1]	Heap out-of-bounds read
movzx ebx, byte [eax+ecx]	Heap out-of-bounds read
movzx ecx, byte [esi+0x1]	Heap out-of-bounds read
movzx ecx, byte [esi]	Heap out-of-bounds read
movzx edi, word [ecx]	NULL pointer dereference
movzx esi, word [edx]	NULL pointer dereference
push dword [ebp-0x8]	Stack overflow (deep / infinite recursion)
push ebp	Stack overflow (deep / infinite recursion)
push ebx	Stack overflow (deep / infinite recursion)
push ecx	Stack overflow (deep / infinite recursion)
push edi	Stack overflow (deep / infinite recursion)
push esi	Stack overflow (deep / infinite recursion)
rep movsd	Heap buffer overflow, Heap out-of-bounds read



# Final results

- Crashes at **39** unique instructions.
  - Many occurring at various points of generic functions such as `memcpy()`, so not the most accurate metric.
  - Quick classification: **18** low severity, **15** medium severity, **6** high severity.
- All reported to VMware on June 15.
- Fixed as part of VMMSA-2016-0014 on September 13 (within 90 days).

Closing thoughts

# Closing thoughts

- Metafiles are complex and interesting files, certainly worth researching further.
  - Supported by a variety of valid attack vectors.
- They can even teach you things about the system API (i.e. the NamedEscape interface).
- As usual, the older and more obscure the format/implementation – the better for the bughunter.
- Inspiration with prior work pays off again.
- The right tool for the right job – manual code auditing vs fuzzing.

# Thanks!



[@j00ru](#)

<http://j00ru.vexillium.org/>

[j00ru.vx@gmail.com](mailto:j00ru.vx@gmail.com)