



Hyntrospect: A Fuzzer for Hyper-V Devices


Diane Dubois - didu@google.com -  @0xdidu

SSTIC

June 2 2021

Whoami?

- Security Engineer at Google
- 20% with Project Zero

- Passionate about vulnerability research on systems
-  @0xdidu

Why Hyper-V?

- Project Zero mission: aims to reduce harm caused by targeted attacks on the Internet
- Hyper-V
 - Hypervisor running Azure, Microsoft Cloud
 - Modern versions of Windows run it (Virtualization-based security)
 - Possibly a high impact if there are 0-days in the wild
- ... and because virtualization is a fun topic
 - Spans multiple layers from hardware to high level software
 - Some complex implementations



Goals

- Instrumenting Hyper-V for vulnerability research
 - A fuzzer called Hyntrospect was developed and open sourced <https://github.com/googleprojectzero/Hyntrospect>
 - Coverage-guided
 - On closed-source binaries
 - Pragmatic approach, using existing Hyper-V features and Windows tools
 - In a real execution environment
 - Its internals and current results (coverage...) are presented
- Finding vulnerabilities and reporting them to Microsoft

Agenda

- Background on Hyper-V
- The Research Target
- Hyntrospect fuzzer
- Current results
- Future endeavours





Microsoft
Hyper-V



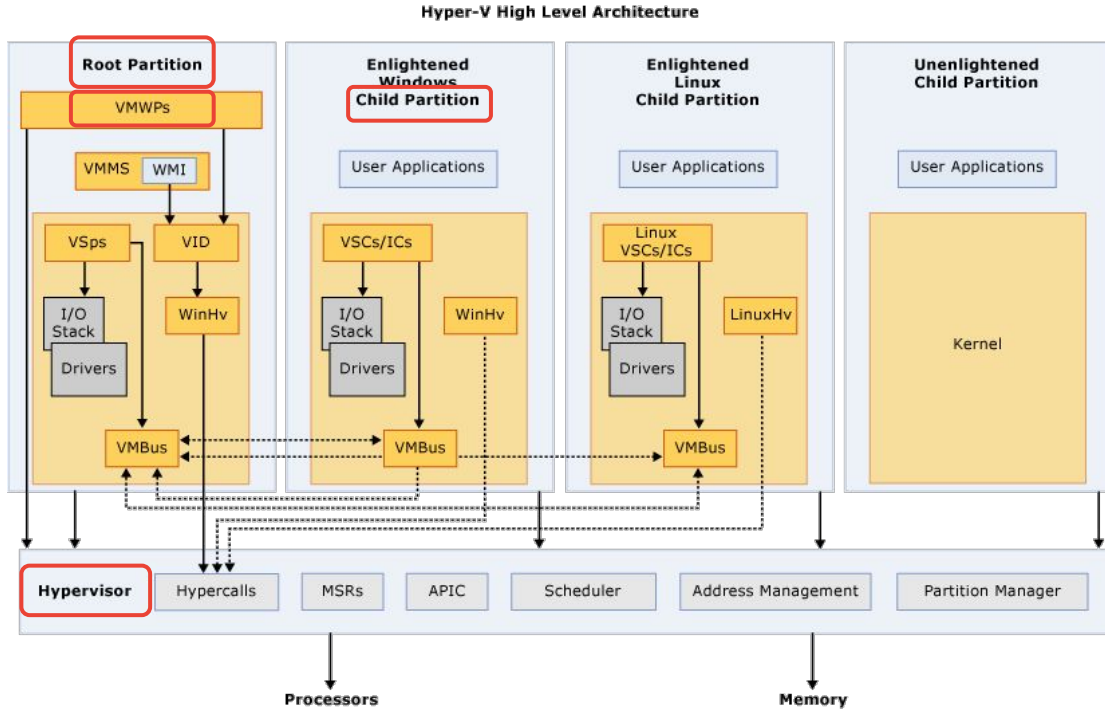
Background on Hyper-V



```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



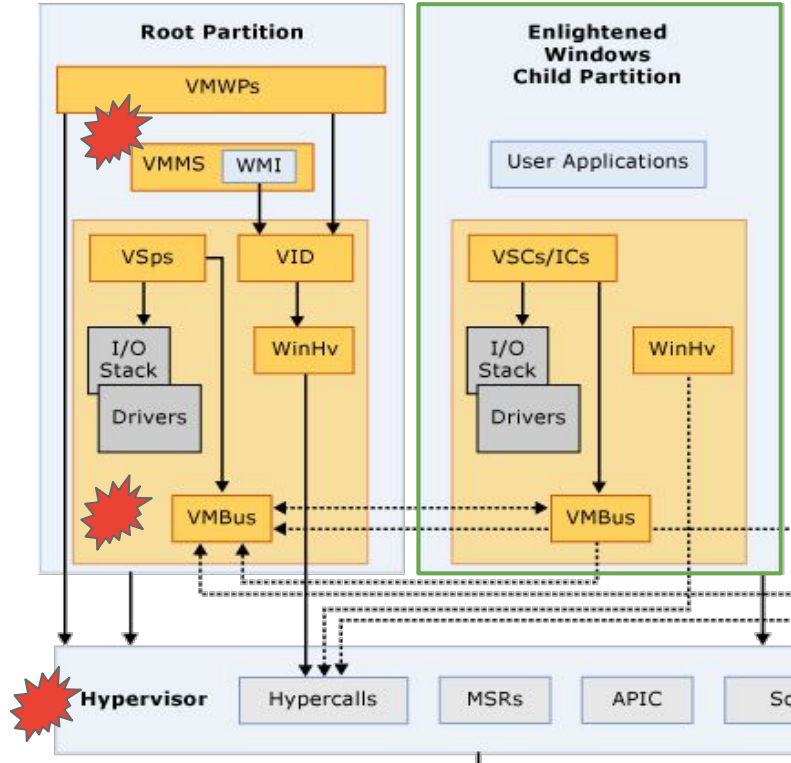
Hyper-V Architecture Overview



Source: Microsoft

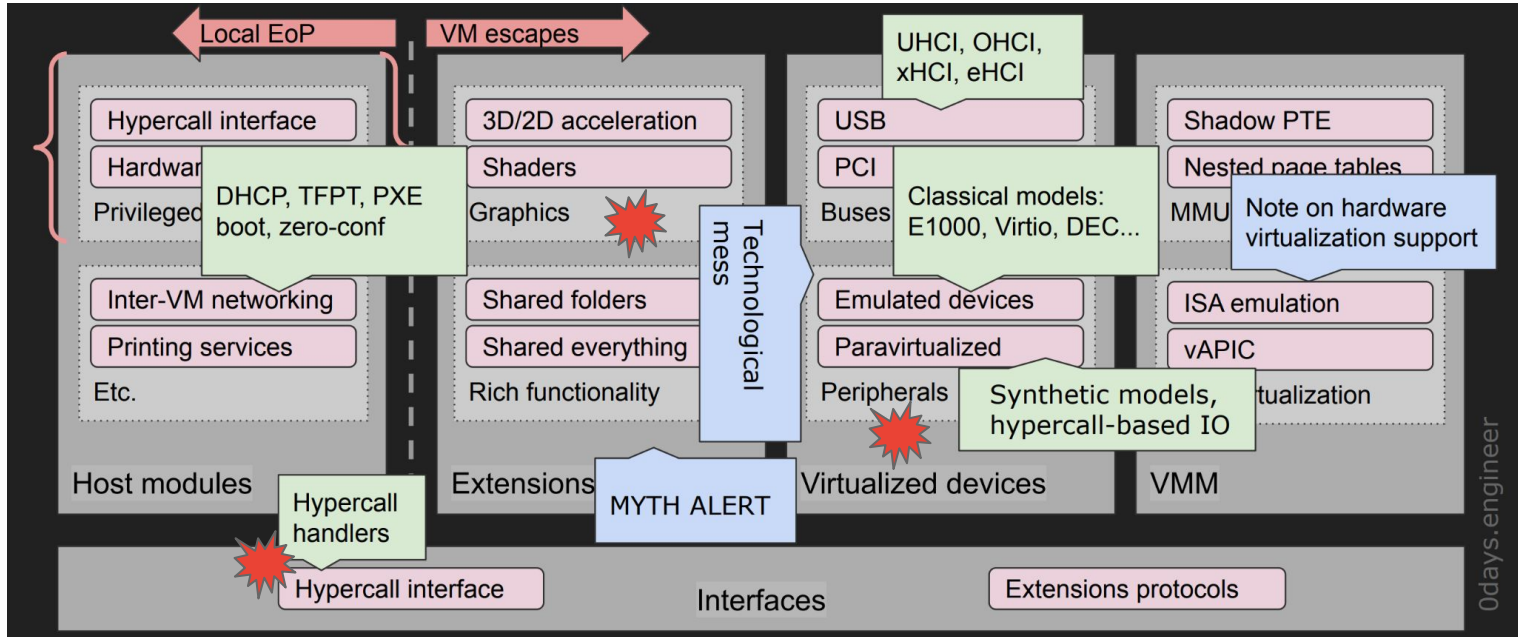
What is a “Guest to Host Escape”?

- Gaining **code execution** on one of the **hypervisor layers** **from a virtual machine**
- On Hyper-V: ambiguous
 - Hypervisor layer
 - Root partition (kernel / userland)
- Other type of attack: **host denial of service**

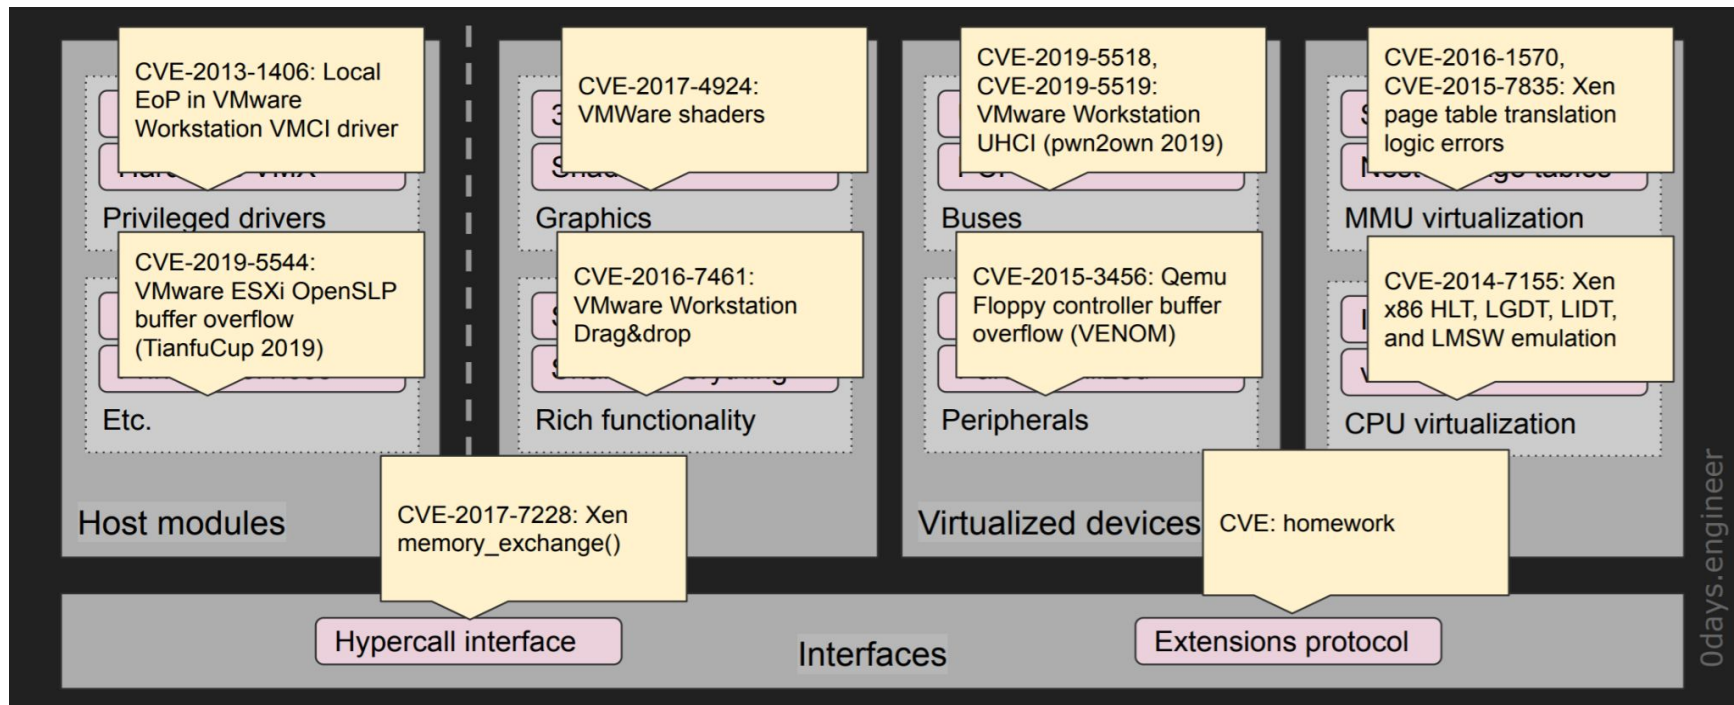


The Hypervisors' Attack Surface

- As defined by Alisa Esage (Zer0Con 2020):



... and in practice



Hyper-V Attack Surface

- Hypervisor
 - Hypercall handlers
 - Faults
 - Instruction emulation
 - Intercepts
 - Register access (MSRs...)
- Root partition kernel attack surface
 - VMBus
- Root partition userland attack surface
 - Emulated devices
 - Integration components
- ... and this list is not exhaustive
- MSRC: [first-steps-in-hyper-v-research](#)

State of the Art on the Research

- MSRC and Microsoft publish on Hyper-V
 - Blog posts to help vulnerability researchers
 - e.g. First Steps in Hyper-V research
 - Posts on Hyper-V components
 - Several presentations at conferences on vulnerabilities found internally
 - e.g. Breaking VSM by Attacking SecureKernel at BlackHatUSA 2020
 - Symbols provided for some key components
- Active external contributors
 - @gerhart_x and his dedicated blog
 - @alisaesage
 - Presentation by Damien Aumaitre on whvp at SSTIC 2020
- And many more (a list can be found on [GitHub/gerhart01/Hyper-V-Internals](https://github.com/gerhart01/Hyper-V-Internals))

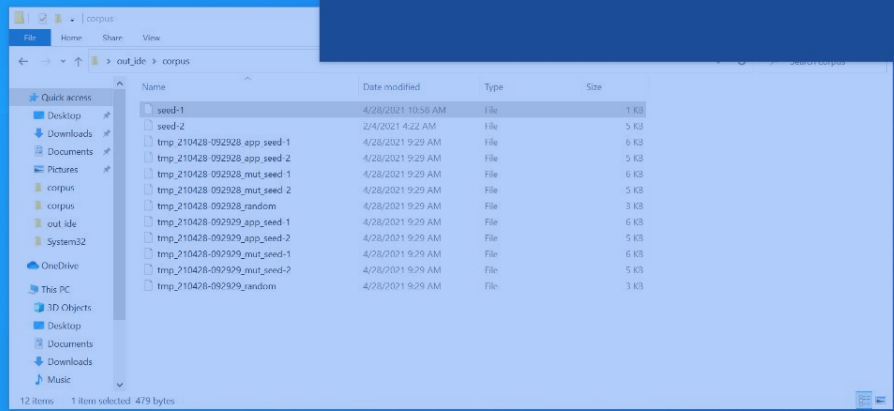
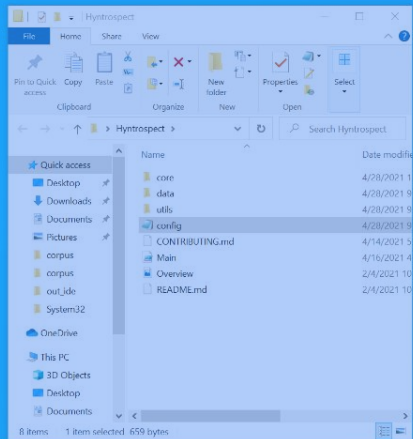
The Research Target

```
Administrator: C:\Users\John\Desktop\Hyntrospect\src\Debug\Hyntrospect.exe
ModLoad: 80007Ffe 35600000 00007Ffe 3567F000 PapiuCLnt
ModLoad: 80007Ffe 13145000 00007Ffe 13145000 symyntos
ModLoad: 80007Ffe 397a0000 00007Ffe 39855000 vmrdrvcore

r11=0000000000000004 r12=0000000000000000 r13=0000000000000000
r14=0000000000000000 r15=0000000000000000
Iop1=0 TBD: flags
cs=0033 ss=002b ds=002b es=002b fs=0053 gs=002b efl=00000246
[00007Ffe43628800] DebugBreakPoint:
00007Ffe43628800 cc int 3

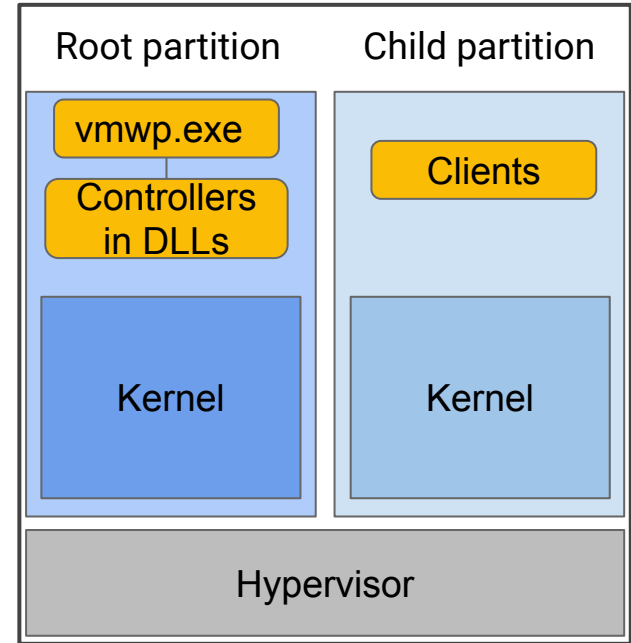
Setting the breakpoints. That operation can take a long time depending on the size of the list (-sec to hour).
Starting the execution.
```

```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\Hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



The Emulated Devices Controllers

- Examples: floppy disks, IDE, PS2
- Called “virtual devices” or “VDEVs” at Microsoft
- Emulation of hardware controller by the hypervisor
 - Real hardware controllers use and access control
 - Resources shared
 - Guest operating systems unmodified
- Implemented for Hyper-V generation 1 VMs
 - Azure mostly uses this generation
- Userland of the root partition
- In DLLs loaded by the worker process

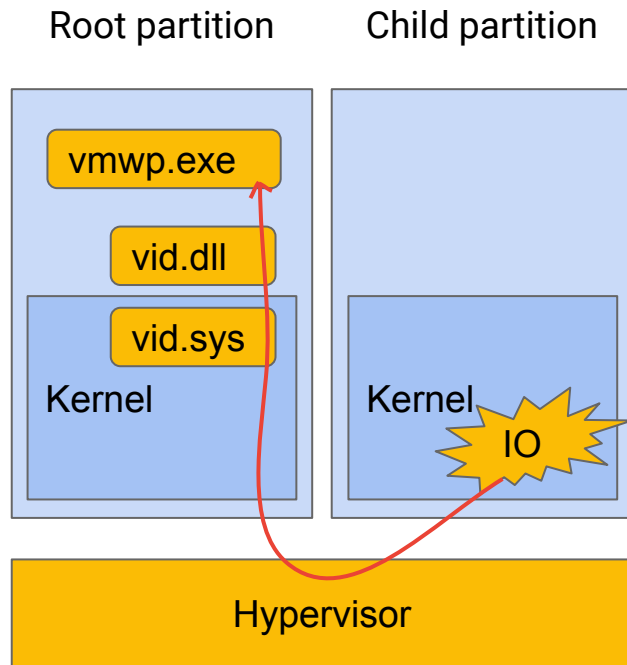


Why Choosing the Emulated Devices?

- Complex (state machines)
 - For example: enabling / disabling ports, updating a status register, waiting for a command
- Several bugs on several hypervisors
- Azure mostly uses Generation 1 VMs
- Hyper-V is developed in C++
- Potential “guest to root partition” escapes

Life of a Request

- Communication through IO ports
 - CPU instructions: IN / OUT
 - “IN EAX, DX”: input from I/O port in DX into EAX
 - “OUT DX, EAX”: output in EAX to I/O port address in DX
 - More details and versions in Intel manuals
- Communication through the hypervisor, the VID, and callbacks
- More on MSRC blogpost “Attacking the VM worker process”
- Some VDEVs are more complex with MMIO handling for instance, or the use of the VMBus



Some Reverse Engineering

- DLL implementing the controllers
- Typical IO handlers
 - \$Device::NotifyIoPortRead
 - \$Device::NotifyIoPortWrite
- + : Symbols available, no particular difficulty (no obfuscation...)
- - : Reversing C++ and its indirect calls
- Example with VmEmulatedStorage.dll



The screenshot shows the IDA Pro interface. On the left, the 'Functions window' lists several functions with their segments and start addresses:

| Function name | Segment | Start |
|--|---------|-------|
| FloppyControllerDevice::NotifyIoPortRead... | .text | 000 |
| FloppyControllerDevice::NotifyIoPortWrite... | .text | 000 |
| IdeControllerDevice::NotifyIoPortRead(u... | .text | 000 |
| IdeControllerDevice::NotifyIoPortWrite(u... | .text | 000 |

The main window displays the Pseudocode-A view for the function `FloppyControllerDevice::NotifyIoPortWrite`. The code is as follows:

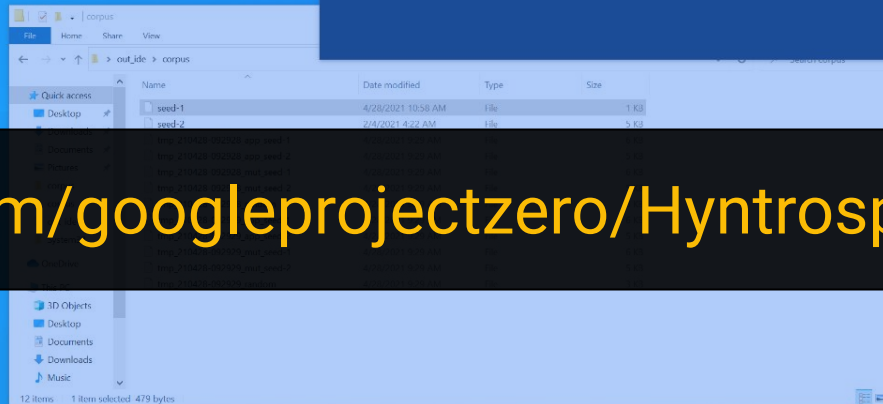

```
; __int64 __fastcall FloppyControllerDevice::NotifyIoPortWrite(FloppyControllerDevice *__hidden this, unsigned __int16, unsigned __int16, unsigned int)
?NotifyIoPortWrite@FloppyControllerDevice@UEAAJGGI@Z proc near

var_38= qword ptr -38h
arg_10= qword ptr 18h

mov     [rsp+arg_10], rbx
push   rbp
push   rsi
push   rdi
```

Hyntrospect: A Fuzzer for the Emulated Devices


```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



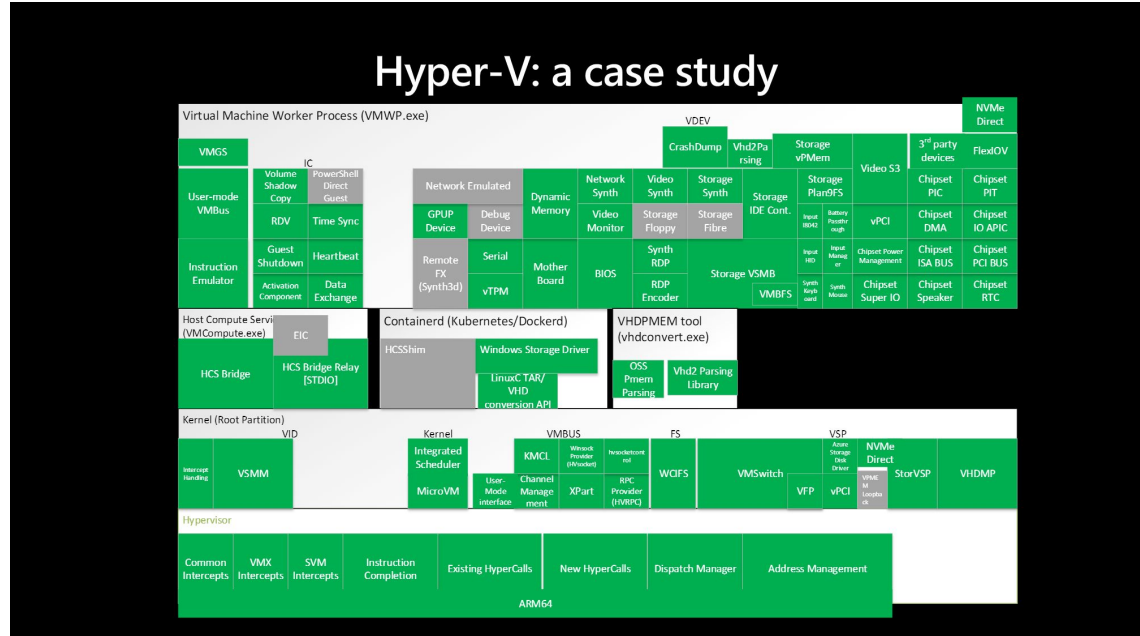
<https://github.com/googleprojectzero/Hyntrospect>

Inspiration

- libFuzzer: coverage-guided approach
- Microsoft publication on their coverage (Keeping Windows Secure - Bluehat IL 2019)
- CVE-2018-0959 + Dedicated MSRC blogpost



How to do the same, closed source?



Existing Tools for Windows Binaries Fuzzing

- **Gathering Coverage**
 - DynamoRIO
 - Intel Pin
 - Intel PT (though this is not a tool like previous two)
 - Mesos
 - QDBI for Windows
 - TinyInst
- **Fuzzers**
 - WinAFL + DynamoRIO
 - Jackalope
 - whvp
- **Memory Corruption Detection**
 - PageHeap

So Why Another Toolchain?

- The target is a DLL
 - This disqualifies all the fuzzers that only apply to executables
- Emulating only the relevant functions is hard
 - All the VM context would be needed
- vmwp binary and the DLL cannot be restarted with instrumentation
 - That would mean restarting the whole VM for each run
- The runtime operations are specific
 - Injecting / mocking the injection of IOs
- Some tools were developed during Hyntrospect development
- Managing all the blocks with a minimal set of languages is hard
- The fuzzer will be ported to similar use cases
 - vSMB, or with some architectural changes the network stack...

Scope

- Windows guest VM
- Intel CPU
- Generation 1 VMs
- Binaries (DLLs/EXEs) in the userland of the root partition

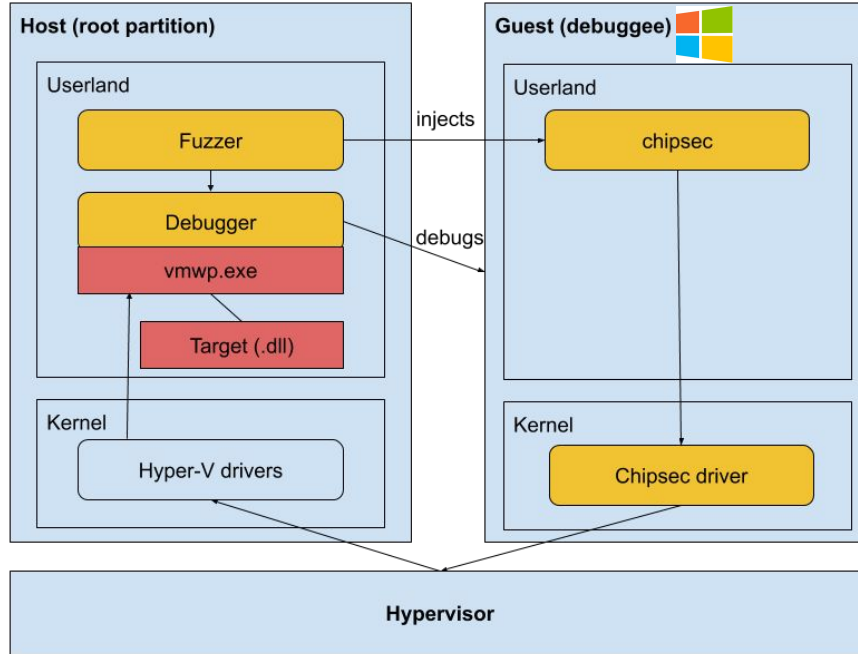
Design Choices at a Glance

| | |
|------------------------------------|---|
| Emulation vs execution | Execution of a VM through a debugger (DbgShell) at runtime |
| Coverage | Tracked with the int3 technique described by @5aelo for TrapFuzz / @gamozolabs mesos |
| Memory corruption detection | Pageheap (gflags) |
| Type of bugs | Memory corruption State machine logic errors [Use after free] Race conditions |

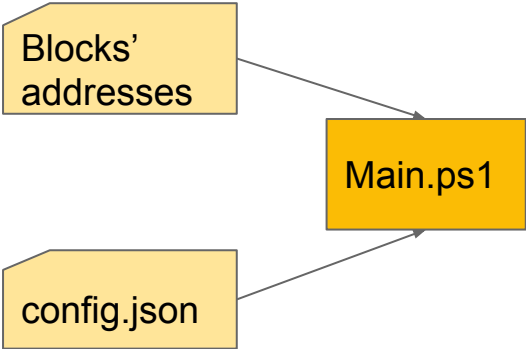
Design Choices at a Glance (2)

| | |
|------------------------------|--|
| Environment reset | Hyper-V checkpoints = snapshots |
| Mutation strategy | Custom |
| Language | PowerShell (except for the IDA scripts) |
| External dependencies | DbgShell, CHIPSEC, [pageheap], [LightHouse], [IDA] |

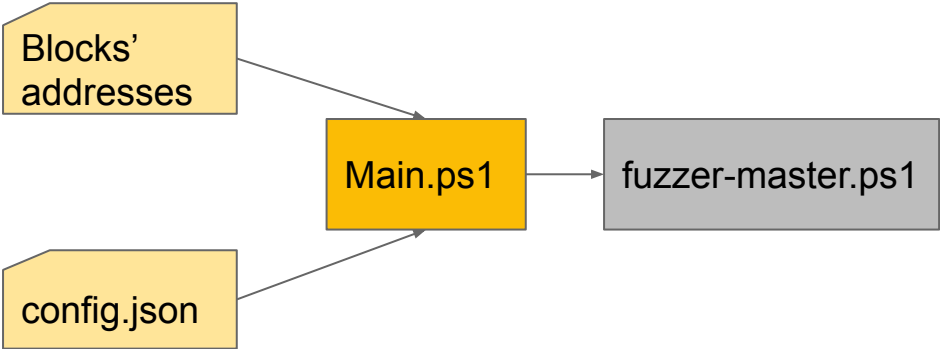
Overview of Hyntrospect



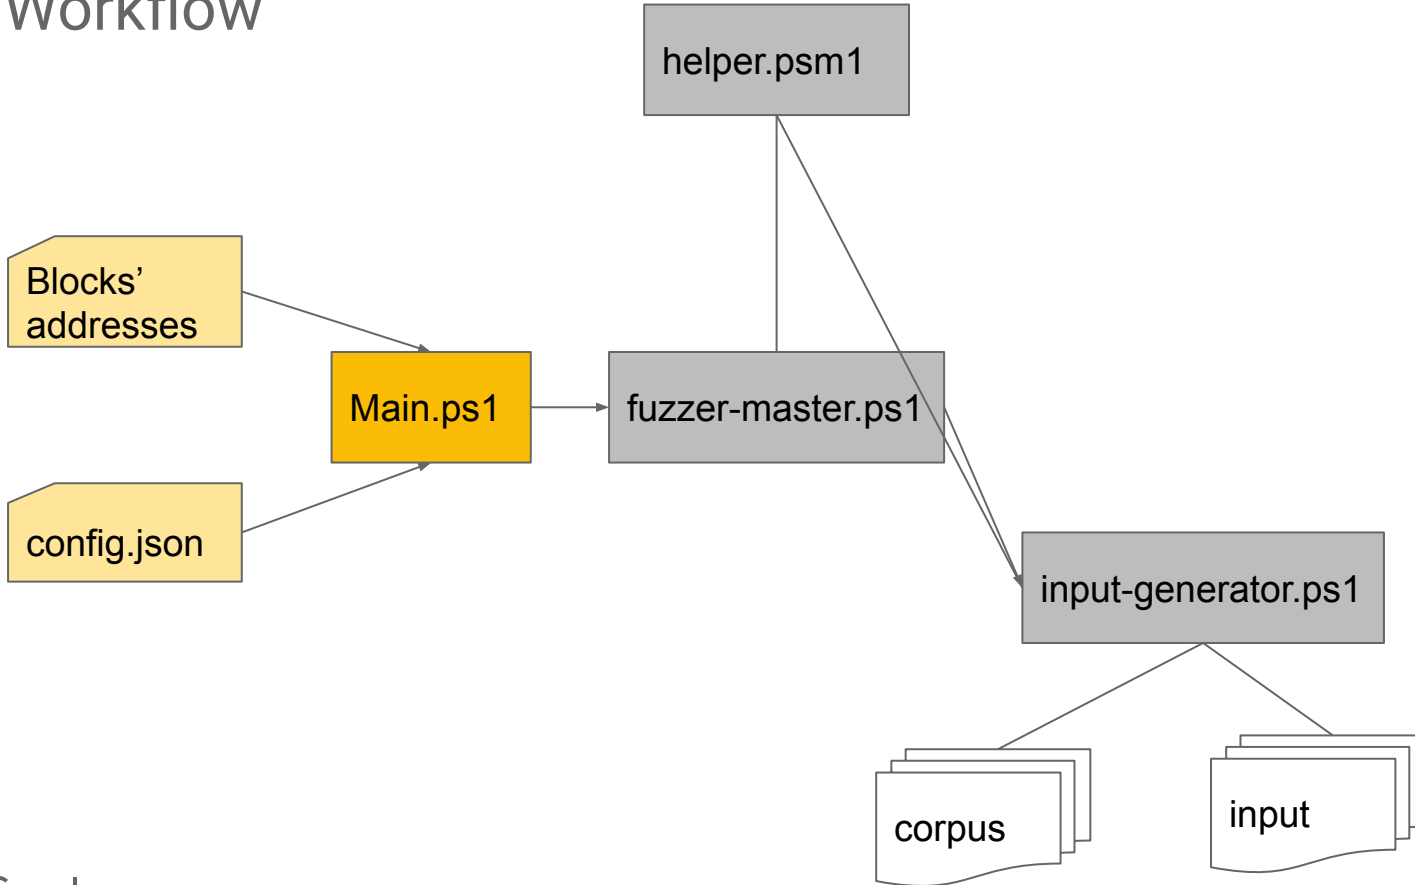
Workflow



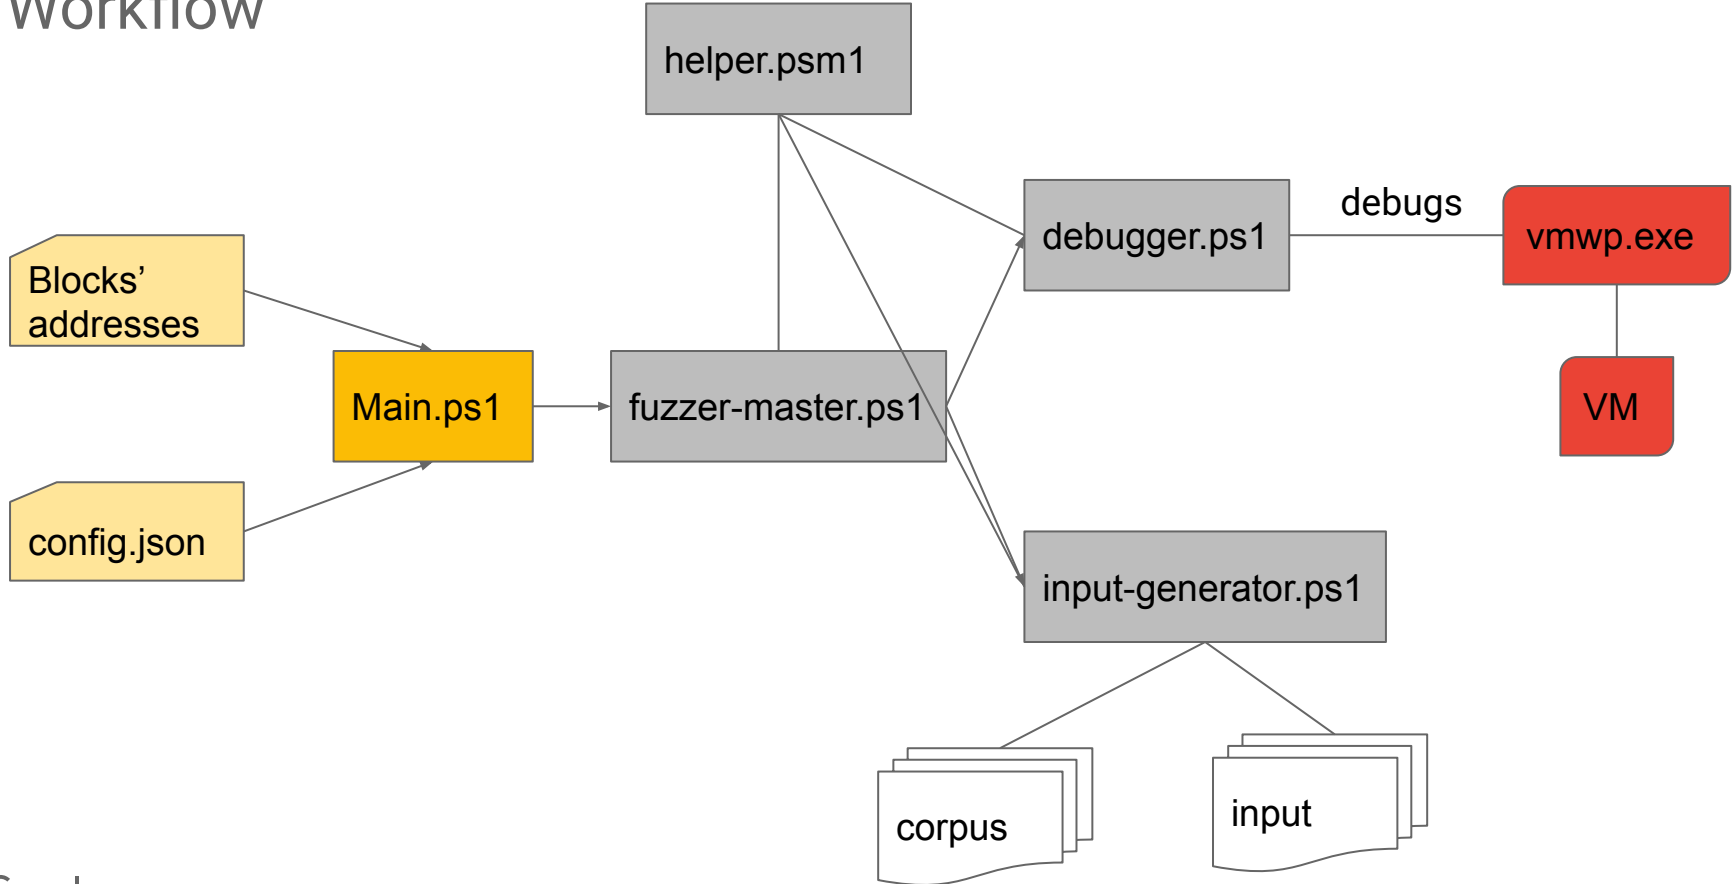
Workflow



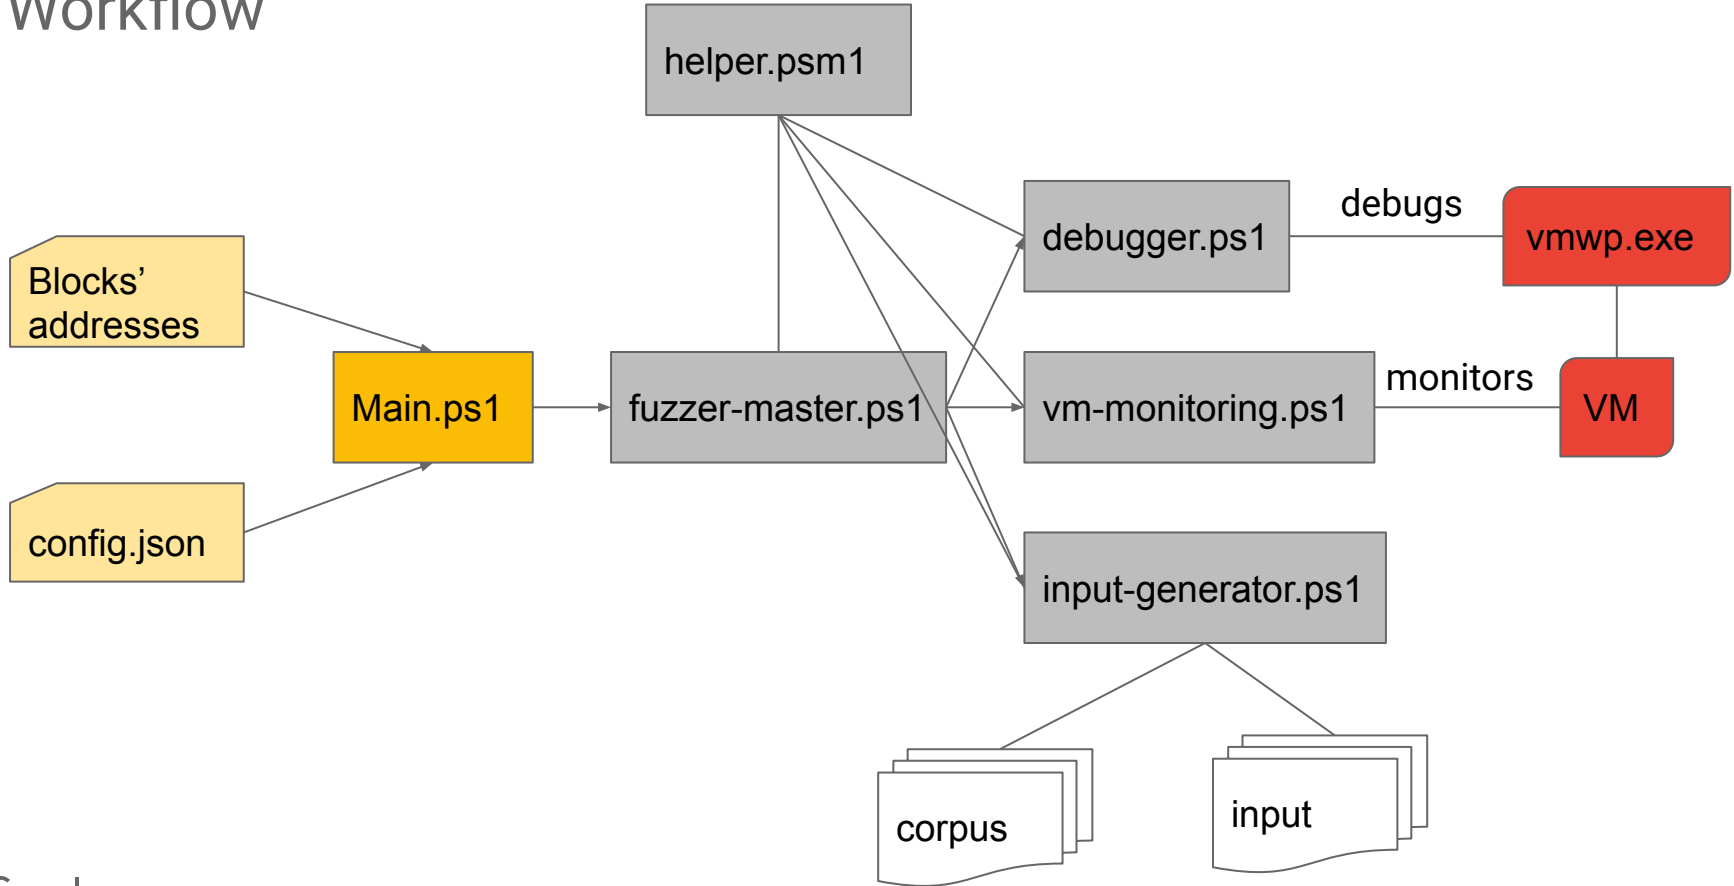
Workflow



Workflow




Workflow




Coverage collection and guidance

- Block coverage




Coverage collection and guidance

- Block coverage




Coverage collection and guidance

- Block coverage




Coverage collection and guidance

- Block coverage
 - Versus edge coverage: easier to implement but does not promote rare paths




Coverage collection and guidance

- Block coverage
 - Versus edge coverage: easier to implement but does not promote rare paths
 - No counter



Coverage collection and guidance

- Block coverage
 - Versus edge coverage: easier to implement but does not promote rare paths
 - No counter
- int3 technique
 - Pre-compute the list of targeted blocks' addresses
 - Set int3 at the beginning of each block
 - Each int3 reached = coverage increase
 - The int3 is removed, input file handled, execution resumes
 - Faster over time

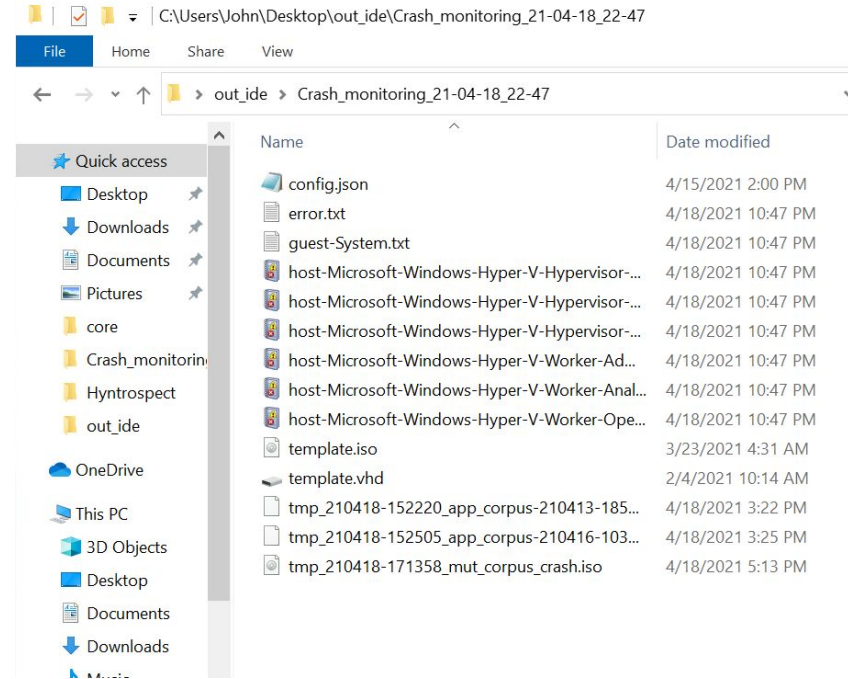


Generation of the Input File

- Record of seeds at the beginning [optional]
 - Record of legitimate traffic
- Corpus of “interesting files”
 - Corpus files = permanent residents
 - Input files = temporary residents to be tested
- Coverage increase -> truncated input file added to the corpus
 - Will influence future runs
- 3 strategies: mutate, append, generate randomly
- Format of input files
 - Byte 0 % 2 -> IN / OUT operation
 - Byte 1 % (number of ports) -> selected IO port
 - Byte 2 % 3 -> length
 - If OUT and based on length -> value

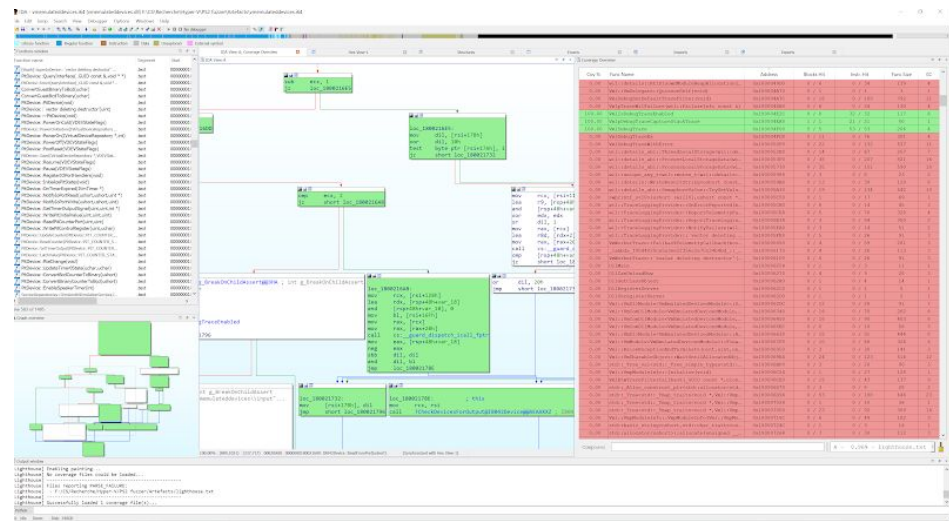
Crash Qualification

- 2 levels of monitoring:
debugger level + monitoring process
 - Tip: the monitoring process can track the VM uptime
 - > avoid while(true)
 - > avoid missing quick status change (up-down-up again) ... as if it was blinking
- Crash folder created with logs and artefacts to re-run the case



Coverage visualization in IDA

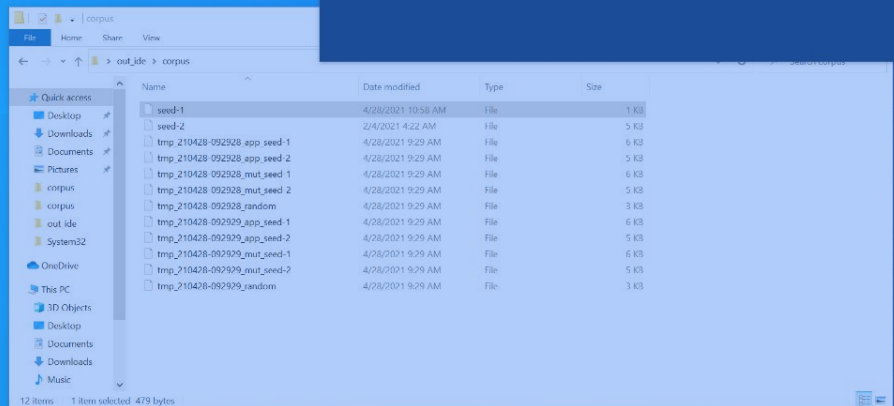

- Optionally, using a helper and IDA+LightHouse



Current Results

```
Administrator: C:\Users\John\Desktop\hyntrospect\corpus\DebugShell.exe
ModLoad: 00007FFE35600000 00007FFE3567F000 PapiuCLnt
ModLoad: 00007FFE131E0000 00007FFE131E5000 symyehntstool
ModLoad: 00007FFE397A0000 00007FFE39855000 vmmdrvcore
ModLoad: 00007FFE41730000 00007FFE417DE000 shcore
...
PopInfo: TBD: flags
cs=0033 ss=002b ds=002b fs=0053 gs=002b efl=00000246
[DebugBreakPoint:
00007FFE43628800 cc int 3
]
Setting the breakpoints. That operation can take a long time depending on the size of the list (-sec to hour).
Starting the execution.
```

```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



Local Runs

- First targets: i8042 (PS/2), videoS3, floppy, IDE
 - Example: I8042 device with IO ports 0x60, 0x61, 0x62, 0x64
- Local setup: dedicated workstation with 32 GB RAM and Intel Core i9 CPU
 - 8 GB per VM, 1 or 2 vCPUs
- Speed limitation
 - Main factor: number of breakpoints
 - Time to set them / update them in DbgShell
 - Not linear
- Next goal: port the fuzzer to GCP

| Number of breakpoints | Time to set up the breakpoints in DbgShell at each iteration |
|-----------------------|--|
| 150 | immediate |
| 500 | 6 seconds |
| 1000 | 20 seconds |
| 2000 | 1 minute 15 seconds |

Coverage (3 days run)

| vmemulateddevices.dll | Current coverage |
|------------------------------|-------------------------|
| VideoS3Device | 42.7% |
| i8042Device | 40% |

| VmEmulatedStorage.dll | Current coverage |
|------------------------------|-------------------------|
| FloppyControllerDevice | 43.3% |
| IdeControllerDevice | 28.8% |

- Start / init / stop functions not called
 - Attaching to a running VM
- Debug strings blocks skipped

Guest VM Crash Found

- On i8042 device
- Reproducible
- BSOD of the VM with different error messages at each run
 - SYSTEM_SERVICE_EXCEPTION (0x3b)
 - PFN_LIST_CORRUPT (4e)
 - ATTEMPTED_WRITE_TO_READONLY_MEMORY (0xbe)
 - KERNEL_SECURITY_CHECK_FAILURE (0x139)
- Memory corruption error

Some More Investigation

- Narrowed down the case
 - Sequence of 2 OUT operations
- State machine, path accessible in 2 steps
- PciBusDevice::HandleA20GateChange
 - Legacy A20 device
 - Updates the memory mapping on the host
 - ... but the guest keeps the same mapping
- Question: possible compromission of VBS?

Follow-up



- In practice, impossible to exploit
 - Not a security bug
 - Shared with MSRC
-
- Validates the behavior of the fuzzer, crash handling and reproduction scripts

Future endeavours

```
Administrator: C:\Users\John\Desktop\hyntrospect.com\DbgShell\64\DbgShell.exe
ModLoad: 00007FFE.35600000 00007FFE.3567F000 PapiCList
ModLoad: 00007FFE.131E0000 00007FFE.131E5000 empyntstool
ModLoad: 00007FFE.397A0000 00007FFE.39855000 vmrdvcore
ModLoad: 00007FFE.41730000 00007FFE.417DE000 shcore

Pop3#0 TBD: flags
cs=0033 ss=002b ds=002b fs=0053 gs=002b efl=00000246
[00007FFE.43628800 cc int 3
00007FFE.43628800 cc
Setting the breakpoints. That operation can take a long time depending on the size of the list (-sec to hour).
Starting the execution.
```

```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



Design Limitations

- Restricted to the userland of the root partition
 - Limits the attack surface as parts of the virtualized stack are in the root partition kernel and hypervisor
- Not optimized for speed
 - More expensive in Cloud as more cycles are needed

Future Work

- **Development of the fuzzer internals**
 - Mutation strategy
 - Userland vs kernel
 - Speed-related updates: minimal debugger?
- **Porting to GCP**
 - Port to new devices
 - Run faster and longer
- **Adapting to other root partition targets**
 - Keeping the frame and “basic blocks”
 - Changing the commands and input consumption



Conclusion

```
Administrator: C:\Users\John\Desktop\hyntrospect\corpus\DbgShell\64\DbgShell.exe
ModLoad: 00007FFE.35600000 00007FFE.3567F000 PapiuCLnt
ModLoad: 00007FFE.131E0000 00007FFE.131E5000 empyntshl
ModLoad: 00007FFE.397A0000 00007FFE.39855000 vmrdvcone
ModLoad: 00007FFE.41730000 00007FFE.417DE000 shcore

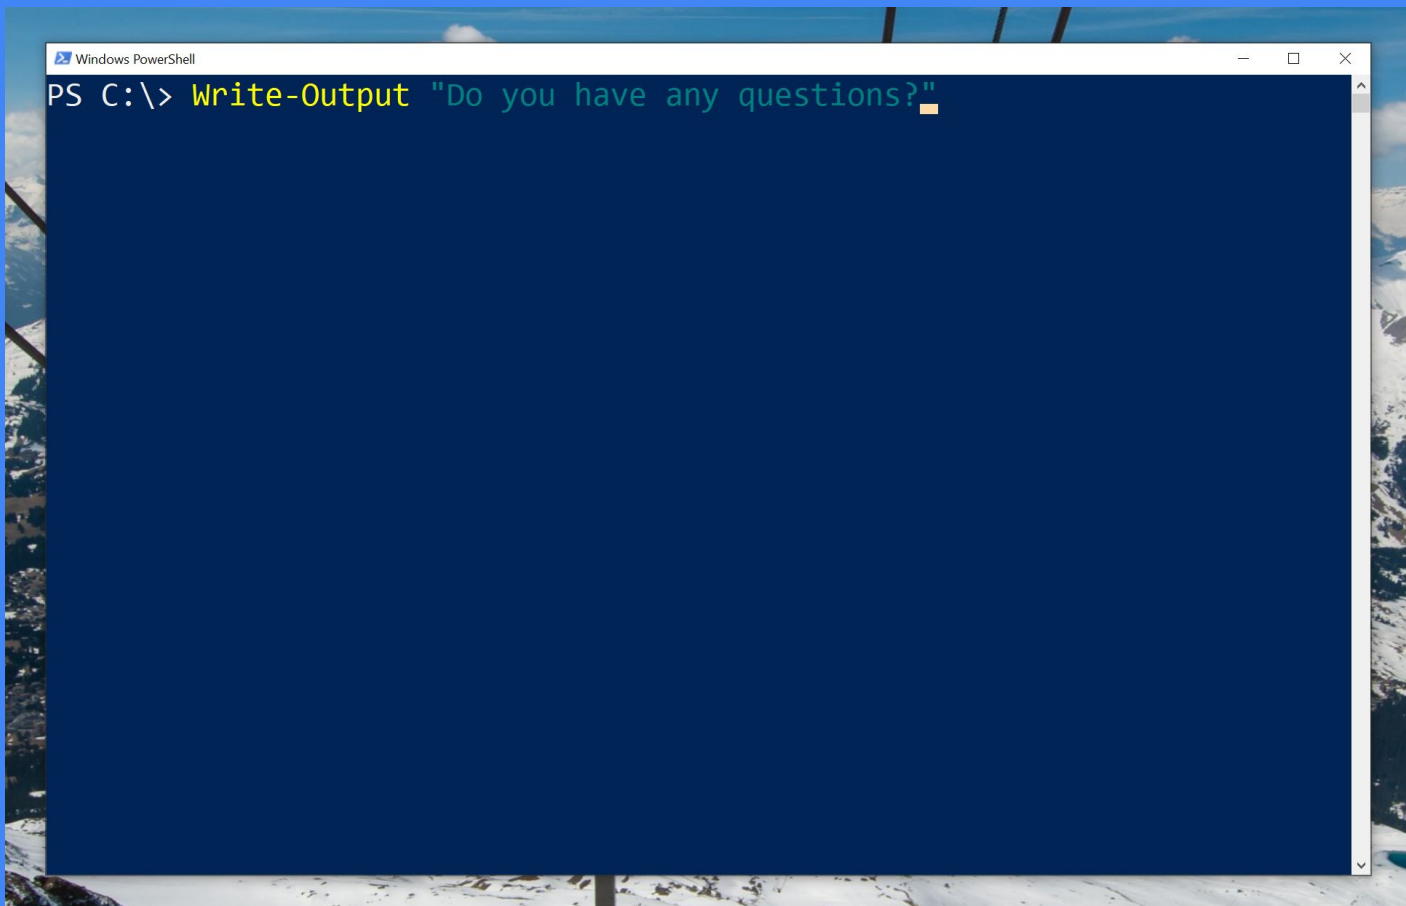
PopInfo TSD: flags
cs=0033 ss=002b ds=002b es=002b fs=0053 gs=002b efl=00000246
[rdi] |DbgBreakPoint:
00007FFE.43628800 cc int 3

Setting the breakpoints. That operation can take a long time depending on the size of the list (-sec to hour).
Starting the execution.
```

```
Administrator: Windows PowerShell
PS C:\Users\John\Desktop\hyntrospect> .\Main.ps1
Starting DbgShell with fuzzer-master parametered through config.json.
C:\Users\John\Desktop\out_ide\corpus\seed-1
```



<https://github.com/googleprojectzero/Hyntrospect>

A screenshot of a Windows PowerShell terminal window. The window title is "Windows PowerShell". The prompt is "PS C:\>". The command entered is "Write-Output 'Do you have any questions?'". The output of the command is "Do you have any questions?". The terminal background is dark blue. The window is set against a background image of a snowy mountain landscape.

```
Windows PowerShell
PS C:\> Write-Output "Do you have any questions?"
Do you have any questions?
```