

# Introducing Blue Pill

Joanna Rutkowska  
**Advanced Malware Labs**  
COSEINC

# About this presentation

- This presentation is based on the research done exclusively for [COSEINC Research \(Advanced Malware Labs\)](#)
- This presentation has been first presented at [SyScan'06 conference](#) in Singapore, on July 21<sup>st</sup>, 2006

# Invisibility by Obscurity

- Current malware is based on a concept...
- e.g. *FU* unlinks EPROCESS from the list of active processes in the system
- e.g. *deepdoor* modifies some function pointers inside NDIS data structures
- ... etc...
- Once you know the *concept* you can write a detector!
- This is boring!

# Imagine a malware...

- ...which does not rely on a concept to remain undetected...
- ...which can not be detected, even though its algorithm (concept) is publicly known!
- ...which can not be detected, even though it's code is publicly known!
  
- Does this reminds you a modern crypto?

# Blue Pill Idea

- Exploit AMD64 SVM extensions to move the operating system into the virtual machine (do it 'on-the-fly')
- Provide thin hypervisor to control the OS
- Hypervisor is responsible for controlling "interesting" events inside guest OS

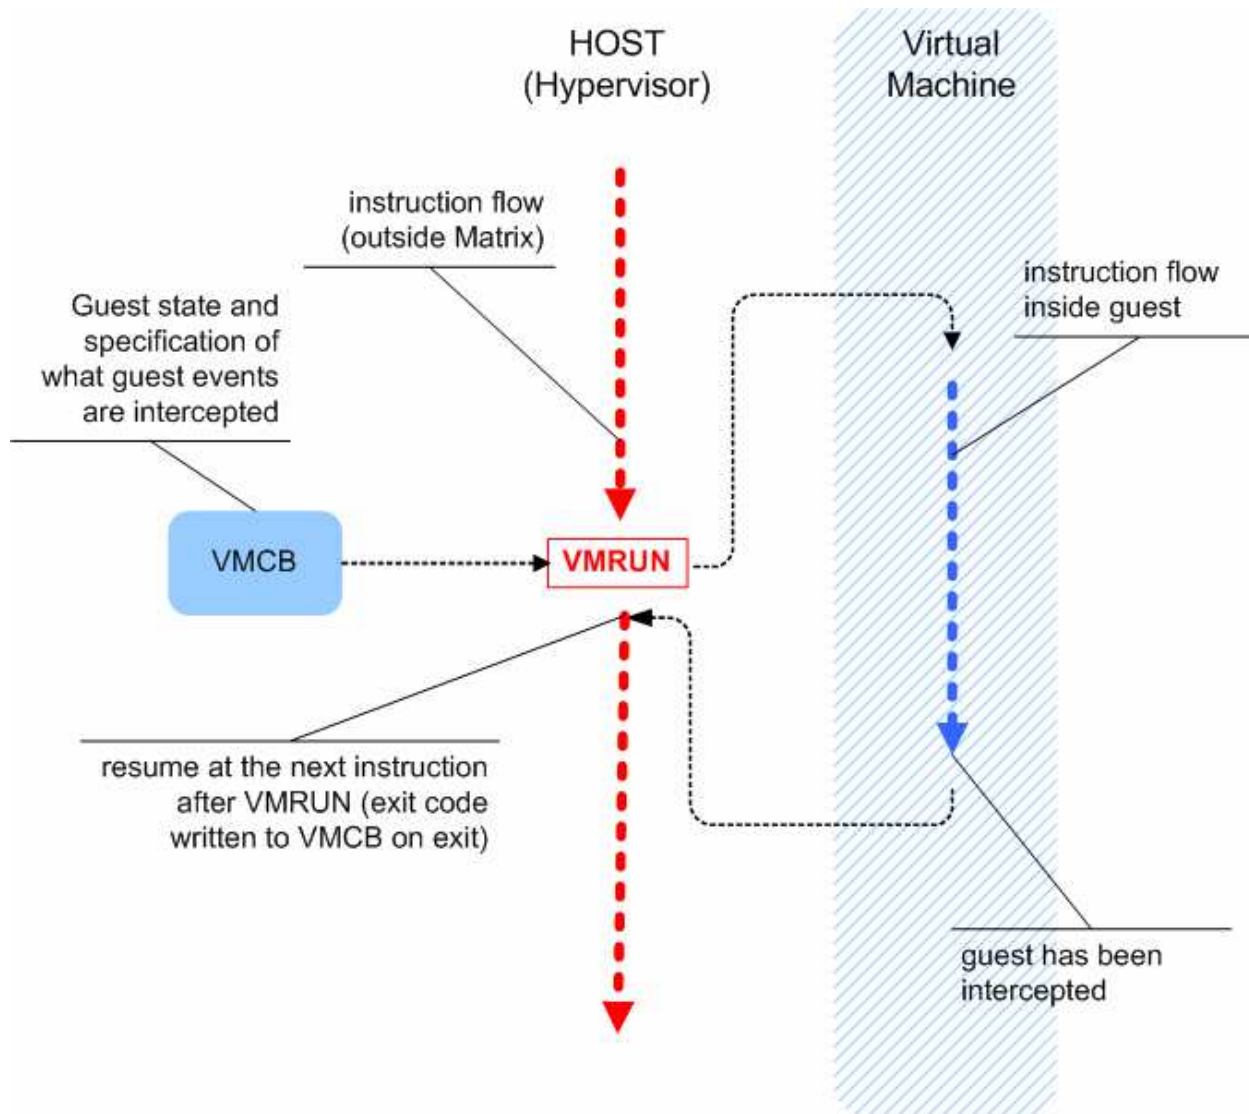
# AMD64 & SVM

- Secure Virtual Machine (AMD SVM) Extensions (AKA Pacifica)
- May 23<sup>rd</sup>, 2006 – AMD releases Athlon 64 processors based on socket AM2 (revision F)
- AM2 based processors are the first to support SVM extensions
- AM2 based hardware is available in shops for end users as of June 2006

# SVM

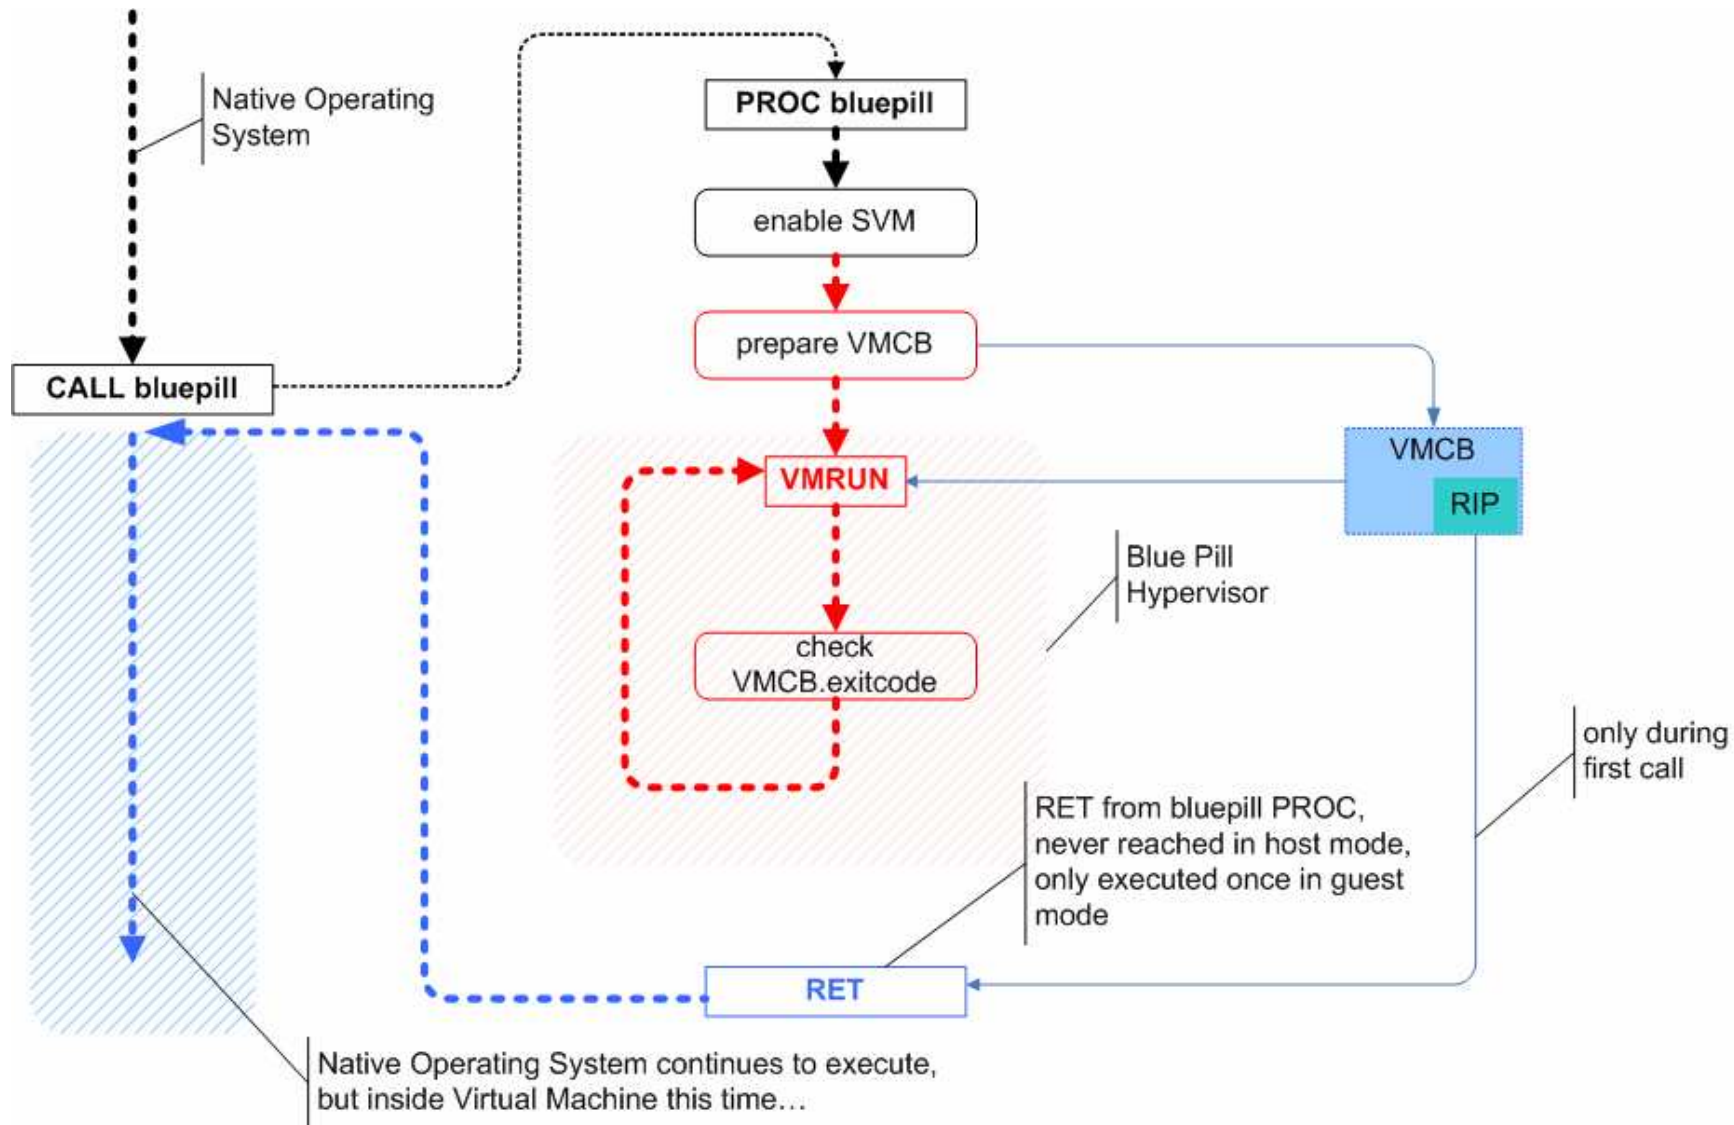
- SVM is a set of instructions which can be used to implement Secure Virtual Machines on AMD64
- MSR EFER register: bit 12 (SVME) controls whether SVM mode is enabled or not
- EFER.SVME must be set to 1 before execution of any SVM instruction.
- Reference:
  - AMD64 Architecture Programmer's Manual Vol. 2: System Programming Rev 3.11
  - [http://www.amd.com/us-en/assets/content\\_type/white\\_papers\\_and\\_tech\\_docs/24593.pdf](http://www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/24593.pdf)

# The heart of SVM: VMRUN instruction





# Blue Pill Idea (simplified)



# BP installs itself ON THE FLY!

- The main idea behind BP is that it installs itself on the fly
- Thus, no modifications to BIOS, boot sector or system files are necessary
- BP, by default, does not survive system reboot
- But this is not a problem:
  - servers are rarely restarted
  - In Vista the 'Power Off' button does not shut down the system – it only puts it into stand by mode!
- And also we can intercept (this has not been yet implemented):
  - restart events (hypervisor survives the reboot)
  - shutdown events (emulated shutdown)

# SubVirt Rootkit

- SubVirt has been created a few months ago by researches at MS Research and University of Michigan
- SubVirt uses commercial VMM (Virtual PC or VMWare) to run the original OS inside a VM

# SubVirt vs. Blue Pill

- SV is permanent! SV has to take control before the original OS during the boot phase. SV can be detected off line.
- SV runs on x86, which does not allow for full virtualization (e.g. SxDT attack)
- SV is based on a commercial VMM, which creates and emulates virtual hardware. This allows for easy detection
- Blue Pill can be installed on the fly – no reboot nor any modifications in BIOS or boot sectors are necessary. BP can not be detected off line.
- BP relies on AMD SVM technology which promises full virtualization
- BP uses ultra thin hypervisor and all the hardware is natively accessible without performance penalty

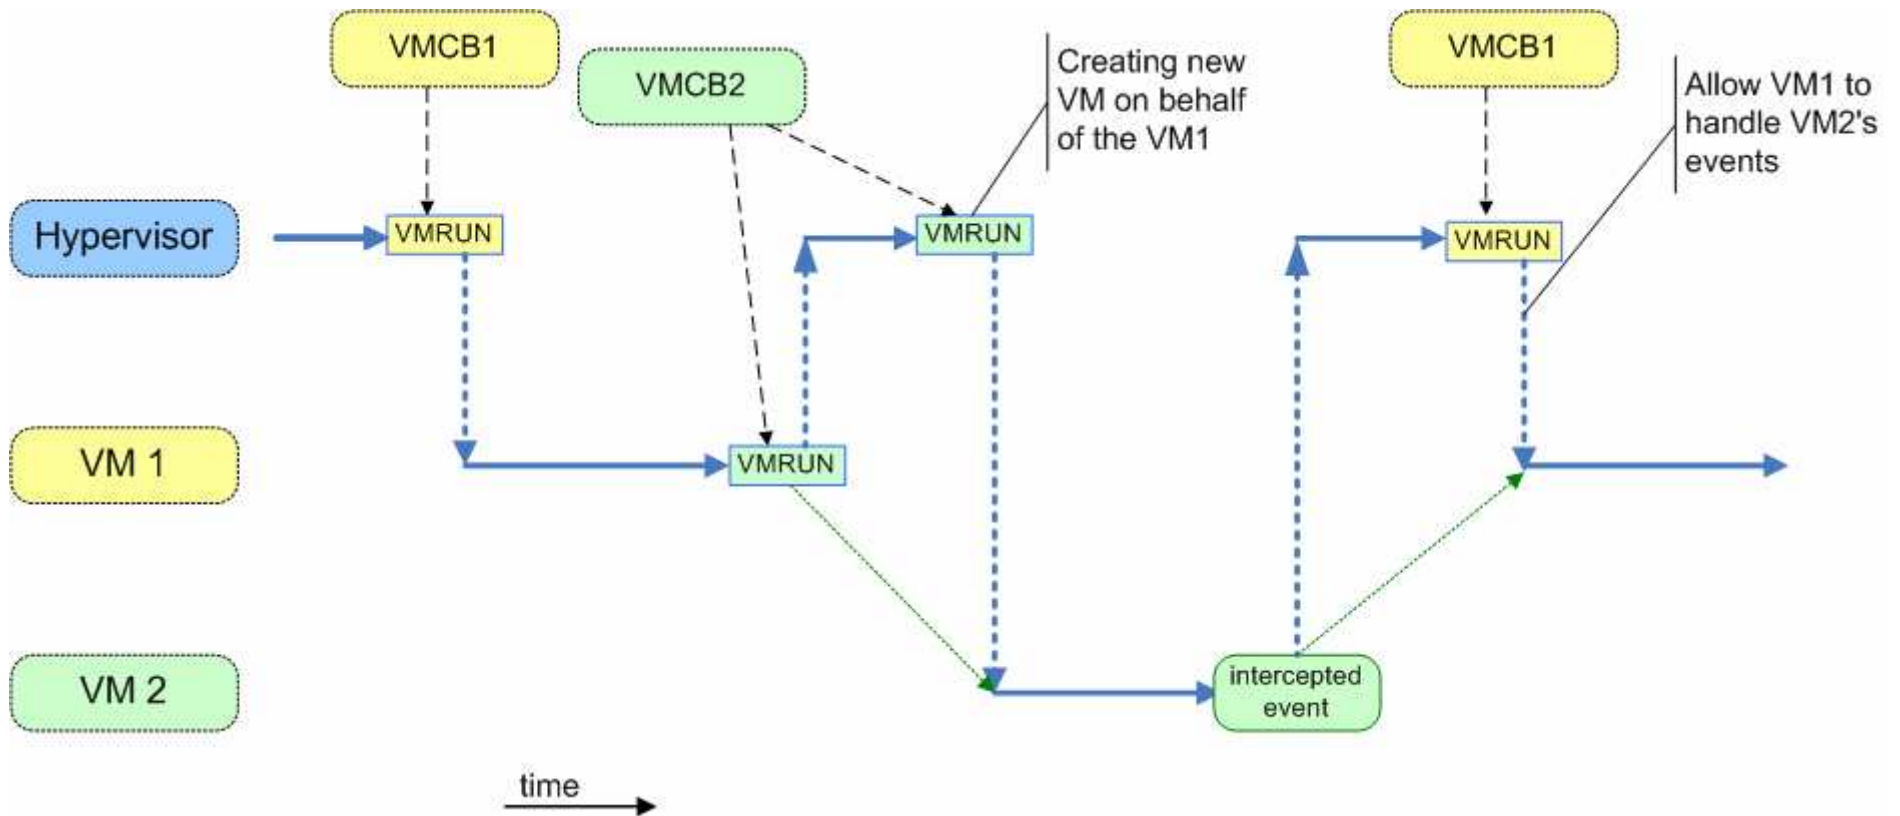
# Blue Pill Demo



# Matrix inside another Matrix

- What happens when you install Blue Pill inside a system which is already bluepilled?
- If nested virtualization is not handled correctly this will allow for trivial detection – all the detector would have to do was to try creating a test VM using a VMRUN instruction
- Of course we can cheat the guest OS that the processor does not support SVM (because we control MSR registers from hypervisor), but this wouldn't cheat more inquisitive users ;)
- So, we need to handle nested VMs...

# Nested VMs



# Detection via timing analysis

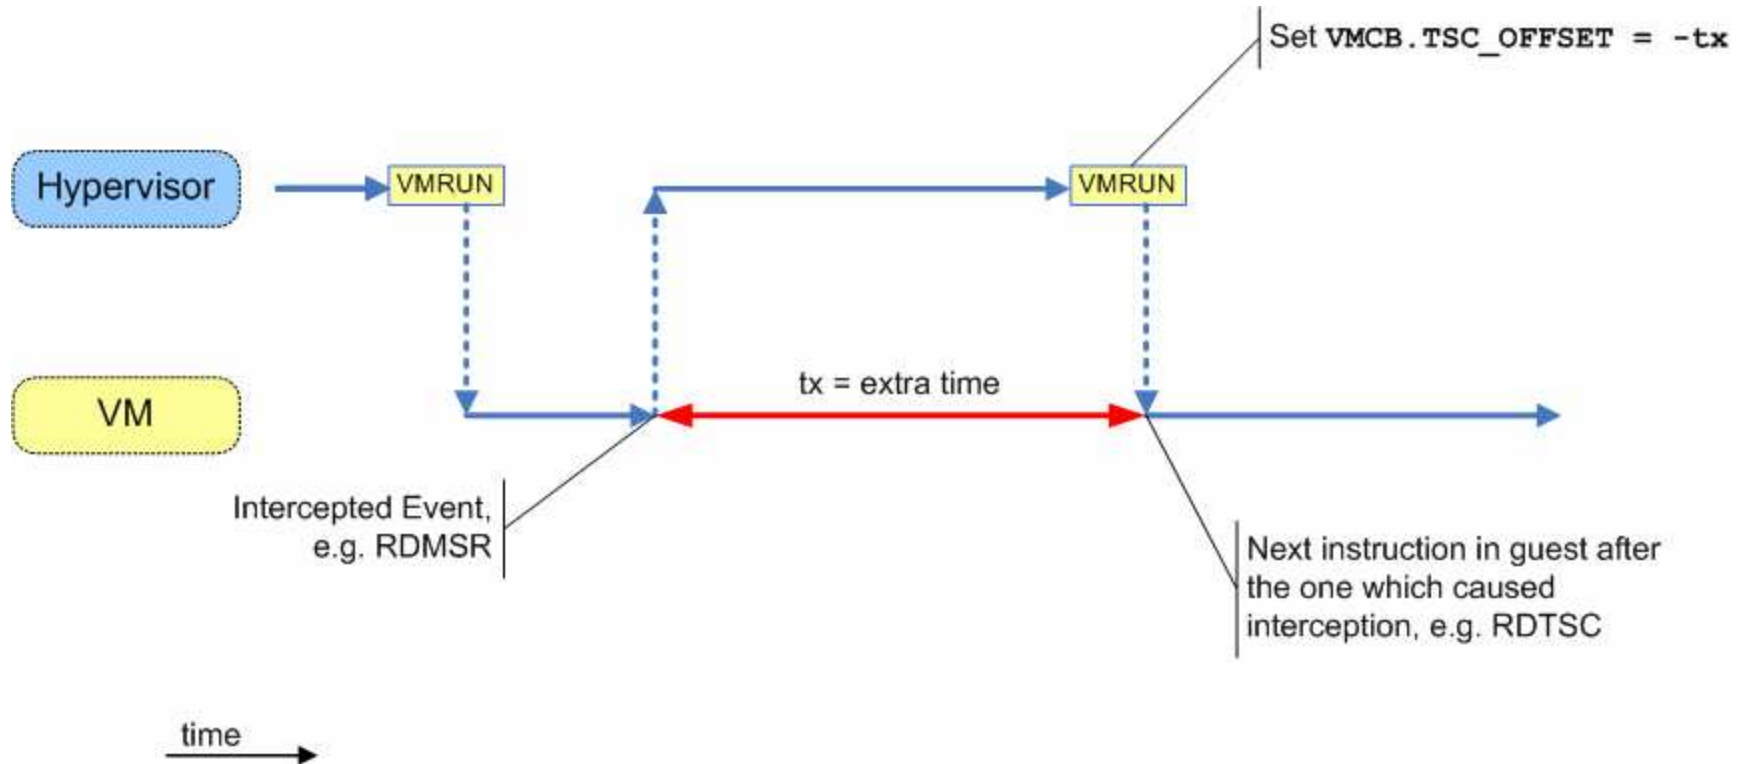
- We can assume that some of the instructions are always intercepted by the hypervisor
  - ~~VMMCALL~~
  - RDMSR – to cheat about the value of `EFER.SVME` bit
- So, not surprisingly, the time needed to execute RDMSR to read the value of EFER would be different (longer) when running from guest
- Detector can execute such instructions a few millions of times and measure the time.



# Cheating timing analysis

- The first problem is that detector (usually) does not have a base line to compare the results with...
- But even if it had – still we can cheat it!
- SVM offers a way to fool the guest's time stamp counter (obtained using RDTSC).
- All we have to do is to adjust VMCB.TSC\_OFFSET accordingly before executing VMRUN (which resumes the guest)

# Time dilatation for guest



# Getting the real time...



# Time profiling in practice

- Now imagine that you need to check 1000 computers in your company using the “external” stopwatch...
- Now imagine that you need to do this a couple of time every day...
- Time dilatation should make it impossible to write a self sufficient detector based on timing analysis!
- The challenge: we need a good ‘calibrating’ mechanism so that we know how much time to subtract.

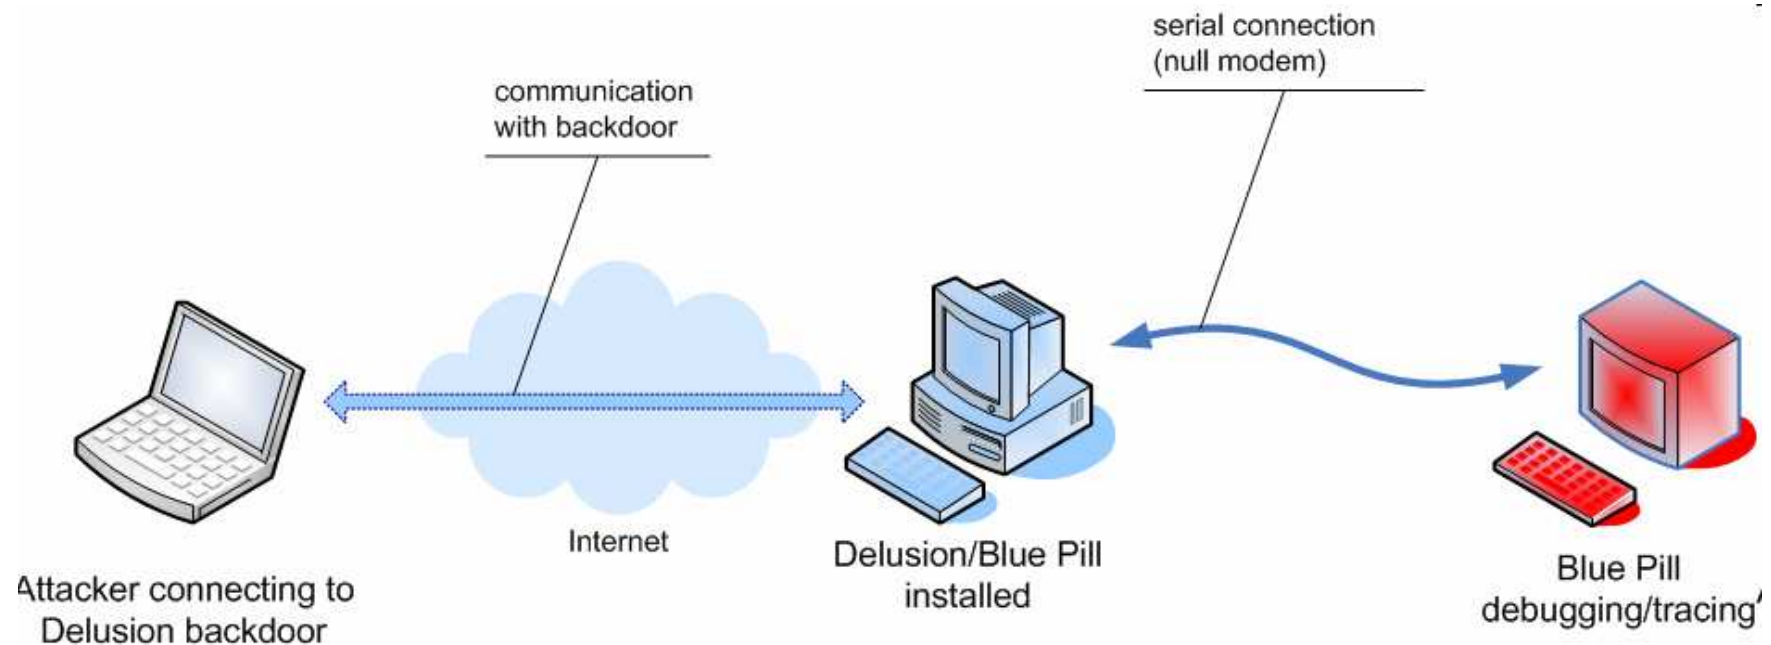
# Blue Pill based malware

- Blue Pill is just a way of silently moving the running OS into Matrix on the fly
- BP technology can be exploited in many various ways in order to create stealth malware
- Basically 'sky is the limit' here :)
- On the next slides we present some simple example:

# Delusion Backdoor

- Simple Blue Pill based network backdoor
- Uses two DB registers to hook:
  - `ReceiveNetBufferListsHandler`
  - `SendNetBufferListsComplete`
- Blue Pill takes care of:
  - handling #DB exception (no need for IDT[1] hooking inside guest)
  - protecting debug registers, so that guest can not realize they are used for hooking
- Not even a single byte is modified in the NDIS data structures nor code!
- Delusion comes with its own TCP/IP stack based on lwIP

# Delusion Demo (Blue Pill powered)



# Blue Pill detection

- Two level of stealth:
  - level 1: can not be detected even though the concept is publicly known (BPL1)
  - level 2: can not be detected even if the code is publicly known (BPL2)
- Level 1 does not requite BP's pages protection
- Level 2 is about avoiding signature based detection
- Level 2 is not needed in targeted attacks
- BPL2 has not been implemented yet!



# Generic BP detection

- If we could come up with a generic program (not based on timing analysis) which would detect SVM virtual mode then...
- it would mean that SVM/Pacifica design/implementation does not support full virtualization!
- To be fair: AMD does not claim full virtualization in SVM documentation – it only says it is ‘Secure VM’... However it’s commonly believed that SVM == full virtualization...

# Blue Pill detection

- We currently research some theoretical generic attacks against BPL1
- It seems that those attacks would only allow for crashing the system if its bluepilled
- It seems that the only attack against BPL2 would be based on timing analysis (or crashing when some special conditions will be met, like e.g. user removing SATA disk in a specific moment during tests)

# Pacifica vs. Vanderpool

- Pacifica (SVM) and Vanderpool (VT-x) are not binary compatible
- However they seem to be very similar
- XEN even implements a common abstraction layer for both technologies
- It *seems* possible to port BP to Intel VT-x

# Blue Pill Prevention

- Disable it in BIOS
  - Its better not to buy SVM capable processor at all! ;)
- Hypervisor built into OS
  - What would be the criteria to allow 3<sup>rd</sup> party VMM (e.g. VMWare or some AV product) to load or not?
  - Or should we stuck with “The Only Justifiable VMM”, provided by our OS vendor? ;)
- Not allowing to move underlying OS *on the fly* into virtual machine
  - How?
  - Besides, would not solve the problem of permanent, “classic” VM based malware
- or maybe another hardware solution...

# Hardware Red Pill?

- How about creating a new instruction – **SVMCHECK** :

```
mov rax, <password>
svmcheck
cmp rax, 0
jnz inside_vm
```

- Password should be different for every processor
- Password is necessary so that it would be impossible to write a *generic* program which would behave differently inside VM and on a native machine.
- Users would get the passwords on certificates when they buy a new processor or computer
- Password would have to be entered to the AV program during its installation.

# Bottom line

- Arbitrary code can be injected into Vista x64 kernel (provided attacker gained administrative rights)
- This could be abused to create Blue Pill based malware on processors supporting virtualization
- BP installs itself on the fly and does not introduce any modifications to BIOS nor hard disk
- BP can be used in many different ways to create the actual malware – Delusion was just one example
- BP should be undetectable in any *practical* way (when fully implemented)
- Blocking BP based attacks on software level will also prevent ISVs from providing their own VMMs and security products based on SVM technology
- Changes in hardware (processor) could allow for easy BP detection

# References

- MS Research and University of Michigan, *SubVirt: Implementing malware with virtual machines* (non-hardware virtualization malware)

# Credits

- Neil Clift for interesting discussions about Windows kernel
- Edgar Barbosa for preparing shellcode for the kernel strike attack
  - Edgar joined COSEINC AML at the end of June!
- Alexander Tereshkin AKA 90210 for thrilling discussions about Blue Pill detection
  - Alex joined COSEINC AML in August!
- Brandon Baker for interesting discussions about Virtualization



**Thank you!**