



Characterization Of The Windows Kernel Version Variability For Accurate Memory Analysis

By

Michael Cohen

Presented At

The Digital Forensic Research Conference

DFRWS 2015 EU Dublin, Ireland (Mar 23rd- 26th)

DFRWS is dedicated to the sharing of knowledge and ideas about digital forensics research. Ever since it organized the first open workshop devoted to digital forensics in 2001, DFRWS continues to bring academics and practitioners together in an informal environment. As a non-profit, volunteer organization, DFRWS sponsors technical working groups, annual conferences and challenges to help drive the direction of research and development.

<http://dfrws.org>

Characterization of the Windows Kernel version variability for accurate Memory analysis.

Michael Cohen

scudette@google.com



DFRWS 2015 Dublin.

Memory Analysis overview

How do we analyse memory?

We need to emulate the way code looks at memory.

Data Structures

```
typedef unsigned char uchar;
```

```
enum {
    OPT1,
    OPT2
} options;
```

```
struct foobar {
    enum options flags;
    short int bar;
    uchar *foo;
}
```

It is generally not possible to predict the memory layout of a C struct without knowing external factors:

- Alignment
- Endianess
- Bit size (64/32 bit)
- Compiler
- Optimizations etc


Unless packed structs.

Example memory analysis technique

- Listing processes
 - Find the global kernel symbol "PsActiveProcessHead"
 - Follow the linked list _EPROCESS. ActiveProcessLinks to find all _EPROCESS structs.
 - Print _EPROCESS.ImageFileName

Kernel Constant

Struct Offset



Struct Offset

Data Structures

```
typedef unsigned char uchar;
```


```
enum {  
    OPT1,  
    OPT2  
} options;
```

```
struct foobar {  
    enum options flags;  
    short int bar;  
    uchar *foo;  
}
```

Debugging symbols contain the exact layout of all data structures.

Can use them to get struct offset AND kernel global constants.

What do PDB files look like?



1. Each time the binary is built, a GUID is generated.
2. The debugging symbols are stored in a PDB file.
3. The executable is shipped.
4. PDB files can be made available publicly on a **symbol server**.

Historical perspective

- Older tools have profiles embedded inside the tool.
 - Profile is pre-generated from an exemplar of an OS version released.
 - e.g. Win7SP1x64
 - Profile is embedded inside the tool
 - We assume profile is applicable to all releases of this version.
 - Profile only contains structs - no use of global offsets from PDB file.
 - Global offsets are deduced by scanning.

Can we always guess kernel globals?



The slide is a presentation slide for Black Hat Europe 2012. It features a dark background with a blue circular logo in the top left corner containing a silhouette of a person wearing a fedora hat. The main title "black hat® EUROPE" is displayed in large white and blue letters. To the right, the event details "March 14-16, 2012" and "NH Grand Krasnapolsky Hotel Amsterdam, Netherlands" are shown, along with the Amsterdam coat of arms. The central text area contains the title of the presentation: "One-byte Modification for Breaking Memory Forensic Analysis" in large white font. Below the title, the authors' names "Takahiro Haruyama / Hiroshi Suzuki" and their affiliation "Internet Initiative Japan Inc." are listed. In the bottom right corner, there is a small map of Europe with a blue dot indicating Amsterdam, and the "black hat" logo. At the bottom left, there is a small note "for submission".

black hat®
EUROPE

March 14-16, 2012
NH Grand Krasnapolsky Hotel
Amsterdam, Netherlands

One-byte Modification for
Breaking Memory Forensic
Analysis

Takahiro Haruyama / Hiroshi Suzuki
Internet Initiative Japan Inc.

for submission

ADD -- Complicating Memory Forensics Through Memory Disarray

Jake Williams and Alissa Torres

In this presentation, we'll present ADD (attention deficit disorder), a tool that litters Windows physical memory with (configurable amounts and types of) garbage to disrupt memory forensics. Memory forensics has become so mainstream that it's catching too many malware authors during routine investigations (making Jake a sad panda). If memory forensics were much harder to perform, then attackers would retain an upper hand. ADD increases the cost of memory forensics by allocating new structures in memory that serve only to disrupt an investigation.

We'll present some basic memory forensics techniques (just to set the stage for those who aren't familiar with the concepts). We'll explain how volatility, a core memory forensics tool, actually performs its analysis. In particular, we'll show how it locates hidden processes, drivers, and modules.

Next, we'll show how running ADD on a machine under investigation completely changes the memory forensics landscape. We'll show how an investigator must weed through astounding numbers of false positives before identifying the investigation targets.

Finally, Alissa will show how all is not lost. Even though ADD may confuse junior analysts, she'll show the invariants in memory that analysts should always be able to come back to complete their forensic analysis.

Jake is the Chief Scientist at CSRgroup where he does lots of offensive and defensive research. He is also a SANS instructor and member of the DFIR author team. Occasionally, CSRgroup still lets Jake do penetration tests (where he feels like a kid in a candy store).

Alissa is a digital forensics examiner and incident response consultant for Sibertor Forensics. Also a SANS Instructor, she teaches hundreds of security professionals a year how to find evil in the form of trace artifacts and hidden processes.

Can we do better?

- Rekall chooses to rely on constants obtained from debugging symbols.
 - Pros
 - Better coverage of symbols, especially ones that are not exported.
 - Cons
 - We need to wait for Microsoft to make a pdb file available for us to use.

Lets evaluate this approach.

- Basic assumption in Volatility:
 - Struct layout does not change between major and minor versions.
 - An exemplar from a particular version will apply to all kernels from that version.
 - Kernel global symbols vary too much between major and minor versions to hard code
 - Therefore we need to scan for them.



This repository ▾

Search or type a command



Explore Gist Blog Help



scudette

google / **rekkal-profiles**

Watch ▾ 3

Star 1

Fork 1

branch: master ▾

rekkal-profiles / v1.0 / nt / GUID / +



Updated repo index, removed obsolete html files.

the80srobot authored on Jun 6


latest commit 0142f0aad1

..


00625D7D36754CBEBAA4533BA9A0F3FE22.gz	Initial push of profiles into the profile repository.	3 months ago
0100FCDAFD4049B8B06005EC07705A1F2.gz	Initial push of profiles into the profile repository.	3 months ago
01DDCBD82AEB46BEAFCDC6A409E3B1D31...	Initial push of profiles into the profile repository.	3 months ago
01DF28C698D84DEBB1A74254C3AF800E2.gz	Initial push of profiles into the profile repository.	3 months ago
03185083233249D9BB747EA777B80C982.gz	Initial push of profiles into the profile repository.	3 months ago
04FB9A156FF44ECCA6EBCAE9617D8DB73.gz	Initial push of profiles into the profile repository.	3 months ago
05A6F49C5DD848FF983459421A78F1232.gz	Initial push of profiles into the profile repository.	3 months ago
06472CCD0ECF43B58D676891C6745DAC2.gz	Initial push of profiles into the profile repository.	3 months ago
0887873AD7FC4115AC9258B9871F81341.gz	Initial push of profiles into the profile repository.	3 months ago
08F4D00C3B5A4B34B2D3AE8402F927802.gz	Initial push of profiles into the profile repository.	3 months ago




Windows Struct variations with Kernel Version (AMD64)



Windows Struct variations with Kernel Version (I386)





Result

- Assumption is mostly validated for the kernel
 - Struct offsets do not vary per version.
 - Kernel constant very wildly.
- So the Volatility approach should work in most cases!



An advanced memory forensics framework

Search projects

[Project Home](#) [Export to GitHub](#)

[New issue](#) [Search](#) [Open issues](#) for [Search](#) [Advanced search](#) [Search tips](#) [Subscriptions](#)

★ Issue 174: [profile offsets per build/revision number] was: Process Owner Info not found for TCP_ENDPOINT scan output on Vista (netscan)

8 people starred this issue and may be notified of changes.

[Back to list](#)

Status: Duplicate

Merged: [issue 521](#)

Owner: michael.hale@gmail.com

Closed: Feb 18

Cc: michael.hale@gmail.com,
mike.auty@gmail.com,
jamie.l...@gmail.com,
labaru...@gmail.com,
atc...@gmail.com

Type-Defect

Priority-Low

Milestone-3.0.x

[Add a comment below](#)

Reported by welcome....@gmail.com, Jan 3, 2012

What steps will reproduce the problem?

1. Running netscan plugin on Vista SP2 dump
- 2.
- 3.

What is the expected output? What do you see instead?

PID, Process name is not found TCPEndPoint scan output. However that information is present in Listener, UDP scan output.

0xcf7c0de8	TCPv4	0.0.0.0:61618	0.0.0.0:0	LISTENING	2016	spoolsv.exe
0x1cedd88	TCPv4	0.0.0.0:57883	0.0.0.0:443	CLOSED	-----	-----
0x14e85e08	TCPv4	10.130.179.251:64286	10.177.226.18:1533	ESTABLISHED	-----	-----

What version of the product are you using? On what operating system?

using volatility 2 standalone python precompiled version on vista sp2.

Please provide any additional information below.

Google |



Project Member #17 michael.hale@gmail.com

May 18, 2012

Here's a little more info.

6.0.6002.18272

Microsoft Windows Server 2008 Standard
6.0.6002 Service Pack 2 Build 6002
push dword ptr [edi+164h]
call ds:_imp__PsGetProcessId@4

6.0.6002.18519

Microsoft Windows VistaT Enterprise
6.0.6002 Service Pack 2 Build 6002
push dword ptr [edi+164h]
call ds:_imp__PsGetProcessId@4

6.0.6002.18005

Microsoft Windows Server 2008 Standard (also seen Microsoft Windows VistaT Ultimate)
6.0.6002 Service Pack 2 Build 6002
push dword ptr [edi+160h]
call ds:_imp__PsGetProcessId@4

As you can see, regardless of whether the OS identifies itself as "Server 2008" or "Vista" if the revision number > 18005 then the offset is 0x164. If the revision number <= 18005 then the offset is 0x160. I don't have tcpip.sys binaries for every revision, so I'm not sure if 18005 is exactly where the line is drawn.

Unfortunately until we can choose vtype offsets based on revision number (in addition to the major, minor, build, and memory model which is already possible), then there's not a good way to handle this. I'm not sure we should close this issue since currently we won't print process information for VistaSP2x86/Win2008SP2x86 whose revision numbers are > 18005. However, we might have to defer to fixing it later. In the meantime, manish, you'll have to use a version of volatility where you can change the vtype offset (and if needed build your own exe from it).

Project Member #18 michael.hale@gmail.com

Mar 7, 2014

(No comment was entered for this change.)

Summary: [profile offsets per build/revision number] was: Process Owner Info not found for TCP_ENDPOINT scan output on Vista (netscan) (was: Process Owner Info not found for TCP_ENDPOINT scan output on Vista (netscan))


Cc: -scude...@gmail.com

There are some problems

- Sometimes even struct layout changes within the major/minor version release cycle.
- Mr. Hale is an expert reverse engineer
 - He can figure out the correct struct layout by looking at the disassembly.
 - He knows how to change the program to account for this version.
 - We don't really know how to reproduce:
 - What function was reversed?
 - Which instructions should we look at?

We need automated reversing

- First approach:
 - Look at the data and surrounding context



```
[1] win7.elf.E01 07:20:02> dump "*win32k!grpWinStaList"
-> dump("*win32k!grpWinStaList")
```

Offset	Data
0xfa800225ef60	00 00 00 00 00 00 00 00 00 70 1a 85 01 80 fa ff ff
0xfa800225ef70	30 4b 2c 02 80 fa ff ff 40 f3 3a 00 60 f9 ff ff OK,.....@..`....
0xfa800225ef80	00 00 00 00 00 00 00 00 00 f0 c9 15 c0 00 f9 ff ff
0xfa800225ef90	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225efa0	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225efb0	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225efc0	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225efd0	00 00 00 00 00 00 00 00 00 20 10 87 02 a0 f8 ff ff
0xfa800225efe0	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225eff0	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0xfa800225f000	00 00 58 02 6d 6f 6e 69 eb 25 8b 02 00 00 00 ..X.moni.%....
0xfa800225f010	10 f0 25 02 80 fa ff ff 70 7e 92 00 80 fa ff ff ..%.....p~.....
0xfa800225f020	b0 c4 f0 00 80 f8 ff ff 02 10 40 03 03 00 00 00@....
0xfa800225f030	70 7e 92 00 80 fa ff ff 1a 0c 01 00 00 00 00 00 p~.....

```
[1] win7.elf.E01 07:29:07> analyze_struct "*win32k!grpWinStaList"
-----> analyze_struct("*win32k!grpWinStaList")
0xfa800225ef60 is inside pool allocation with tag 'Win\xe4' (0xfa800225eed0)
Offset      Content
-----
0x0 Data:0x0
0x8 Data:0xfa8001851a70 Tag:Win\xe4 @0xfa8001851a70
0x10 Data:0xfa80022c4b30 Tag:Des\xeb @0xfa80022c4b30
0x18 Data:0xf960003af340 Const:win32k!gTermIO
0x20 Data:0x0
0x28 Data:0xf900c015c9f0 Tag:Uskb @0xf900c015c9f0
0x90 Data:0x0
0x98 Data:0x0
0xa0 Data:0x6f6d02580000
0xa8 Data:0x28b25eb
0xb0 Data:0xfa800225f010 Empty Tag:moni @0xfa800225f010
0xb8 Data:0xfa8000927e70 Tag:FxDr @0xfa8000927e70
0xc0 Data:0xf88000f0c4b0
0xc8 Data:0x303401002
0xd0 Data:0xfa8000927e70 Tag:FxDr @0xfa8000927e70
0xd8 Data:0x10c1a
0xe0 Data:0xfa80009437b0 _LIST_ENTRY @0xfa80009437b0 Tag:moni @0xfa80009437b0
0xe8 Data:0xfa8000920460
0xf0 Data:0x0
0xf8 Data:0xfa8000919e70 Tag:moni @0xfa8000919e70
```

```
[1] win7.elf.E01 07:21:00> win32k_autodetect
-----> win32k_autodetect()
DEBUG:root:Listed 41 processes using PsActiveProcessHead
DEBUG:root:Listed 37 processes using CSRSS
DEBUG:root:Listed 41 processes using PspCidTable
DEBUG:root:Listed 39 processes using Sessions
DEBUG:root:Listed 40 processes using Handles
DEBUG:root:Switching to process context: System (Pid 4@0xfa80008959e0)
DEBUG:root:Switching to process context: svchost.exe (Pid 236@0xfa80024f85d0)
DEBUG:root:Checking tagWINDOWSTATION at 0xfa800225ef60
DEBUG:root:Unhandled field 0x0, ['Data:0x0']
DEBUG:root:Detected field rpwinstaNext: ['Data:0xfa8001851a70', 'Tag:Win\xe4', '@0xfa8001851a70'] @ 0x8
DEBUG:root:Detected field rpdeskList: ['Data:0xfa80022c4b30', 'Tag:Des\xeb', '@0xfa80022c4b30'] @ 0x10
DEBUG:root:Detected field pTerm: ['Data:0xf960003af340', 'Const:win32k!gTermIO'] @ 0x18
DEBUG:root:Unhandled field 0x20, ['Data:0x0']
DEBUG:root:Unhandled field 0x28, ['Data:0xf900c015c9f0', 'Tag:Uskb', '@0xf900c015c9f0']
DEBUG:root:Unhandled field 0x30, ['Data:0x0']
DEBUG:root:Unhandled field 0x38, ['Data:0x0']
DEBUG:root:Unhandled field 0x40, ['Data:0x0']
DEBUG:root:Unhandled field 0x48, ['Data:0x0']
DEBUG:root:Unhandled field 0x50, ['Data:0x0']
DEBUG:root:Unhandled field 0x58, ['Data:0x0']
DEBUG:root:Unhandled field 0x60, ['Data:0x0']
DEBUG:root:Unhandled field 0x68, ['Data:0x0']
DEBUG:root:Unhandled field 0x70, ['Data:0x0']
DEBUG:root:Detected field pGlobalAtomTable: ['Data:0xf8a002871020', 'Tag:AtmT', '@0xf8a002871020'] @ 0x78
DEBUG:root:Unhandled field 0x80, ['Data:0x0']
DEBUG:root:Unhandled field 0x88, ['Data:0x0']
DEBUG:root:Unhandled field 0x90, ['Data:0x0']
DEBUG:root:Unhandled field 0x98, ['Data:0x0']
DEBUG:root:Unhandled field 0xa0, ['Data:0xf6d02580000']
DEBUG:root:Unhandled field 0x8, ['Data:0x28b25eb']
DEBUG:root:Unhandled field 0xb0, ['Data:0xfa800225f010', 'Empty', 'Tag:moni', '@0xfa800225f010']
DEBUG:root:Unhandled field 0xb8, ['Data:0xfa8000927e70', 'Tag:FxD', '@0xfa8000927e70']
DEBUG:root:Unhandled field 0xc0, ['Data:0xf88000f0c4b0']
DEBUG:root:Unhandled field 0xc8, ['Data:0x303401002']
DEBUG:root:Unhandled field 0xd0, ['Data:0xfa8000927e70', 'Tag:FxD', '@0xfa8000927e70']
DEBUG:root:Unhandled field 0xd8, ['Data:0x10c1a']
DEBUG:root:Unhandled field 0xe0, ['Data:0xfa80009437b0', '_LIST_ENTRY', '@0xfa80009437b0', 'Tag:moni', '@0xfa80009437b0']
DEBUG:root:Unhandled field 0xe8, ['Data:0xfa8000920460']
DEBUG:root:Unhandled field 0xf0, ['Data:0x0']
DEBUG:root:Unhandled field 0xf8, ['Data:0xfa8000919e70', 'Tag:moni', '@0xfa8000919e70']
DEBUG:root:Unhandled field 0x100, ['Data:0xfa8000919e90', '_LIST_ENTRY', '@0xfa8000919e90', 'Tag:moni', '@0xfa8000919e90']
DEBUG:root:Unhandled field 0x108, ['Data:0xfa8000919e90']
```

***** Struct tagWINDOWSTATION *****

field	offset	Definition
rpwinstaNext	0x8	['Pointer', {'target': 'tagWINDOWSTATION'}]
rpdeskList	0x10	['Pointer', {'target': 'tagDESKTOP'}]
pTerm	0x18	['Pointer', {'target': 'tagTERMINAL'}]
pGlobalAtomTable	0x78	['Pointer', {'target': '_RTL_ATOM_TABLE'}]

***** Struct tagTHREADINFO *****

field	offset	Definition
pEThread	0x0	['Pointer', {'target': '_ETHREAD'}]
GdiTmpTgoList	0x50	[_LIST_ENTRY]
ppi	0x158	['Pointer', {'target': 'tagPROCESSINFO'}]
pq	0x160	['Pointer', {'target': 'tagQ'}]
spklActive	0x168	['Pointer', {'target': 'tagKL'}]
rpdesk	0x178	['Pointer', {'target': 'tagDESKTOP'}]
PtiLink	0x260	[_LIST_ENTRY]

***** Struct tagDESKTOP *****

field	offset	Definition
rpdeskNext	0x18	['Pointer', {'target': 'tagDESKTOP'}]
rpwinstaParent	0x20	['Pointer', {'target': 'tagWINDOWSTATION'}]
_hsectionDesktop	0x78	['Pointer', {'target': '_SECTION_OBJECT'}]
PtiList	0xa8	[_LIST_ENTRY]

2: Reverse Engineering code

- We need a reproducible and robust reverse engineering method
 - Expert makes the initial reversing analysis
 - Machine parseable method of documenting the finding.
 - Repeatable analysis on similar code variants.

tagDESKTOP:

PtiList:

- - Disassembler

- rules:

- MOV \$var1, *grpdeskRitInput
- TEST \$var1, \$var1
- MOV \$var1, [\$var1+\$rpwinstaParent]
- MOV \$pdesk, [\$var1+\$rpdeskList]
- LEA *, [\$pdesk+\$out]

start: win32k!SetGlobalCursorLevel

target: Pointer

max_separation: 300

pheapDesktop:

- - Disassembler

- rules:

- MOV \$var1, [*+\$out]
 - CALL *RtlAllocateHeap
- start: win32k!DesktopAlloc
target: Pointer

rpdeskNext:

- - Disassembler

- rules:

- MOV \$var1, [\$var2+\$out]
- TEST \$var1, \$var1
- JZ *

- MOV *CX, \$var1

- CALL *ObQueryNameInfo


start: win32k!ParseDesktop
target: Pointer

```
[1] win7.elf.E01 09:39:47> print session.profile.tagDESKTOP()
ERROR:root:Failed to find match for tagDESKTOP.rpdeskNext.
DEBUG:root:Unable to find tagDESKTOP.rpdeskNext via Disassemble win32k!ParseDesktop
DEBUG:root:Found match for tagDESKTOP.rpdeskNext
DEBUG:root... 0xf960001c4545      0x4d 488b4e18      MOV RCX, [RSI+0x18]
DEBUG:root... 0xf960001c4549      0x51 4885c9      TEST RCX, RCX
DEBUG:root... 0xf960001c4553      0x5b ff15c7balaa00    CALL QWORD [RIP+0x1abac7]  0xfffffff80002695fe0 win32k!imp_0bfDereferenceObject ->
Object
DEBUG:root:Found match for tagDESKTOP.PtiList
DEBUG:root: 0xf96000206452      0x6 488b05579cla00    MOV RAX, [RIP+0x1a9c57]  0xfffffa80022c4b30 win32k!grpdeskRitInput
DEBUG:root: 0xf9600020645b      0xf 4885c0      TEST RAX, RAX
DEBUG:root: 0xf96000206460      0x14 488b4020      MOV RAX, [RAX+0x20]
DEBUG:root: 0xf96000206464      0x18 488b5010      MOV RDX, [RAX+0x10]
DEBUG:root... 0xf9600020646a      0x1e 4c8d82a8000000    LEA R8, [RDX+0xa8]
ERROR:root:Failed to find match for tagDESKTOP.rpwinstaParent.
DEBUG:root:Unable to find tagDESKTOP.rpwinstaParent via Disassemble win32k!SetGlobalCursorLevel
DEBUG:root:Found match for tagDESKTOP.pheapDesktop
DEBUG:root.. 0xf960001924a5      0x11 488b8980000000    MOV RCX, [RCX+0x80]
DEBUG:root.. 0xf960001924b1      0x1d ff1521e01d00    CALL QWORD [RIP+0x1de021]  0xfffff8000265e840 win32k!imp_RtlAllocateHeap -> nt!Rt
[tagDESKTOP tagDESKTOP] @ 0x00000000
0x00 rpwinstaParent <None Pointer to [0x00000000] (rpwinstaParent)>
0x18 rpdeskNext   <None Pointer to [0x00000000] (rpdeskNext)>
0x80 pheapDesktop <None Pointer to [0x00000000] (pheapDesktop)>
0xA8 PtiList     <None Pointer to [0x00000000] (PtiList)>
```

Repeat analysis with undocumented structures



Win32k.sys constants variations with Version (AMD64)



-
- Scanning for constants in undocumented code is extremely difficult.
 - Often requires the automated disassembly of function call sites.
 - Often loses context (e.g. which session does this win32k object belong to?).
 - Quite slow.
 - We want to be able to use known versions
 - Just extract constants from the PDB files.

What does a Rekall profile look like?

```
{ "$CONSTANTS": {  
    "CmNtCSDVersion": 718856,  
    ...  
    "$ENUMS": {  
        "BUS_QUERY_ID_TYPE": {  
            "0": "BusQueryDeviceID",  
            "1": "BusQueryHardwareIDs",  
            ...  
        }  
    }  
    "$FUNCTIONS": {  
        "ADD_MAP_REGISTERS": 606670,  
        ...  
    }  
    "$METADATA": {  
        "ProfileClass": "Nt",  
        "arch": "I386"  
        ...  
    }  
    "$STRUCTS": {  
        "BATTERY_REPORTING_SCALE": [8, {  
            "Capacity": [4, ["unsigned long", {}]], ...  
        }]  
    }  
}
```

- File is a JSON data structure.
- Divided into Sections:
 - \$CONSTANTS
 - \$FUNCTIONS
 - \$METADATA
 - \$STRUCTS

\$STRUCT section.

```
{ "_EPROCESS": [624, {  
    "AccountingFolded": [548, ["BitField", {  
        "end_bit": 2,  
        "start_bit": 1,  
        "target": "unsigned long"  
    }]],  
    "ActiveProcessLinks": [160, ["_LIST_ENTRY", {}]],  
    "ActiveThreads": [376, ["unsigned long", {}]],  
    "AddressCreationLock": [232, ["_EX_PUSH_LOCK", {}]],  
    "AddressSpaceInitialized": [552, ["BitField", {  
        "end_bit": 12,  
        "start_bit": 10,  
        "target": "unsigned long"  
    }]],  
    "AffinityPermanent": [548, ["BitField", {  
        "end_bit": 19,  
        "start_bit": 18,  
        "target": "unsigned long"...  
    }]]  
},  
    "Struct Size": [624],  
    "Member Offset": [548],  
    "Member Type": ["BitField"],  
    "Arguments to the member type": [{}]  
}
```

\$CONSTANTS and \$FUNCTIONS

```
"NtAlpcSendWaitReceivePort": 2207436,  
"NtAlpcSetInformation": 1805611,  
"NtApphelpCacheControl": 2308968,  
"NtAreMappedFilesTheSame": 2367400,  
"NtAssignProcessToJobObject": 1912487,  
"NtBuildGUID": 411132,  
"NtBuildLab": 410688,  
"NtBuildLabEx": 410912, ...
```

Constant name

Constant offset (Relative to
the kernel base).

- These addresses come directly from Microsoft Debugging symbols.
 - Identical to the way the kernel debugger works.
 - No need to scan, guess or otherwise deduce symbol addresses.

Rekall Profiles - JSON files

- A profile file is a data structure which represents all the information needed to parse OS specific memory.
 - Files are stored in the public profile repository:
 - <http://profiles.rekall-forensic.com>
 - Windows Profiles are identified by GUID.

Revision c39b14f8dca9: /nt/GUID

[[Project Page](#)]

- ..
- [00625D7D36754CBEBA4533BA9A0F3FE22.gz](#)
- [0100FCDAFD4049B8B06005EC07705A1F2.gz](#)
- [01DDCBD82AEB46BEAFCDC6A409E3B1D31.gz](#)
- [01DF28C698D84DEBB1A74254C3AF800E2.gz](#)
- [03185082233249D9BB747EA777B80C982.gz](#)
- [04FB9A156FF44ECCA6EBCAE9617D8DB73.gz](#)
- [05A6E49C5DD848FF983459421A78E1232.gz](#)

Profiles for nt kernel
are stored here.

Every single kernel
build has a unique
GUID.

Profile Indexes


Problem statement: Set membership:

- Given a set of binaries, is this binary in the set, and which one is it?

Solution:

- Generate a sample of data points in each binary and build a decision tree.

Conceptual overview



Potential Complications

- Not all pages in binary are always readable:
 - Must discount comparison points in unreadable pages -> weakens signatures.
- How many points should we use?
- Sometimes we do not have the actual binary - we only have the binary GUID.
 - Deduce data in binary purely from symbol information:
 - Functions have known preamble.
 - String constants have debug symbols.

path: nt
symbols:

-
 name: "str:KernelSpace"
 data:
 - "str:KernelSpace"

Symbol Name (Usually mangled)

-
 name: "str:ZwQueryInformationFile"
 data:
 - "str:ZwQueryInformationFile"

Probably what the address contains.

name: ExEnumHandleTable
data: ["90"]
shift: -1

-
 name: FsRtlAllocateFileLock
 data: ["90"]
 shift: -1

Before each function - NOP slide.

Result index

```
{  
  "$INDEX": {  
    "nt/GUID/00625D7D36754CBEBA4533BA9A0F3FE22": [[2038160, ["4b65726e656c5370616365"]],  
    [3601204, ["4952505f4d4e5f51554552595f4445564943455f54455854"]], [253980, ["410\\  
    0500050005f004e0041004d004500"]], [120086, ["90"]], [2559256, ["90"]], [137962, ["90"]], [2500200,  
    ["90"]], [2569084, ["90"]], [206055, ["90"]], [630264, ["90"]], [10\\  
    4589, ["90"]]],  
    "nt/GUID/0100FCDAFD4049B8B06005EC07705A1F2": [[463544,  
    ["5a775175657279496e666f726d6174696f6e46696c65"]], [2197832,  
    ["4952505f4d4e5f51554552595f4445564943455f544558\\  
    54"]], [206384, ["4100500050005f004e0041004d004500"]], [299937, ["90"]], [1174215, ["90"]],  
    [113888, ["90"]], [1339912, ["90"]], [1746854, ["90"]], [236722, ["90"]], \\  
    [264355, ["90"]]]]  
  }  
}
```

Profile Name

Offset in binary.


Expected data.

Conclusions

- Validating assumptions about kernel versions.
 - Will our analysis work in every case?
 - Maybe not
- Develop methods for automated reverse engineering
 - Helps to document expert effort.
 - Helps to repeat on many samples.
- Towards fully automated Linux profile generation!
 - Given a binary kernel image, calculate the correct profile automatically.

<http://www.rekall-forensic.com/>

Sorry, Quaid. Your whole life is just a dream.



See you at the party, Richter!