Security of Cloud Computing

Fabrizio Baiardi
f.baiardi@unipi.it

Syllabus

- Cloud Computing Introduction
  - Definitions
  - Economic Reasons
  - Service Model
  - Deployment Model
- Supporting Technologies
  - Virtualization Technology
  - Scalable Computing = Elasticity
- Security
  - New Threat Model
  - New Attacks
  - Countermeasures ← Introspection

VMI

- Virtual Machine Introspection
- Techniques and tools to monitor VM behavior
- Inspect a VM from the outside to assess what's happening on the inside
- Possible for security tools
  - Virus scanners
  - Intrusion detection systems
- Observe and respond to VM events from a "safe" location outside the monitored machine

Virtualization Overview

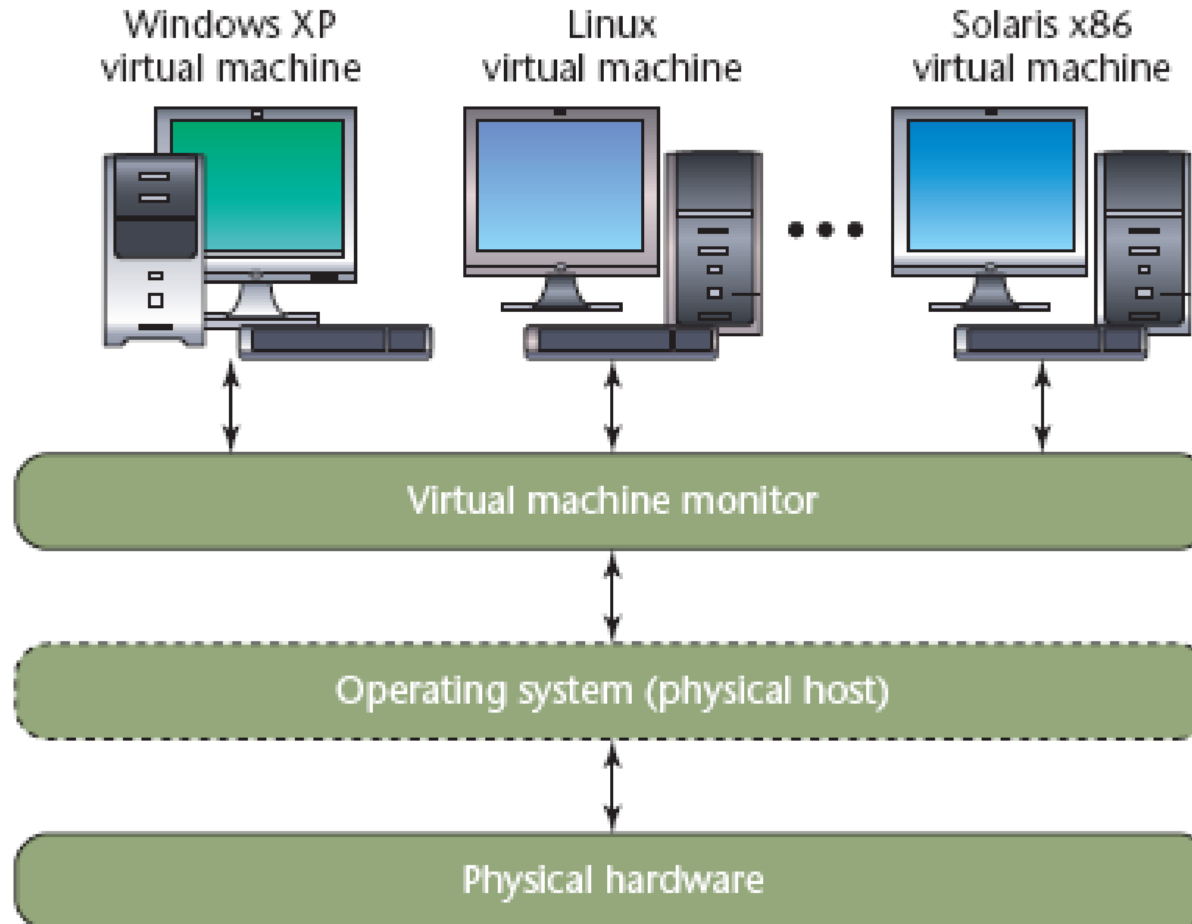

Memory Mapping

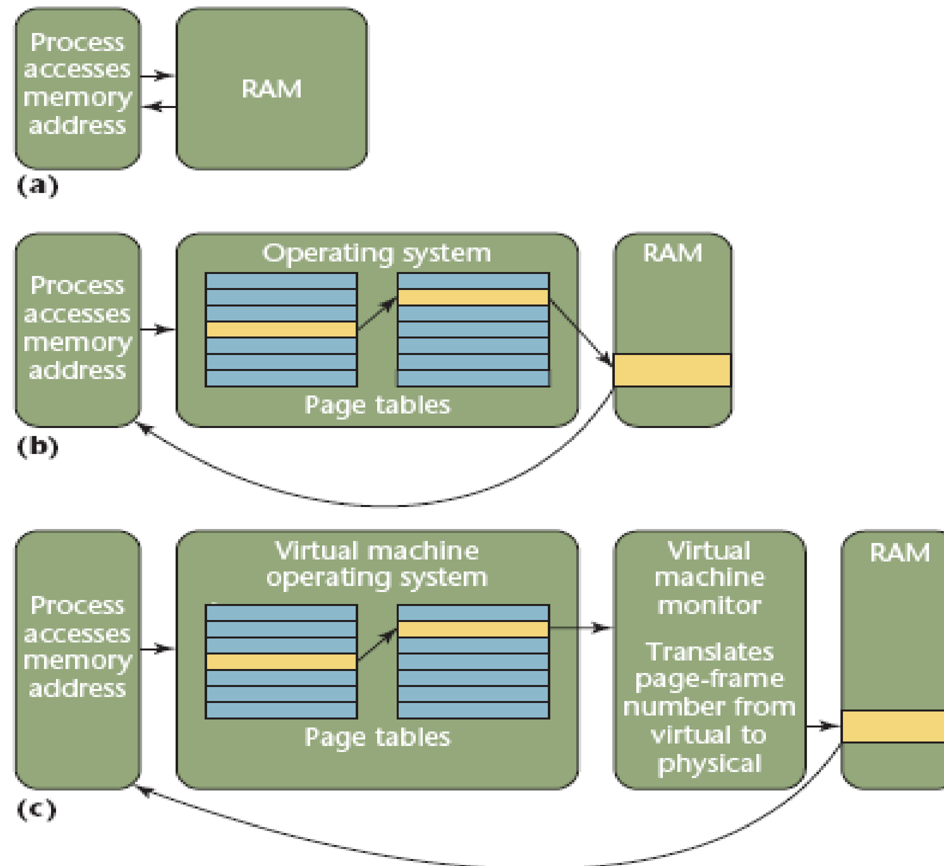

Memory Mapping

- A process perspective
  - Request results in direct access to the memory address
- The OS layer has an active role in providing memory location access
  - Access the page table to map the logical memory address to a physical memory address
- VMM provides an abstraction layer between
  - Each VM OS's memory management
  - The underlying physical hardware
- VMM translates the VM-requested page frame number into a page frame number for the physical hardware
- Gives the VM access to that page

VMM Memory Accesses

- VMM accesses memory pages assigned to each VM directly by
  - VMM's active involvement in this process
  - Its elevated privileges
- Without the VM actually requesting the page
- Can also make those pages accessible to other VMs

Virtual Machine Introspection -1

- By implementing a physical machine through a virtual one, we can check the integrity of any component of the physical machine by evaluating a predicate on the state of the virtual one = on some memory subset of the physical one
- This task can be delegated to the VMM but this strongly increases the complexity of the VMM itself together with the probability of a successful attack
- If the VMM has not been successfully attacked, then the same task can be delegated to another VM, the introspection one
- This may be seen as a particular kind of dynamic, or semantic, attestation where the Introspection VM can give some assurance about the status of another VM
- Bootstrap = the Introspection VM assures the integrity of a component on a VM that, in turn, assures the integrity of the VM

Virtual Machine Introspection -2

- There are several ways of implementing VMI
  - Asynchronous: the introspection VM evaluates some invariant that should hold independently of the actions executed by the VM
  - Synchronous: the introspection VM monitors the execution of the other VM and, at some predefined moments,
    - freezes the execution of the VM
    - evaluates a condition on the status of the VM
    - resume the execution or kills the VM
- Synchronous is more complex because it involves a synchronization between the two VMs
- In any case a semantic gap arises: the Introspection VM access single memory positions while the condition/assertion is defined at a higher abstraction level

Virtual Machine Introspection - 3

- Controls are more expensive but even more robust, wrt those implemented between two processes sharing some memory, because of the separation that the VMM implements between
  - The environment to be monitored, monitored VM
  - The monitoring environment, introspection VM
- To minimize the control cost, a chain of trust can be used where
  - some components in the execution VM implement some control
  - the introspection VM checks the integrity of these components
- In any case, the controls requires the formalization of a process self to be compared against the actual process behaviour

VM Introspection: the modular solution

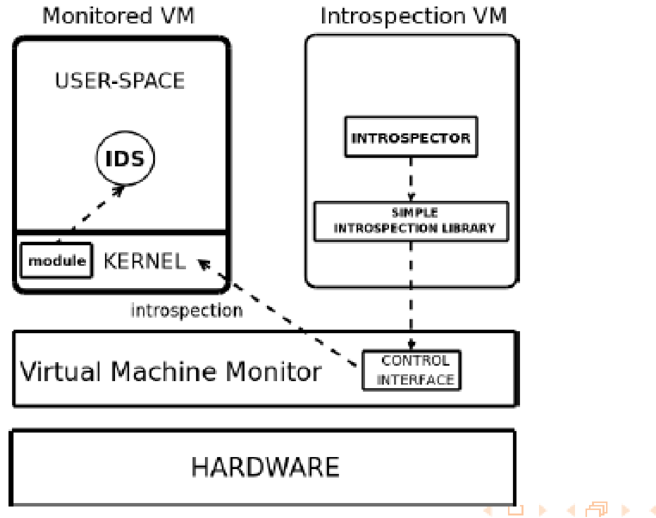

- A simple introspection library to access the memory of the Monitored VM
- A module in the kernel that checks the integrity of the IDS on the Monitored VM
- The integrity of the kernel of the Monitored VM is protected by the Introspector in the Introspection VM
- Definition of the Introspector depends upon that of the module in the kernel
- Checks can be implemented anytime a given nmber of kernel invocation has occurred

Chain of Trust

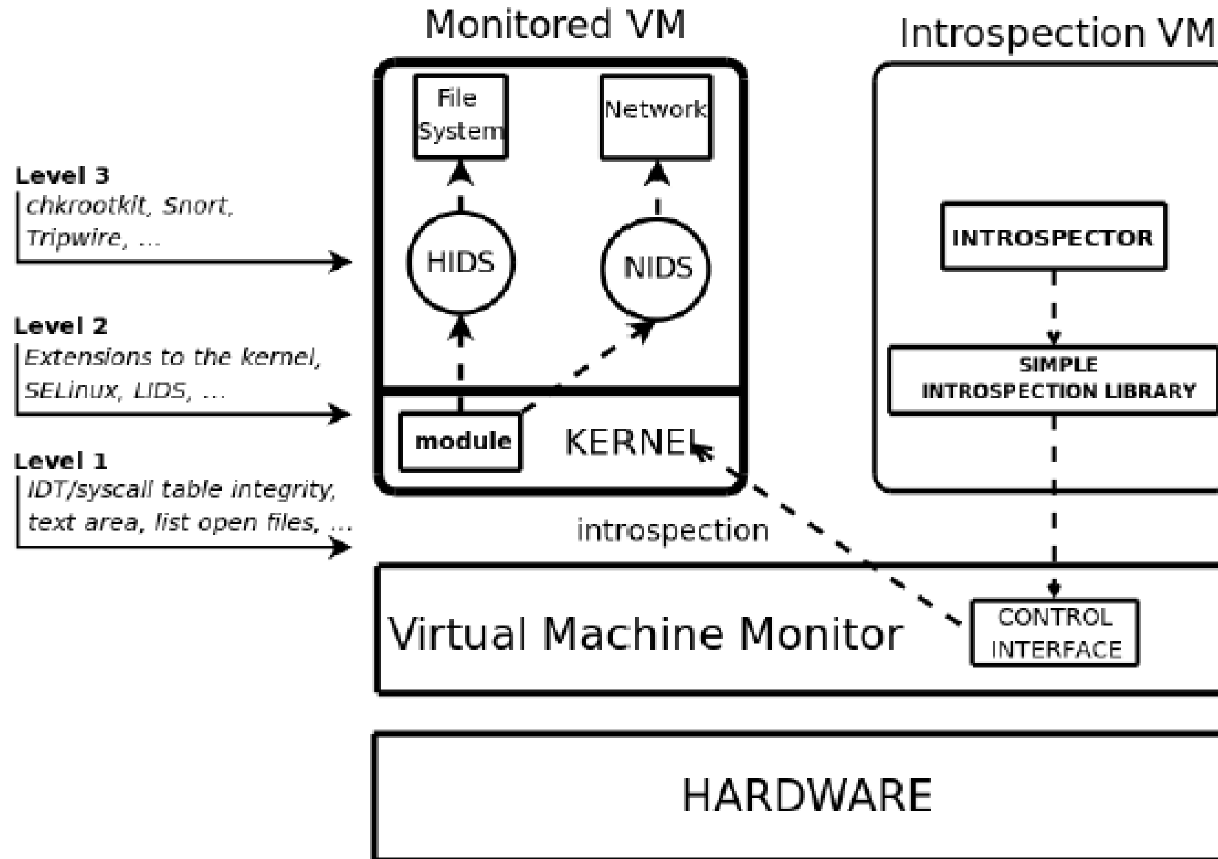

Further advantages of VMI

- Full visibility of the system running inside the Monitored VM: the Introspection VM can access every Monitored VM component, such as the main memory or the processor's registers.
- Transparency: the security checks can be implemented without modifying the software on the Mon-VM and they are almost invisible
  - The kernel has to be modified but not the application running on the Monitored VM
  - If the underlying architecture fully support virtualization, no software on the Monitored VM has to be updated

A full HIDS: Introspection and Alerts

Introspection VM: monitors all the VMs.

- The introspector protects kernel integrity.
- The director:
  1. collects the alerts;
  2. executes actions in response to an alert: stops a VM.

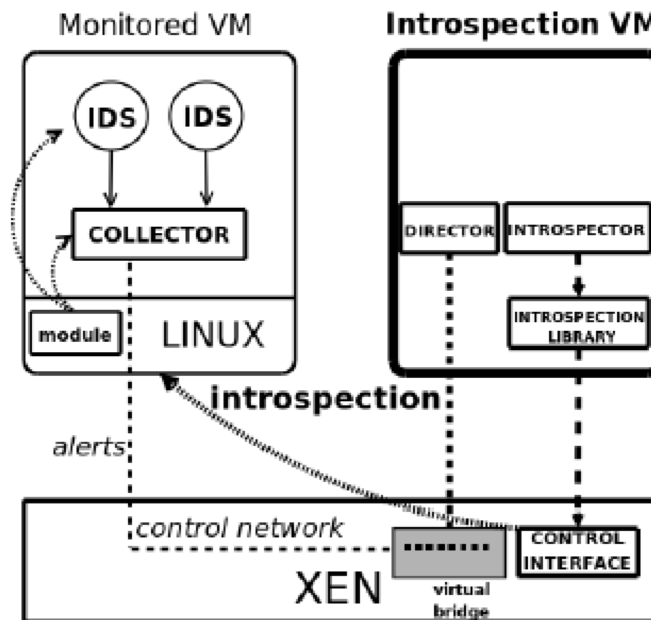

A more general case

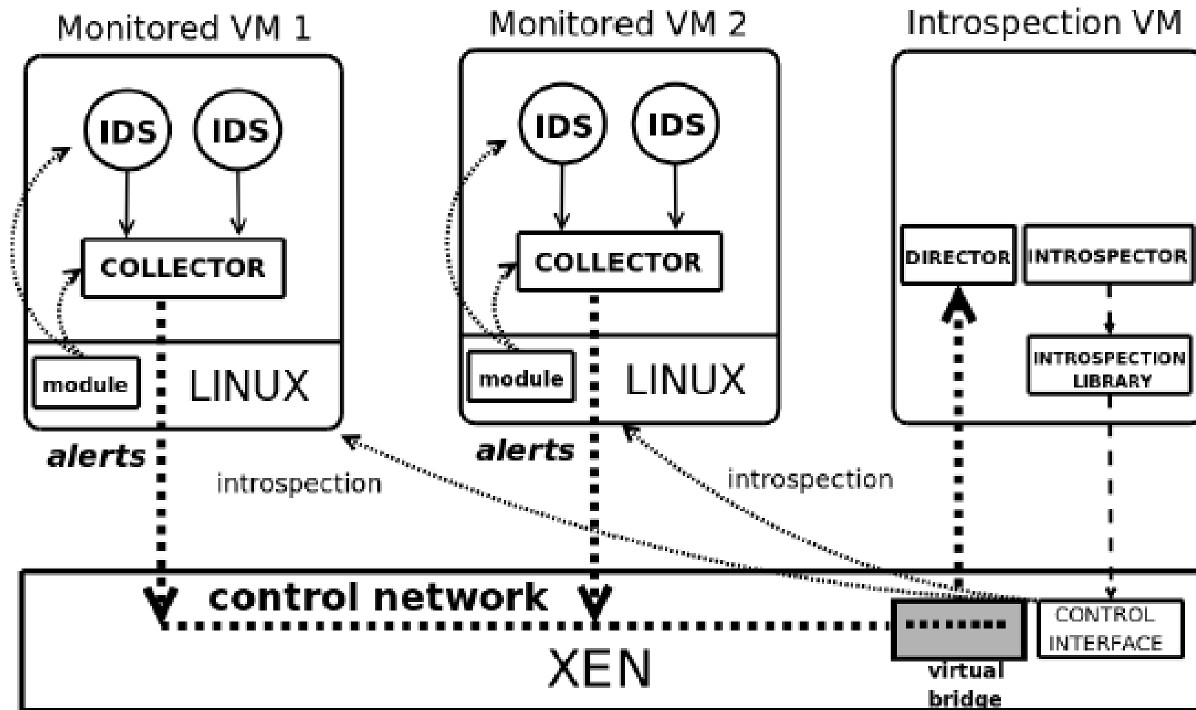

Semantic Integrity and Introspection

A trivial attack classification

- Attacks against user-level processes:
  - the attacker injects some code into a process
  - the attacker diverges the original control-flow to execute the injected code.
- Attacks against the kernel:
  - modify some kernel functionalities
  - modify the kernel behavior to hide any sign of the attack.
- User level attacks are the first step of a complex attack that, after increasing the privilege of the attacker results in the execution of an attack against the kernel or against the kernel and then the VMM

Process Self

- Process Self = The properties of a process that determine its run-time behavior
- The process self can be approximated through static analysis.
- Axiom

  if the process current behavior deviates from the process self then the process code has been altered by an attack.
- Measuring the semantic integrity:
  - the approximation of the process self
  - the monitoring the actual process behavior to assure that it is coherent with the process self.
- If P is a generic process that we want to protect.
  - Self (P) refers to the process self of P
  - SourceCode(P) is the source code of P program= syntactic integrity

All the relations

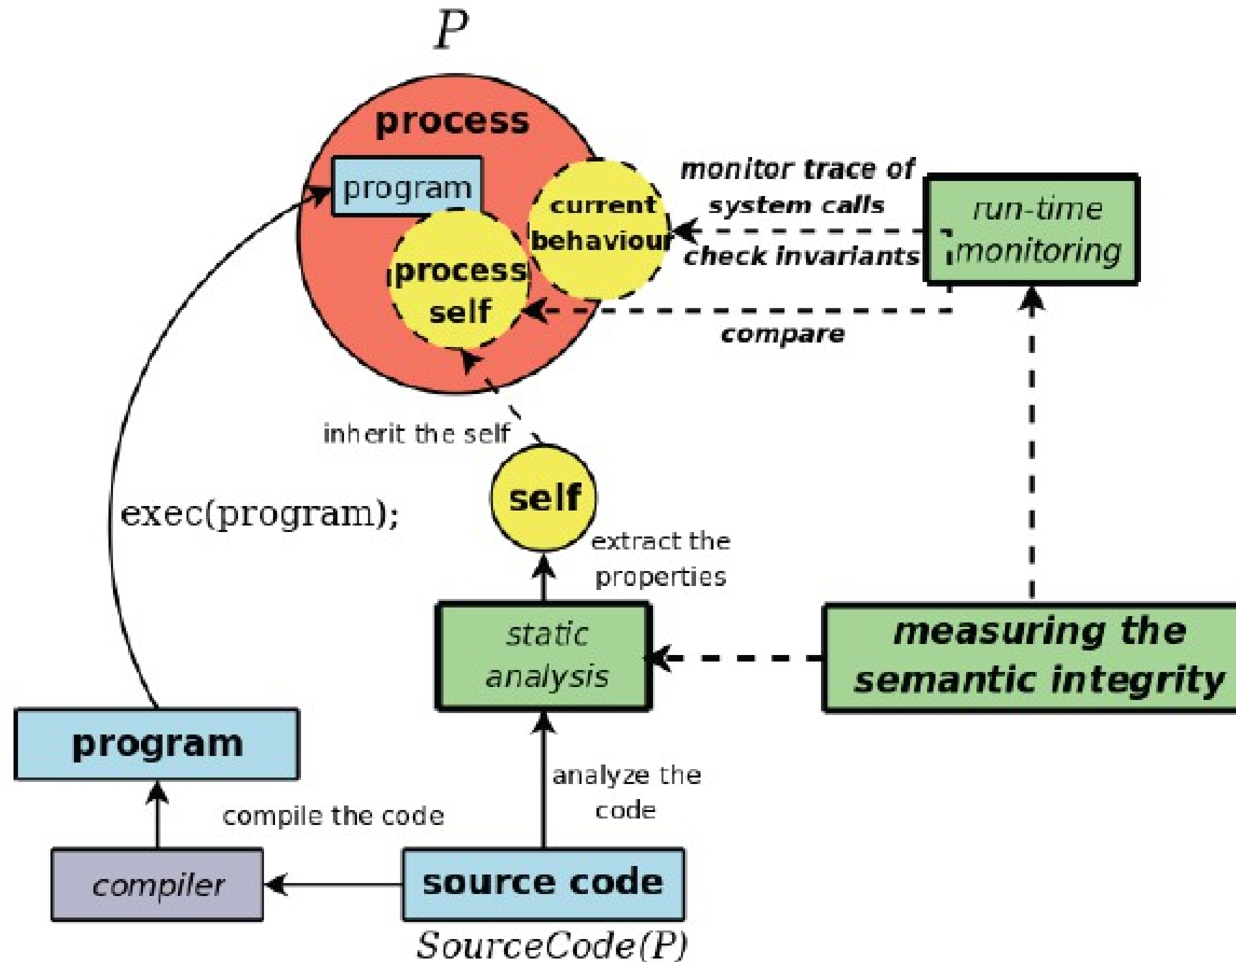

Self and OS calls

- It is widely accepted that an abstract description of a process self should consider the OS calls issued by the process
- Any attempt to violate the security policy, hide the trace of an attack, avoid intrusion detection mechanisms involves some interaction with the OS
- Hence the process self should be defined in terms of the OS calls

Self: Alternative Descriptions

Default Allow

- Forbidden Calls: the set of system calls that P cannot issue
- Forbidden Parameters: the set of system calls that P cannot issue and assertions on the parameters it cannot transmit to a call

Default Deny

- Hashing or Memory Invariants; memory invariants to be evaluated anytime P issues a given system call
- Allowed Calls: the set of system calls that P can issue and assertions on their parameters
- Enriched Traces: the sequence of system calls that P issues in one execution; each call may be coupled with an assertion on the process memory

Enriched Traces

- A set of enriched traces fully describes alternative legal behaviors of P
- Proper static tools may be designed to map SourceCode(P) into Self (P) described through enriched traces = <CFG(P), IT(P) >
- CFG(P) =
  - context-free grammar that defines the system call traces that P may issue during its execution
  - a set of strings on an alphabet with a symbol for each system call
- IT(P)= a set of invariants {I(P, 1), …, I(P, n) }, each associated with a program point i, $1 \leq i \leq n$, where P invokes a system call.

Grammar Generation Algorithm - 1

- A static tool can generate CFG(P) while traversing AST(P), the abstrax syntax tree of P
- CFG(P) = < T, F, S, R > where
  - T is a set of terminal symbols with one symbol for each distinct system call in SourceCode(P)
  - F is a set of non-terminal symbols, one for each function defined in SourceCode(P); each symbol corresponds to a subset of T.
  - S is the starting symbol, which corresponds to main;
  - R is the set of production rules $X \rightarrow B$ where
    - X is a non-terminal symbol
    - B a sequence of terminal and non-terminal symbols.

Grammar Generation Algorithm - 2

- GGA analyzes AST(P) and for each function *fun* defined inSourceCode(P) it inserts into F a new non-terminal symbol S_{fun} and a new rule Rnew into R with S_{fun} as its left-hand-side
- To generate the right-hand side of the rule, GGA linearly scans the definition of fun in SourceCode(P)
- Distinct production rules may be generated, according to the type of statements in the body of *fun*.
- For each statement, GGA generates a new rule and adds a new symbol to the right-hand side of Rnew .
- In this way, CFG(P) represents the system calls that *fun* can invoke and the ordering among the invocations in the body of *fun*.

Grammar Generation Algorithm - 3

```
1   f(){
2       open();
3       read();
4       g();
5       close();
6   }
7
8   g(){
9      getpid();
10  }
```

$\langle F \rangle \longrightarrow$ **open read** $\langle G \rangle$ **close**;
$\langle G \rangle \longrightarrow$ **getpid**;

```
1   f(){
2       open();
3       if(x)
4           read();
5   }
```

$\langle F \rangle \longrightarrow$ **open** $\langle ST_1 \rangle$;
$\langle ST_1 \rangle \longrightarrow$ **read** $\mid \epsilon$;

May result in a false negative

```
1   f(){
2       open();
3       if(x)
4           read();
5       else
6           close();
7   }
```

$\langle F \rangle \longrightarrow$ **open** $\langle IFEL_1 \rangle$;
$\langle IFEL_1 \rangle \longrightarrow \langle STIF_2 \rangle \mid$
$\langle ELSE_3 \rangle$;
$\langle STIF_2 \rangle \longrightarrow$ **read**;
$\langle ELSE_3 \rangle \longrightarrow$ **close**;

May result in a false negative

Assertion Generator -1

- The Assertion Generator traverses AST(P) and analyzes the variables, functions and language statements to build the invariant table (IT (P)).
- To simplify the analysis, we restrict to:
  - integer variables: only files and socket descriptors to express relations among these variables and the system calls;
  - string variables: in case of arrays of char statically declared, functions to manipulate strings are treated like assignments;
  - struct members: only integer or string type field.

Assertion Generator - 2

Any assertion is the composition of any of the followings:

- Parameters assertions. They express data-ow relations among parameters of distinct calls, e.g. the file descriptor in a read call is the result of a previous open call.
- File Assertions. To prevent symlink and race condition attacks, they check, as an example, that the real file-name corresponding to the le descriptor belongs to a known directory.
- Buffer length assertions. They check that the length of the string passed to a vulnerable function is not larger than the local buer to hold it.
- Conditional statements assertions. They prevent problems due to impossible paths by relating a system call and the expression in the guard of a conditional statement (important difference wrt self described as CFG only)

The Analyst - 1

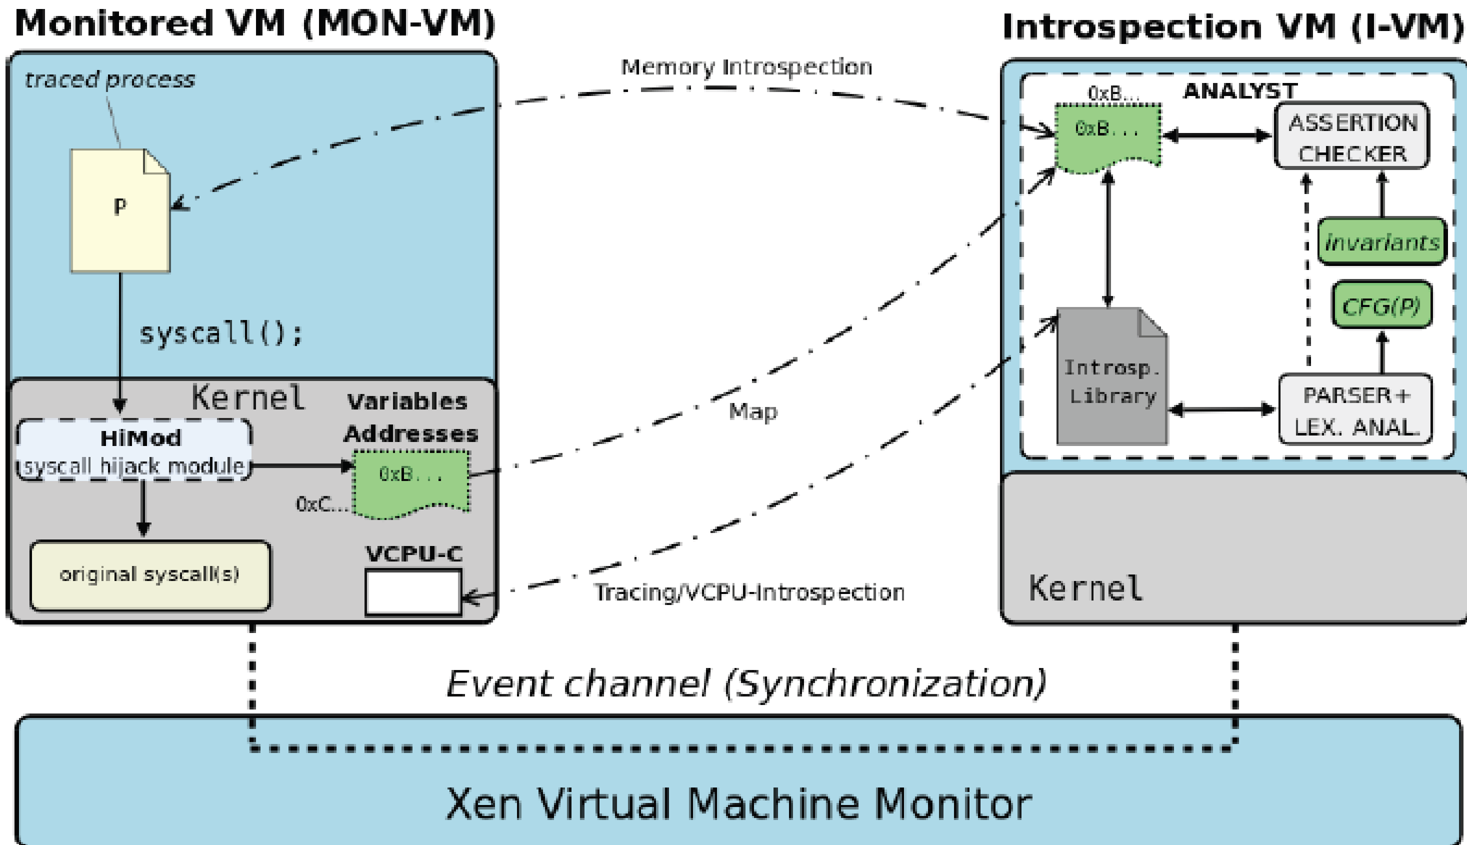

The Analyst - 2

The Analyst in the I-VM verifies the integrity of the self of P through:

- Lexical Analyzer: it verifies that the system call that P wants to issue belongs to the set of system calls returned by the static analysis of $SourceCode(P)$;
- Parser: it checks that the current trace of system calls issued by P is coherent with $CFG(P)$, i.e. it is a prefix of a word allowed by $CFG(P)$;
- Assertion Checker: it checks whether the invariant coupled with the current system-call holds.

Invariant Evaluation - 1

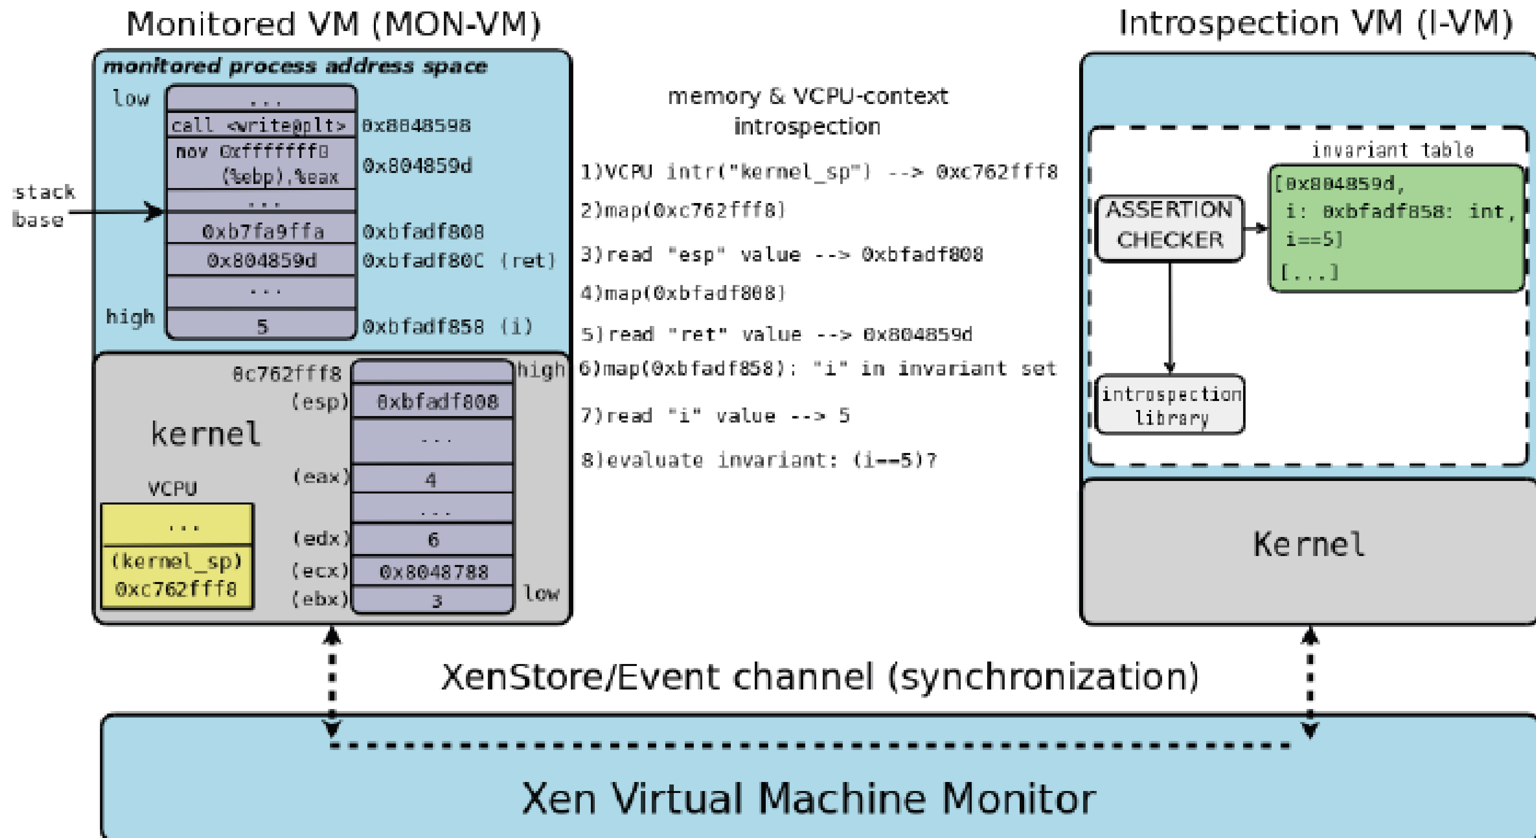

Invariant Evaluation - 2

- The Introspection VM runs an Assertion Checker that evaluates invariants on P memory status
- Access to the variables of P and to the CPU of the Monitored VM is implemented through an Introspection Library
- Every time P issues a system call the Introspection VM:
  (i) retrieves the system call number and the value of its parameter;
  (ii) determines the invariant coupled with the issued system call;
  (iii) retrieves the values of the variables that the invariant refers to;
  (iv) evaluates the invariant and:
    - kills P if the invariant is false
    - otherwise it resumes the execution of P.

Memory Monitoring Implementation

Monitoring Memory on Production-Level Systems

(1) **Passive Monitoring**: Viewing memory in A_i from S without any timing synchronization between the two virtual machines
(2) **Active Monitoring**: Viewing memory in A_i from S with event notification being sent from A_i to S to permit monitoring at relevant times
(3) **Locating Valuable Data**: Applying formal models or obtained from supervised learning to find critical data structures within the raw memory view

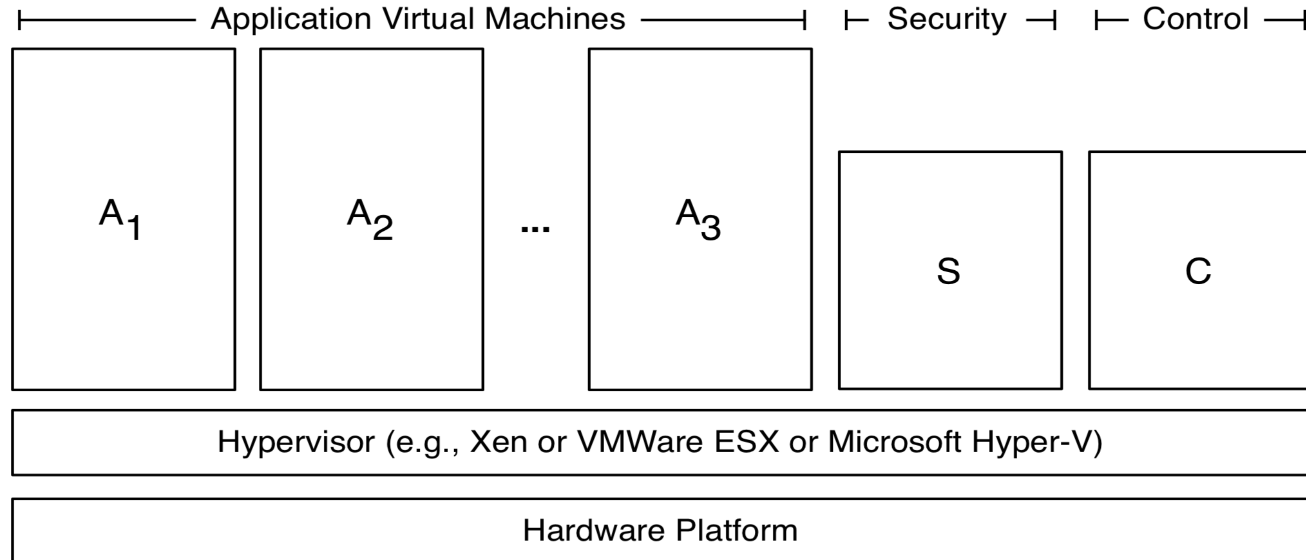

Xen overview

- Runs directly on the physical hardware
- Special management domain is called Dom0 to provide a management interface
- The VMM gives Dom0 system access to a control library
  - create, destroy, start, pause, stop, and allocate resources to VMs from Dom0
- Provides drivers for the host’s physical hardware
- Can also request that memory pages allocated to unprivileged VMs

Xen overview

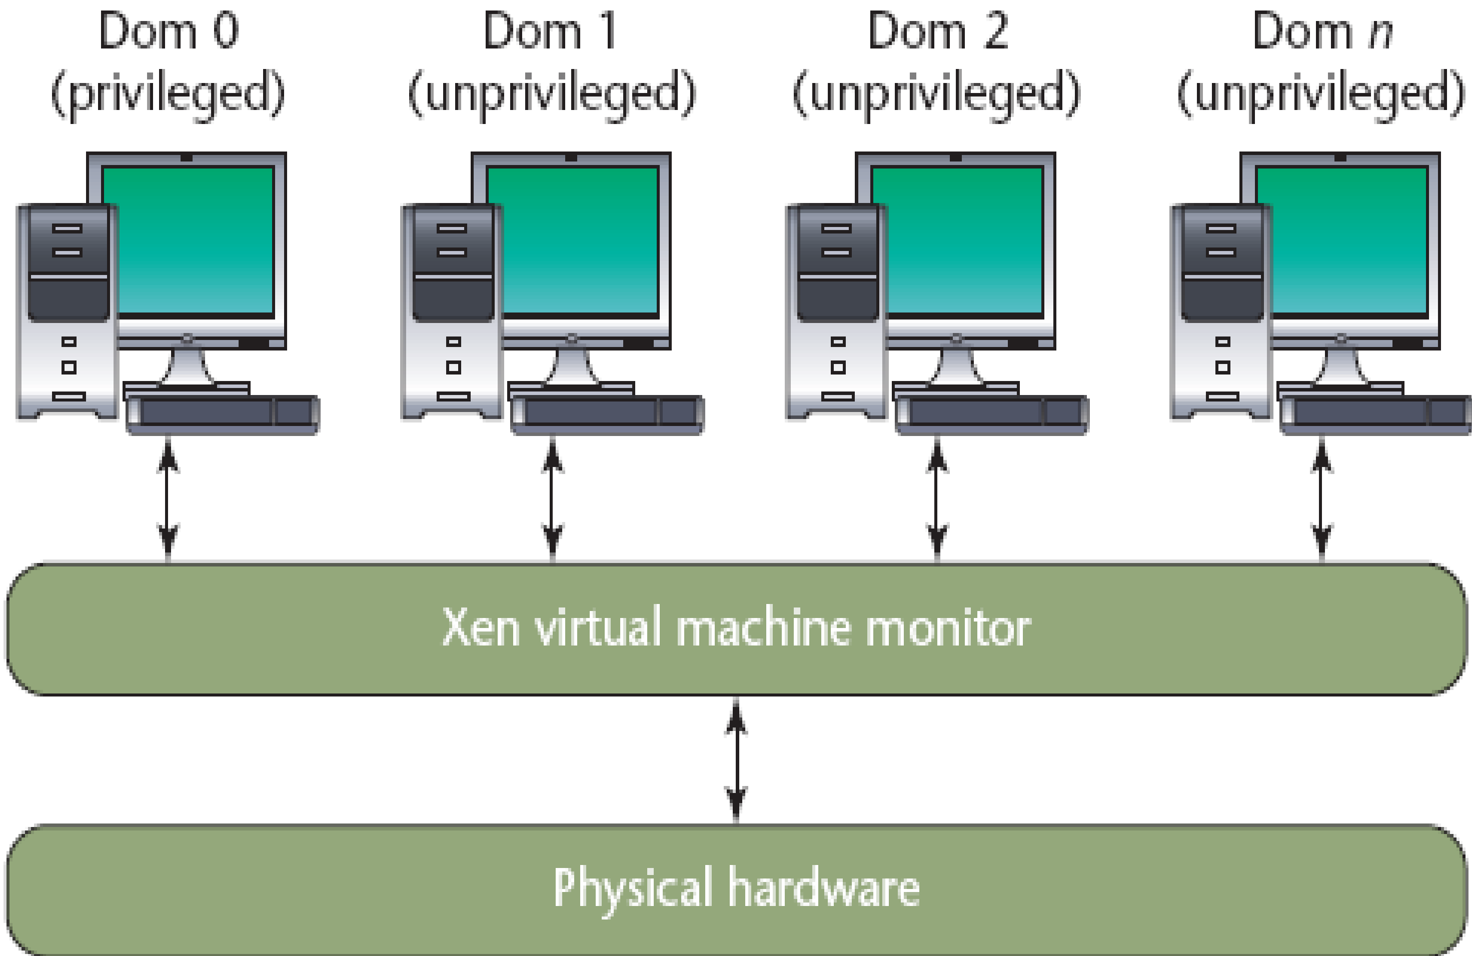

The XenAccess Library

- An open source VM introspection library
- Access to virtual addresses, kernel symbols, and more
- Works with Xen and dd-style memory image files
- Released in Spring 2006
- Maintained by Georgia Tech Inf. Sec. Center to encourage more research
- http://www.xenaccess.org

```
root@bluemoon:/home/bdpayne/
File  Edit  View  Terminal  Tabs  Help
[root@bluemoon examples]# ./process-list 1
[    4] System
[  420] smss.exe
[  468] csrss.exe
[  496] winlogon.exe
[  540] services.exe
[  552] lsass.exe
[  700] svchost.exe
[  760] svchost.exe
[  828] svchost.exe
[  876] svchost.exe
[  924] svchost.exe
[ 1220] spoolsv.exe
[ 1792] alg.exe
[ 1876] wscntfy.exe
[ 1952] explorer.exe
[  140] ctfmon.exe
[ 1924] procexp.exe
[root@bluemoon examp
```

```
/* initialize the xen access library */
xa_init(dom, &xai);

/* get the head of the list */
xa_read_long_sym(&xai, "PsInitialSystemProcess", &list_head);
memory = xa_access_virtual_address(&xai, list_head, &offset);
memcpy(&next_process, memory + offset + ActiveProcessLinks_OFFSET, 4);
list_head = next_process;

/* print out the first process */
name = (char *) (memory + offset + ImageFileName_OFFSET);
memcpy(&pid, memory + offset + UniqueProcessId_OFFSET, 4);
printf("[%5d] %s\n", pid, name);
munmap(memory, xai.page_size);

/* walk the process list */
while (1){
    /* follow the next pointer */
    memory = xa_access_virtual_address(&xai, next_process, &offset);
    memcpy(&next_process, memory + offset, 4);

    /* if we are back at the list head, we are done */
    if (list_head == next_process){
        break;
    }

    /* print out the next process */
    name = (char *) (memory + offset + ImageFileName_OFFSET -
           ActiveProcessLinks_OFFSET);
    memcpy(&pid, memory + offset + UniqueProcessId_OFFSET -
           ActiveProcessLinks_OFFSET, 4);
    printf("[%5d] %s\n", pid, name);
    munmap(memory, xai.page_size);
}

/* cleanup */
xa_destroy(&xai);
```

Passive Monitoring

To monitor application memory of another virtual machine we have to map the memory into an address of the monitoring one

Mapping “raw memory view” to virtual addresses and symbols requires the steps shown in figure below.

Address and symbol mapping can be performed by a VM introspection library (e.g., XenAccess)

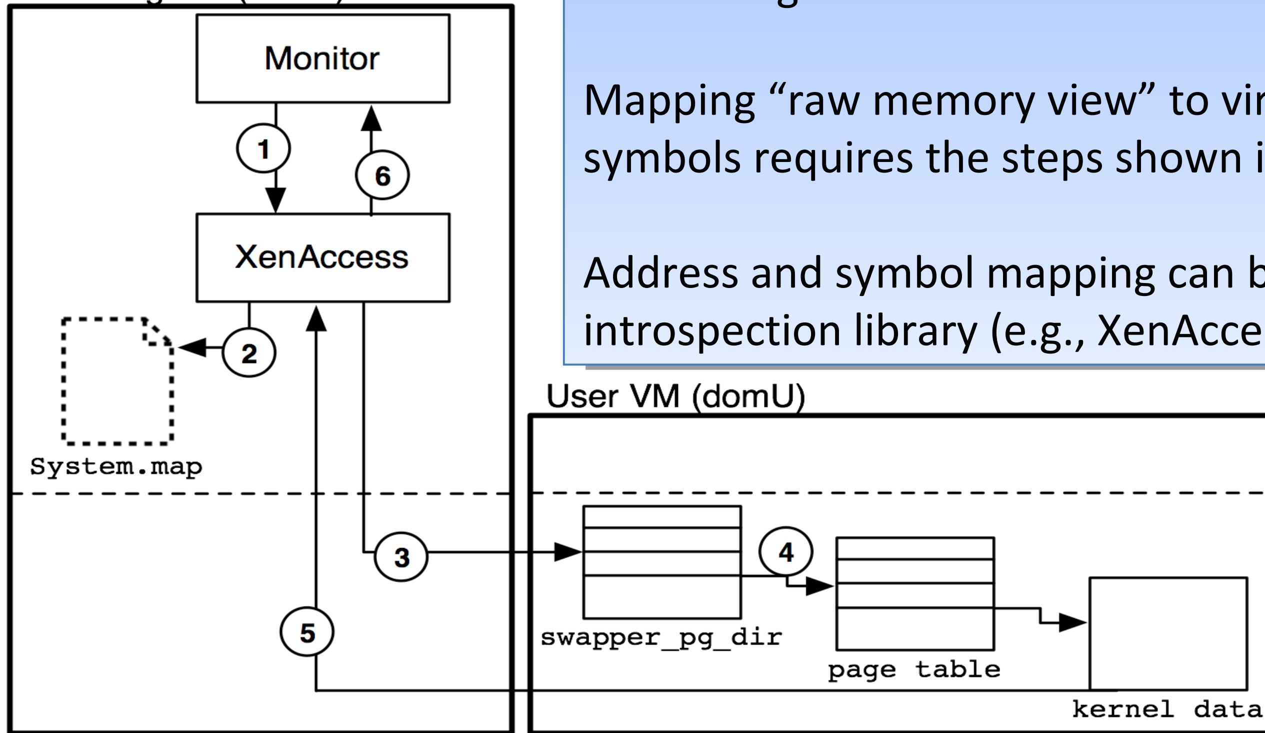

BD Payne, M Carbone, and W Lee. *Secure and Flexible Monitoring of Virtual Machines*. In ACSAC 2007.

Steps for Passive Monitoring

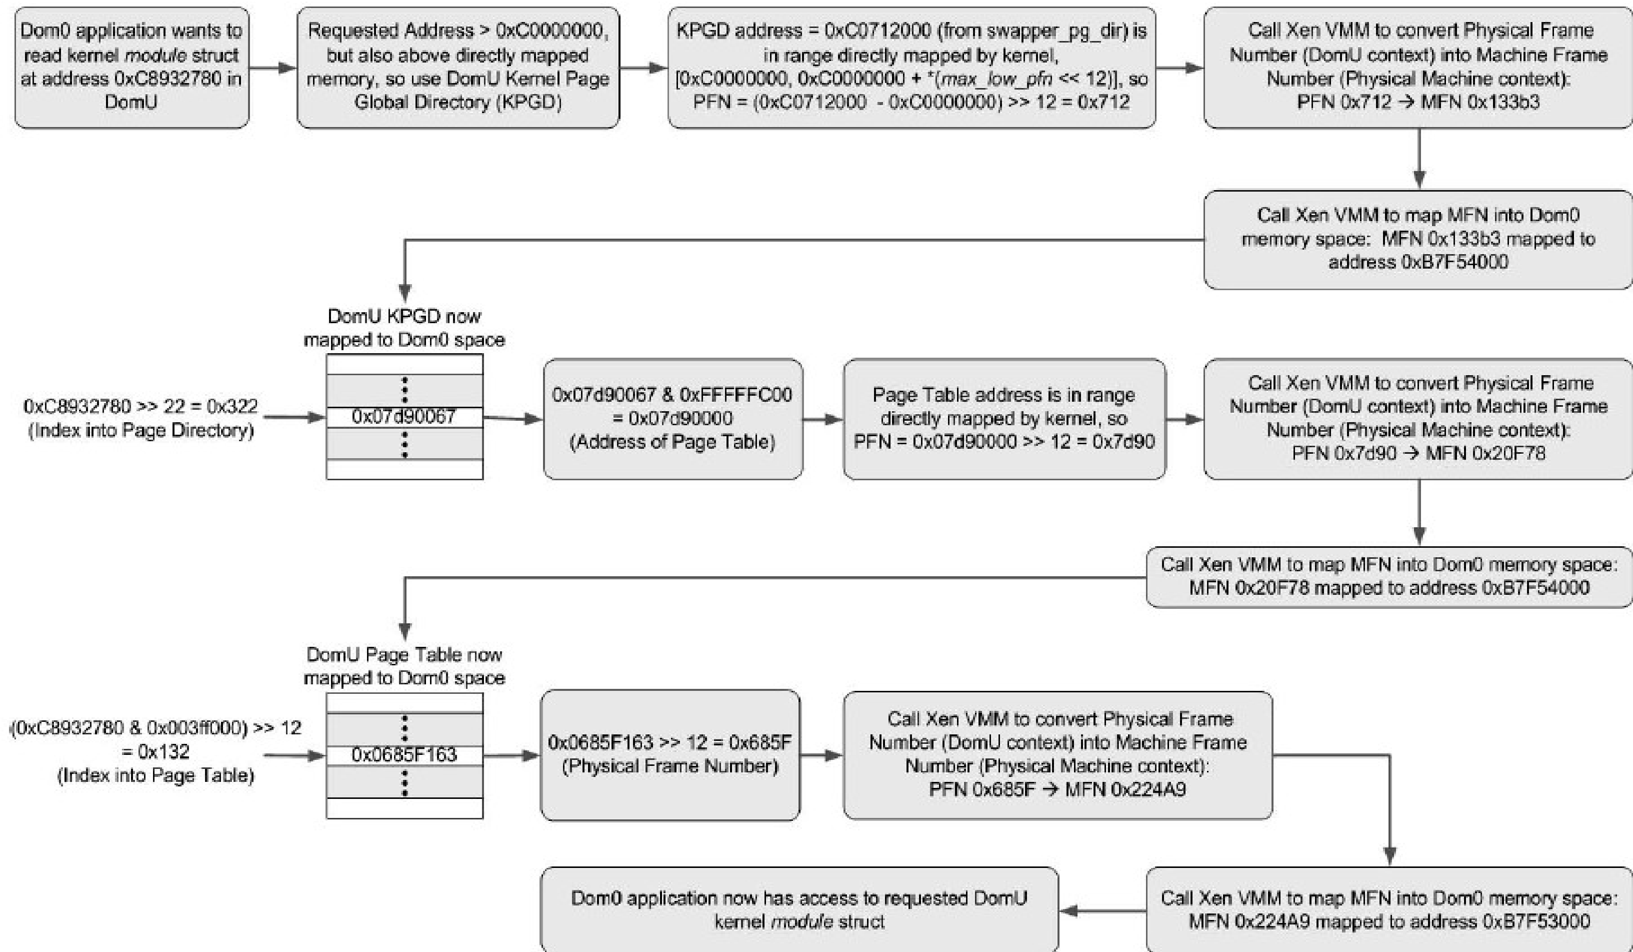

Active Monitoring

Monitoring application receives event notification from Guest VM when code execution reaches one of the hooks installed in the Guest VM kernel.

Hooks and all associated code are protected from tampering using hypervisor-enforced memory protections (i.e., User VM can not modify these security-critical components).

Hooks invoke trampoline, which transfers control to the security application.

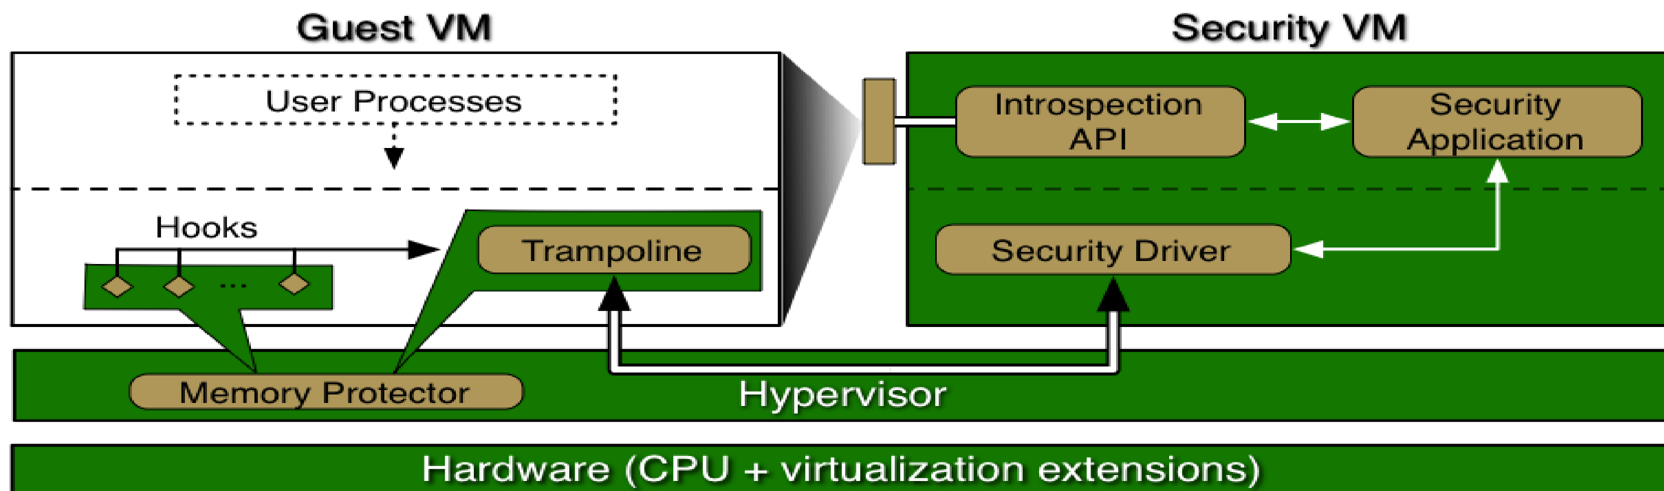

BD Payne, M Carbone, M Sharif, and W Lee. *An Arch for Secure Active Monitoring Using Virtualization*. In Oakland 2008.

Ether

- Use Intel VT hardware virtualization extensions to provide instruction execution on actual hardware
- Extend the Xen hypervisor to leverage Intel VT for malware analysis
- Provides for both instruction-by-instruction examination of malware, and also coarser grained system call-by-system call examination
- System Diagram:

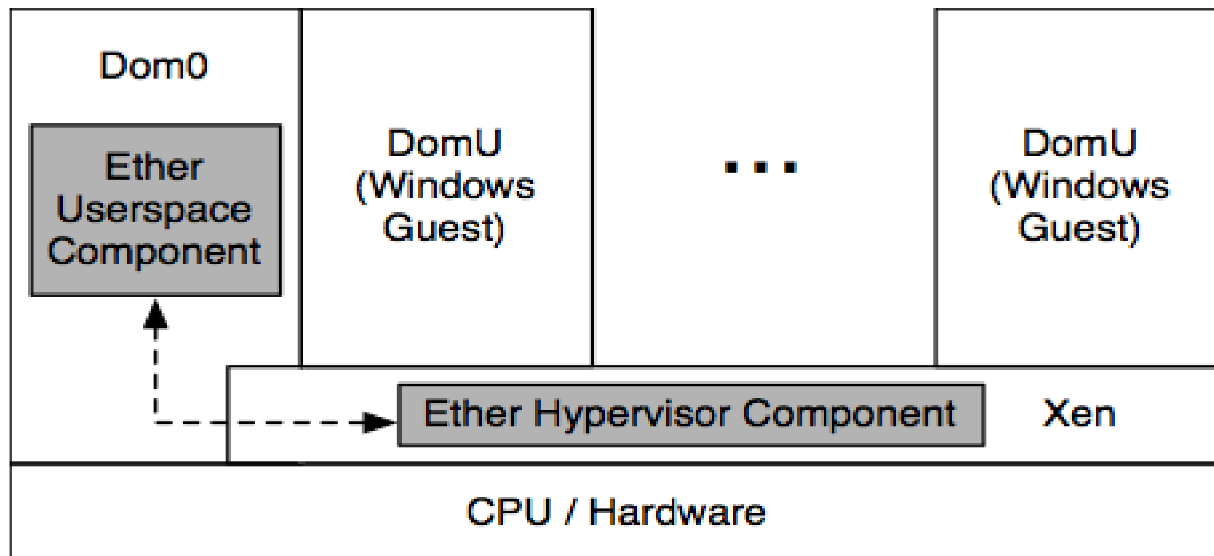

Ether: Experiments

- Two tools to test the Ether framework:
  - EtherUnpack: extracts hidden code from obfuscatd malware
  - EtherTrace: Records system calls executed by obfuscated malware

- Evaluation
  - EtherUnpack: how well current tools extract hidden code by obfuscating a test binary and looking for a known string in the extracted code
  - EtherTrace: a test binary which executes a set of known operations obfuscated and then observe if these operation were logged by the tool

Ether: EtherUnpack Results

Packing Tool	PolyUnpack	Renovo	EtherUnpack
Armadillo	no	no	yes
Aspack	no	yes	yes
Asprotect	yes	yes	yes
FSG	yes	yes	yes
MEW	yes	yes	yes
MoleBox	no	yes	yes
Morphine	yes	yes	yes
Obsidium	no	no	yes
PECompact	no	yes	yes
Themida	no	yes	yes
Themida VM	no	no	yes
UPX	yes	yes	yes
UPX Scrambled	yes	yes	yes
WinUPack	no	yes	yes
Yoda's Protector	no	yes	yes

Ether: EtherUnpack Results

PolyUnpack = Approach is based on the observation that sequences of packed or hidden code in a malware instance can be made self-identifying when its runtime execution is checked against its static code model.

Renovo = An approach based on the observation that sequences of packed or hidden code in a malware instance can be made self-identifying when its runtime execution is checked against its static code model.

Ether: EtherTrace Results

Packing Tool	Norman Sandbox	Anubis	EtherTrace
None	yes	yes	yes
Armadillo	no	no	yes
UPX	yes	yes	yes
Upack	yes	yes	yes
Themida	yes	yes	yes
PECompact	yes	yes	yes
ASPack	yes	yes	yes
FSG	yes	yes	yes
ASProtect	yes	no	yes
WinUpack	yes	yes	yes
tElock	yes	no	yes
PKLITE32	yes	yes	yes
Yoda's Protector	no	yes	yes
NsPack	yes	yes	yes
MEW	yes	yes	yes
nPack	yes	yes	yes
RLPack	yes	yes	yes
RCryptor	yes	yes	yes

VIX

- Virtual Introspection for Xen
- Place in the privileged Dom0 VM
- Interact through a stable API
- Reduce the application's ability to perform inline processing (requests in real time)

How VIX works

- Pauses operation of the target VM
- Maps some of its memory into the Dom0
- Acquires and decodes the memory pages
- Resumes operation of the target VM
- Reference task_struct data structures
  - process ID, process name, memory map, and execution time
- Traverses the list of task_structs

List of task_structs

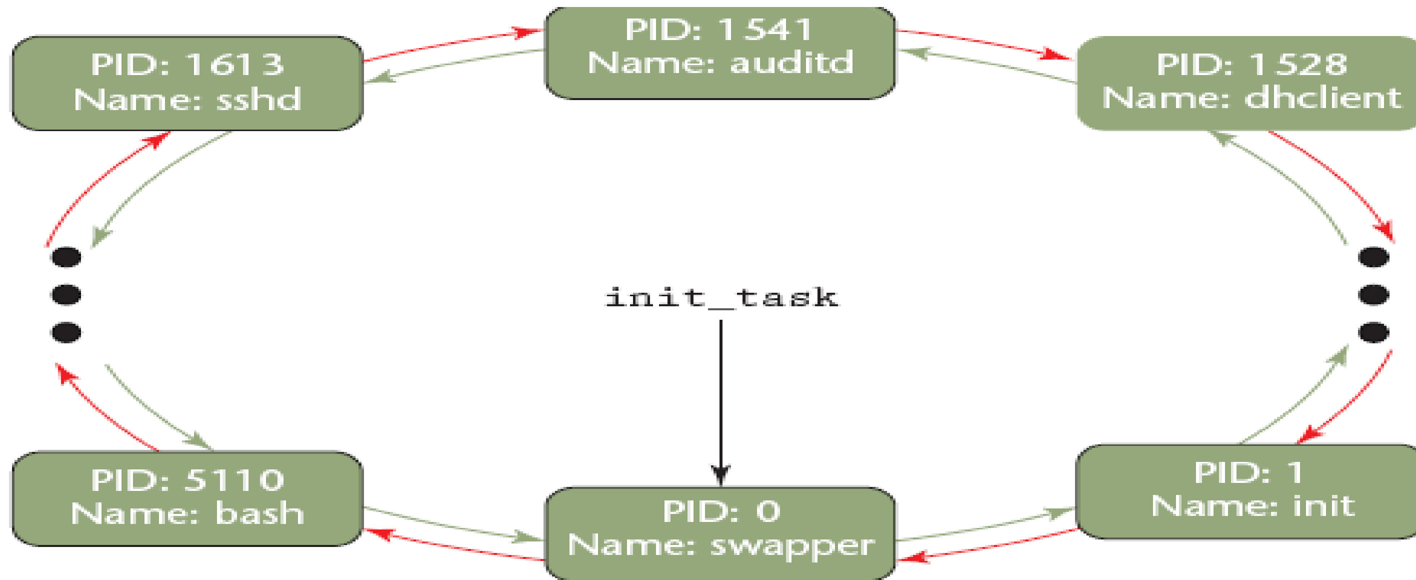

Linux stores this list as a circular double-linked list
Each kernel version has an associated memory address for the first process

VMI Functionality

- Not depend on any VM OS functionality for information
- VIX application
  - vix-ps, vix-netstat, vix-lsof, vix-pstrings, vix-lsmod, vix-pmap, and vix-top
- vix-ps
  - Traverse the entire task list
  - Output as the ps command

VM Introspection - VMware Initiatives

Security API's

- Designed for security productization
- Agent runs within a VM
- Capabilities
  - Memory access events
  - Selected CPU events
  - VM lifecycle events
  - Access to VM memory & CPU state
  - Page Table walker

Security APIs (VMsafe)

Goals

- Better than physical
  - Exploit hypervisor interposition to place new security agent
  - Provide security coverage for the VM kernel (and applications)
- Hypervisor as a Base of Trust
  - Divide responsibilities between the hypervisor and in-VM agent
  - The hypervisor covers the VM kernel, the rest is done from within the VM
  - Insure in-VM security agent execution and correctness
- Security as an infrastructure service
  - “Agent less” security services for VMs
  - Flexible OS independent solutions

Verify-Before-Execute Flow

Sample Introspection Agents

Verify-Before-Execute

- Utilize memory introspection to validate all executing pages

Flow

- Trace all pages for execution access

- On execution detection
  - Trace for page modification
  - Verify if page contain malware
  - Remove execution trace

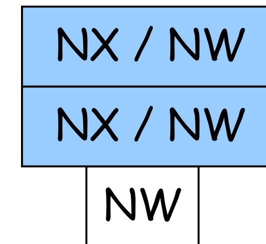

Is bad?

- On modification detection
  - Trace for execution
  - Remove modification trace

Security APIs – Use cases

VM Kernel coverage

- Detect infection in early boot process
  - Device opt ROM attacks
  - Boot loader
  - Boot records
  - OS image
- Detect code injection due to kernel vulnerabilities
- Detect self modifying code in kernel
- Lock kernel after initialization

Case Study - Microsoft Patch Guard

- Goal
  - Prevent patching of (x64 based) kernels
  - Force Indipendent Software Vendors to behave nicely
  - Prevent Root-kits ??
- Implementation
  - Obfuscated Invocation
  - Obfuscated Persistence
  - Evolving (Thanks to the awesome work from uninformed.org)
- What's The Problem?
  - Circumventable
  - Complicated
  - Only for x64 based Windows Systems

Security APIs – Use cases cont'

- Watch dog services
  - Liveness check for in-VM security agent
    - Detect agent presence
    - Verify agent periodic execution
    - Protect agent code and static data

TPM vs. Introspection

TPM

- Root of trust rely on hardware
- Passive device
- Platform and software stack decide what to measure
- Need software update to change measurement coverage
- Can not detect compromise in software stack since verification

VM Introspection

- Root of trust rely on hypervisor
- Introspection agent have the initiative
- Security vendor / policy dictate what to measure
- Coverage is content, and can change independently of VM
- Designed to continuously scan VMs and to detect compromise

VMsafe – Network Introspection

- Capabilities
  - Place an inline network agent on any VM virtual nic
  - Allow reading, injecting, modifying, and dropping packets.
- Benefits
  - Efficiently monitor inter-VM network communication
  - Integrated support for live migration
- Virtualization only applications
  - Correlate VM internals with network policy. (using CPU/ Memory inspections one can learn OS version, patch level, configuration etc)
  - Build a trusted distributed firewall.

Retrospective Security

- Motivation
  - Detect whether you have been attacked in the past
  - Detect if you might be still compromised by a past attack
- Method
  - VMware Record & Replay allow for a deterministic replay of VM using recorded logs
  - Potentially the recordings have captured an attack
  - The security API's are detached from the recorded VM (unlike in-VM agent) and can attach to a replay session

Retrospective Security

- What is it good for?
  - Run more aggressive policies that will not be acceptable in production environments
  - Discover 0days used to exploit your system
  - Learn how the malware / attacker have navigated your system
  - Use data tainting technique to detect any side effects that still exist on your system
  - Possibly clean the finding from last step on your production VM.
  - Learn about the scope of the damage done to your system, i.e. what is the extent of data leakage

Security vs. Hardware Virtualization

1st Generation – SVM, VT-X
- VMM no longer need to run the VM kernel under binary translation
- Security Trade off – Code Breakpoint, Guest code patching (while translating), Control flow visibility

2nd Generation – NPT, EPT
- VMM no longer need to have software based MMU
- Security Trade off – Tracking LA->PA mapping is becoming expensive, resulting with inability to operate on linear addresses.

3rd Generation – IO MMU, VT-D
- VMM can assign physical devices to VMs without worry of VM escape or hypervisor corruption
- Security Trade off – Avoids interposition on the pass-thru device

1 generation

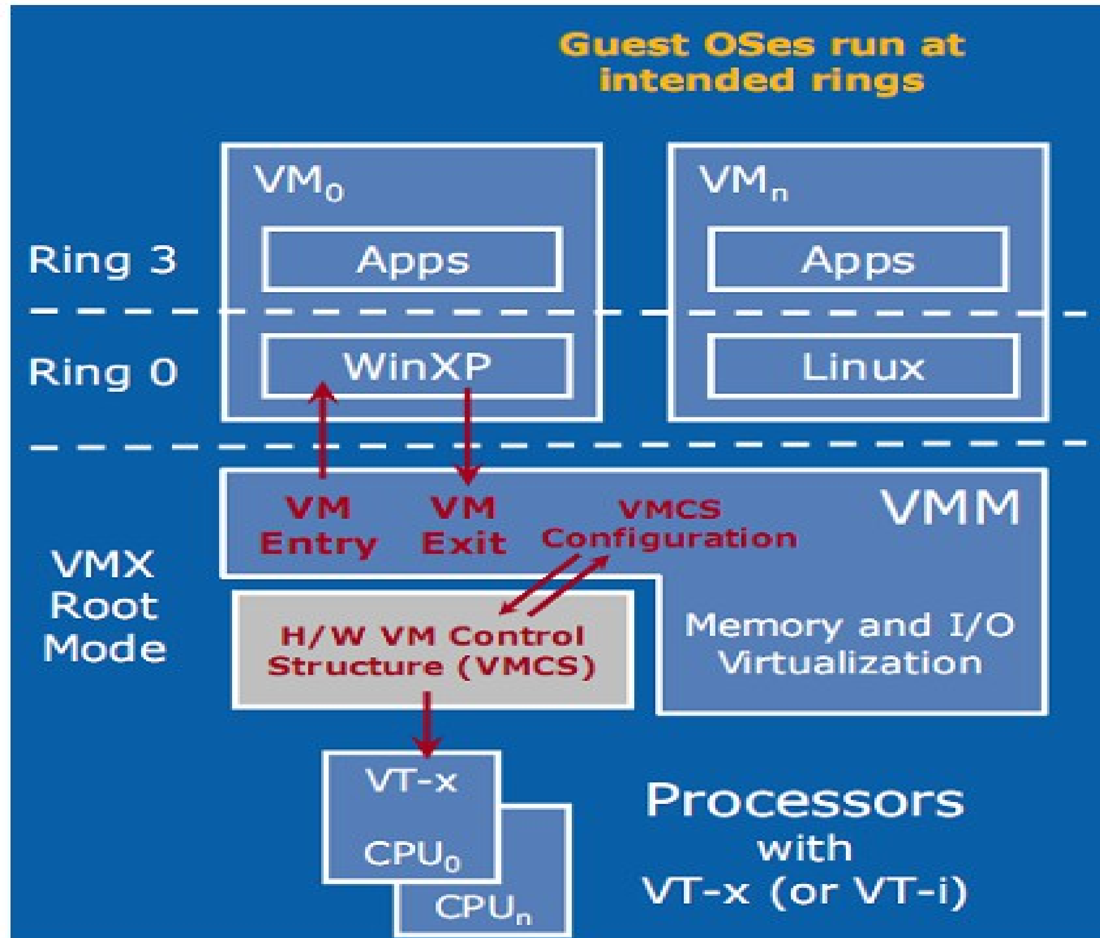

With HW virtualization, the guest OS is back where it belongs: ring 0.

Shadow page table

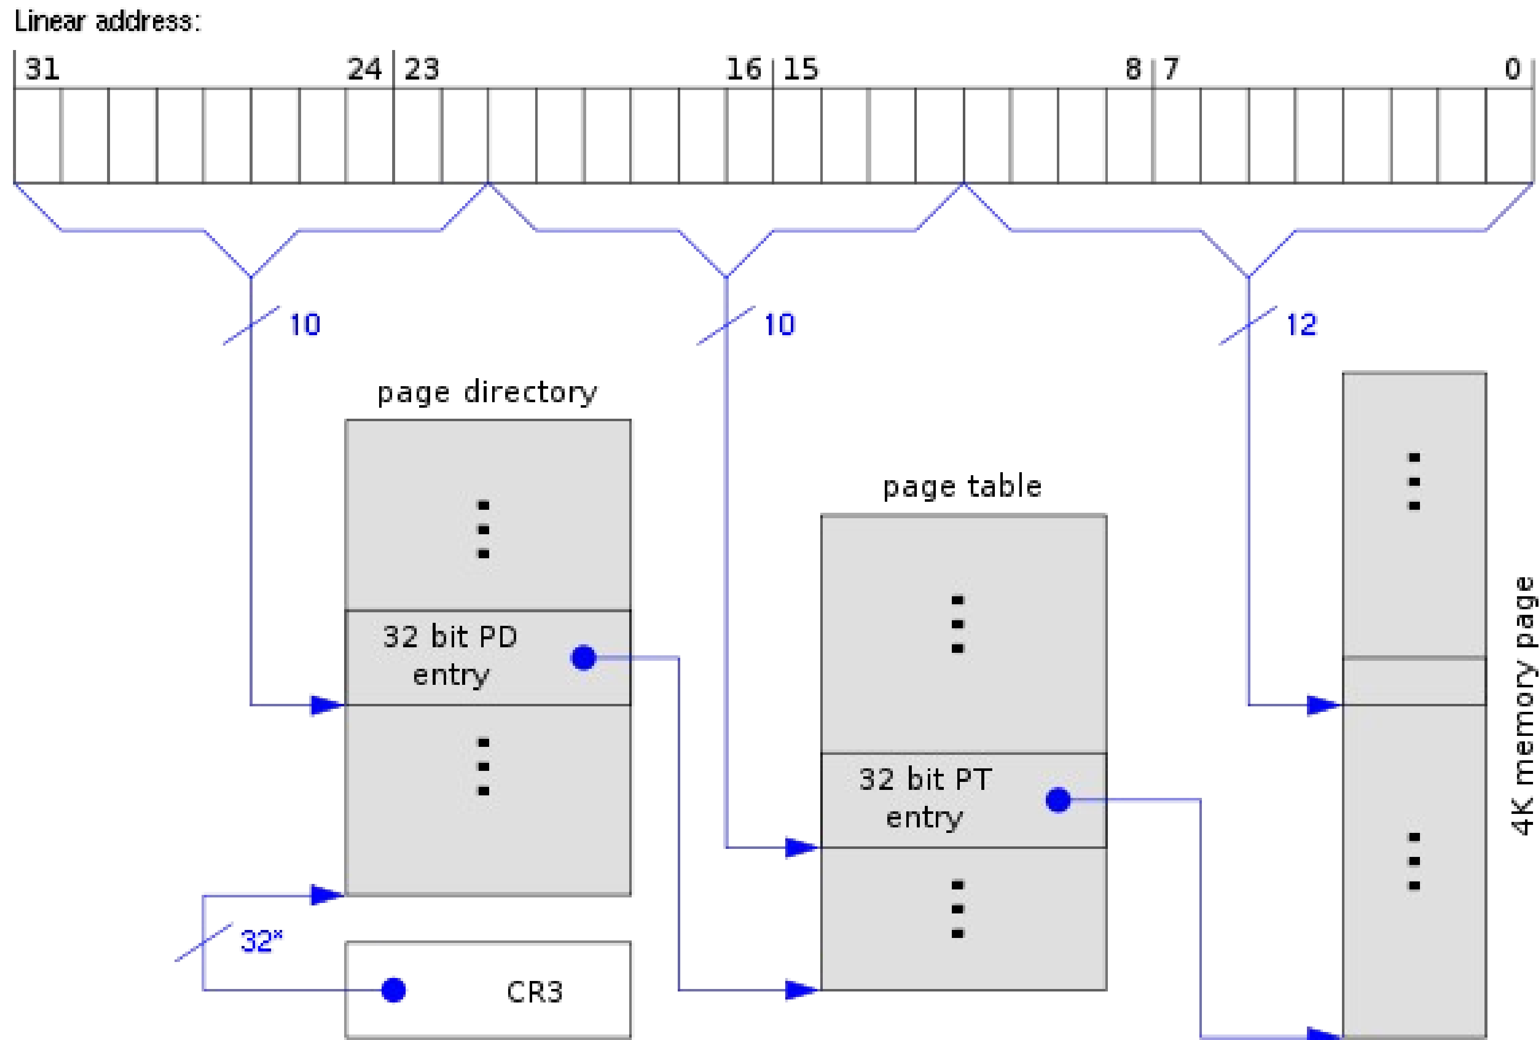

Shadow page table

- Three abstractions of memory
  - Machine: actual hardware memory
    - 2 GB of DRAM
  - Physical: abstraction of hardware memory managed by OS
    - If a VMM allocates 512 MB to a VM, the OS thinks the computer has 512 MB of contiguous physical memory
    - (Underlying machine memory may be discontiguous)
  - Virtual: virtual address spaces you know and love
    - Standard 232 address space
- In each VM, OS creates and manages page tables for its virtual address spaces without modification
- But these page tables are not used by the MMU hardware

Shadow page table

VMM creates and manages page tables that map virtual pages directly to machine pages

- These tables are loaded into the MMU on a context switch
- VMM page tables are the shadow page tables
- VMM needs to keep its $V \rightarrow M$ tables consistent with changes made by OS to its $V \rightarrow P$ tables
- VMM maps OS page tables as read only
- When OS writes to page tables, trap to VMM
- VMM applies write to shadow table and OS table, returns
- Also known as memory tracing

Shadow page table

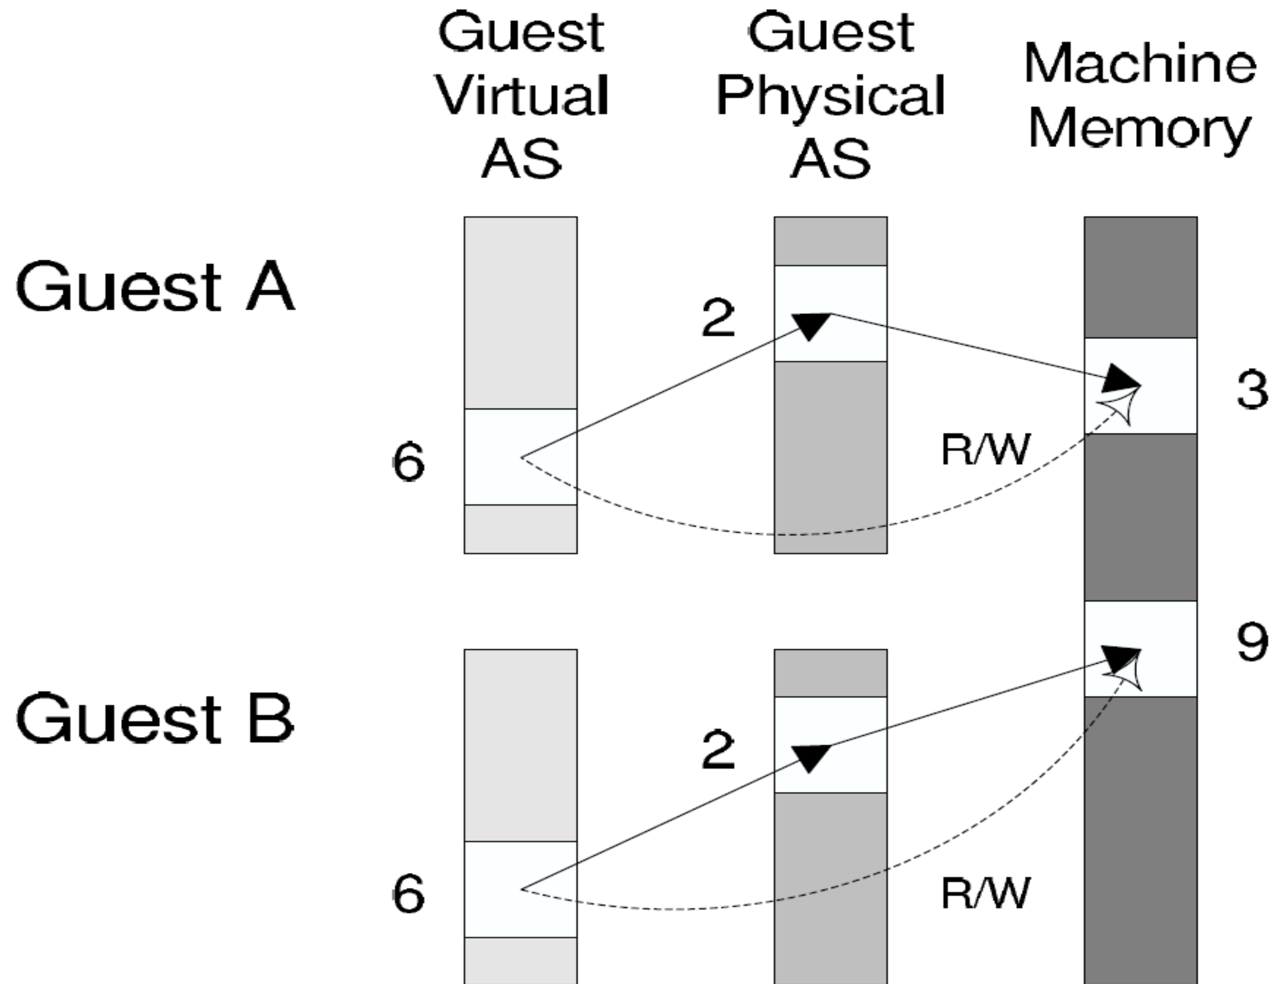

2 generation

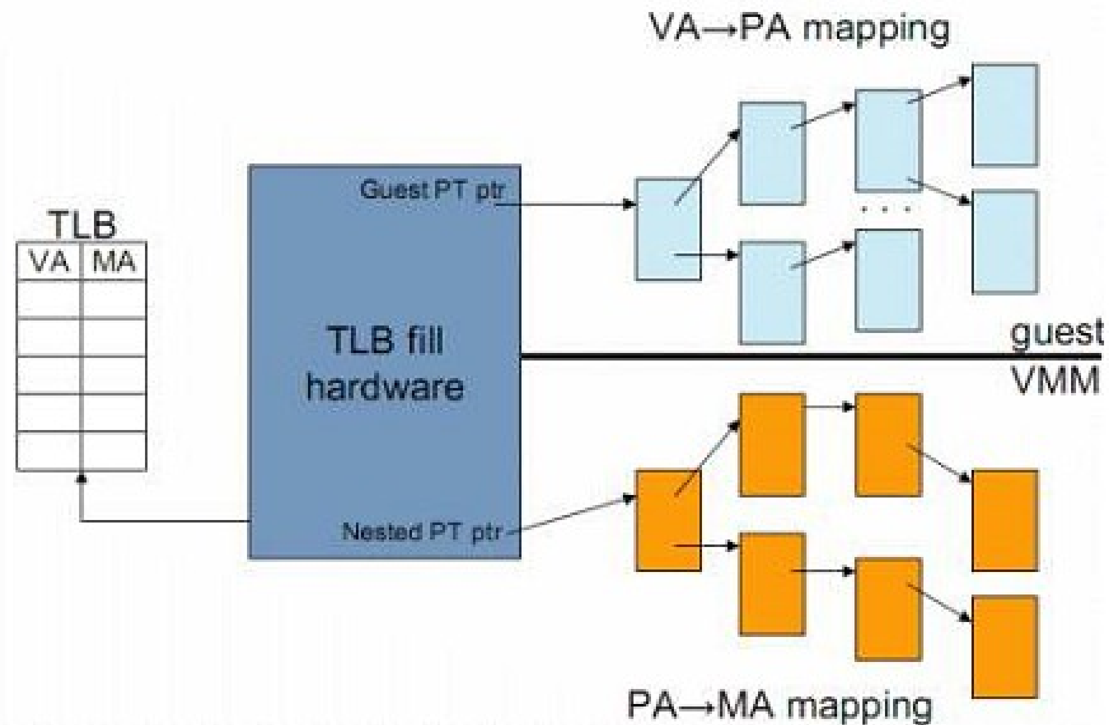

EPT or Nested Page Tables is based on a "super" TLB that keeps track of both the Guest OS and the VMM memory management.

Second generation: Intel EPT & AMD NPT

- Eliminating the need to shadow page table

Future Extensions: EPT

EPT: Overview

CR3
EPT Base Pointer
Guest Linear Address
Intel® 64 Page Tables
Guest Physical Address
EPT Page Tables
Host Physical Address

- Intel® 64 page tables
  - Map guest-linear to guest-physical (translated again)
  - Can be read and written by guest
- New EPT page tables under VMM control
  - Map guest-physical to host-physical (accesses memory)
  - Referenced by new EPT base pointer
- No VM exits due to page faults, INVLPG, or CR3 accesses

intel

8

Third generation: Intel VT-d & AMD IOMMU

- I/O device assignment
  - VM owns real device
- DMA remapping
  - Support address translation for DMA
- Interrupt remapping
  - Routing device interrupt

Threat Monitoring/Interfering

- Other approaches are possible
- An important classification is
  - Monitor subject
  - Interfere with subject
- Only monitor subject behavior
  - Livewire
  - Monitor a system can only detect and report problems
- Interfere with subject behavior
  - LycosID, µDenali
  - Can actually respond to a detected threat
  - Might terminate the relevant processes or VM
  - Might reduce the resources available to the VM (starve the attacker)

Livewire

- An early host-based intrusion detection system
- Monitors VMs to gather information and detect attacks
- Merely reports it rather than interfering

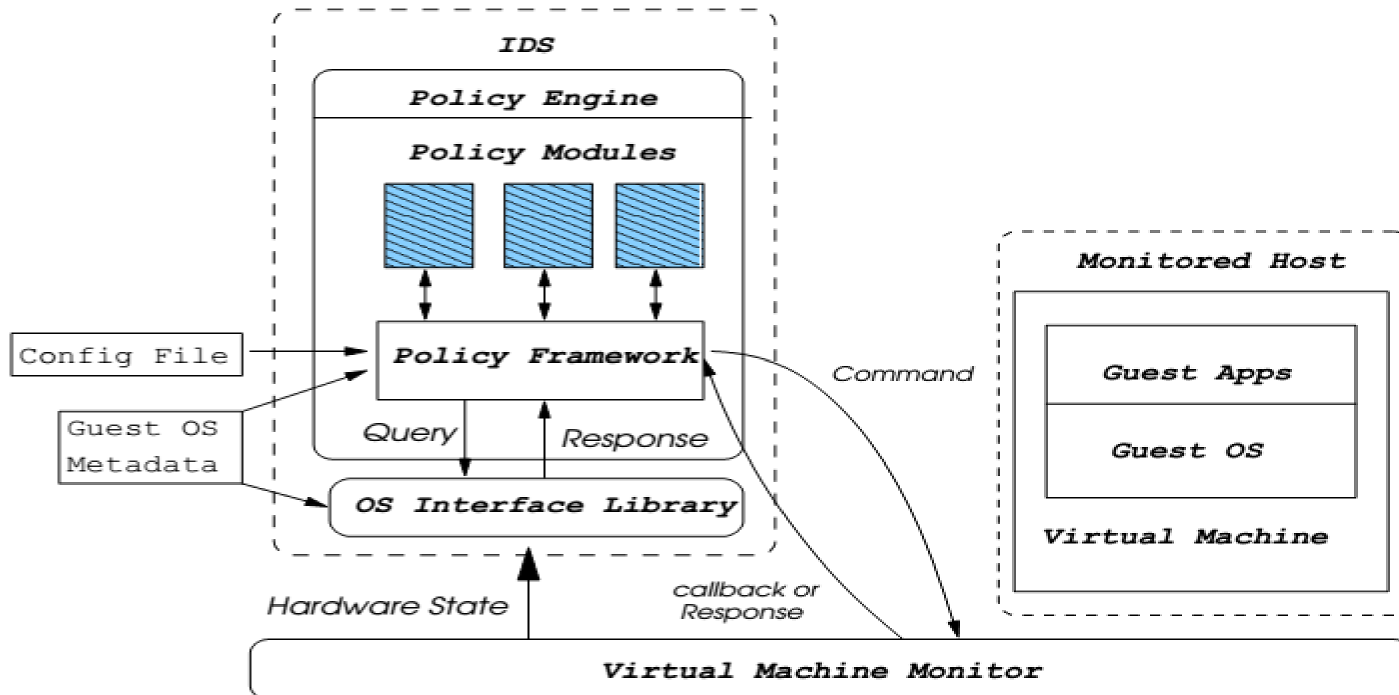

LycosID

- Uses crossview validation techniques to compare running processes
- Patches running code to enable reliable identification of hidden processes

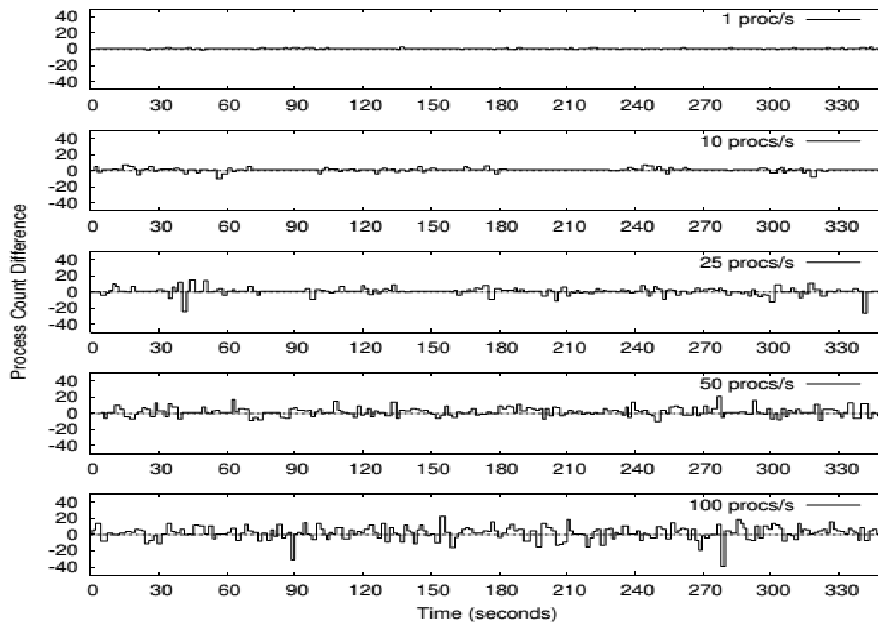

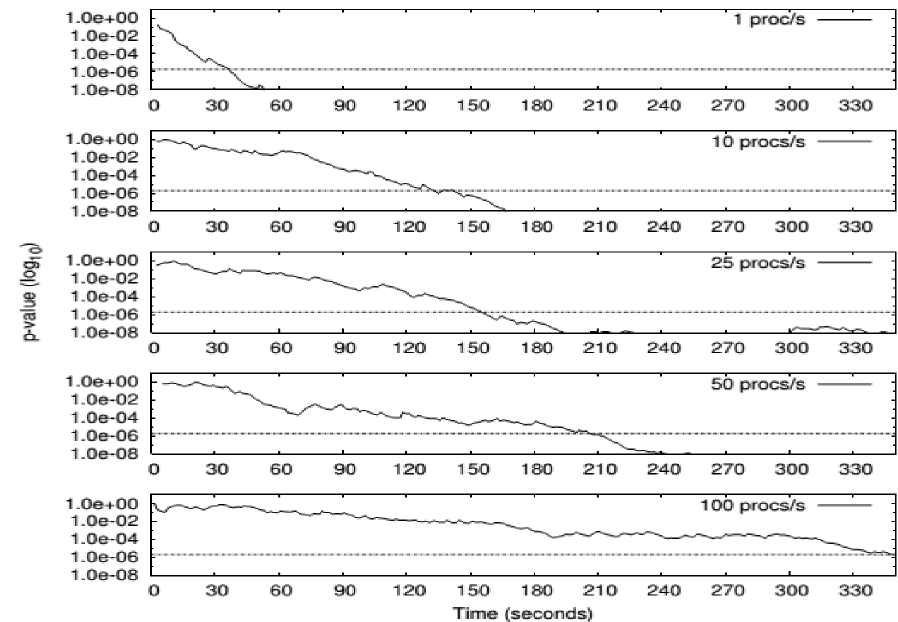

Manitou

- A VMI designed to detect malware
- Compares known instruction-page hashes with memory-page hashes at runtime before starting a program
- The instruction-page is corrupted and nonexecutable for no match
- A self attestation model

µDenali

- Acts as a switch for network requests to a set of VMs
- Can force a VM reboot
- Its first goal is designing and implementing mechanisms for lightweight VMMs, virtual machines, and guest operating systems, so that 100s or 1000s can concurrently execute
- An ancillary challenge implied by this is resource management across virtual machines: to fully isolate one VM from another

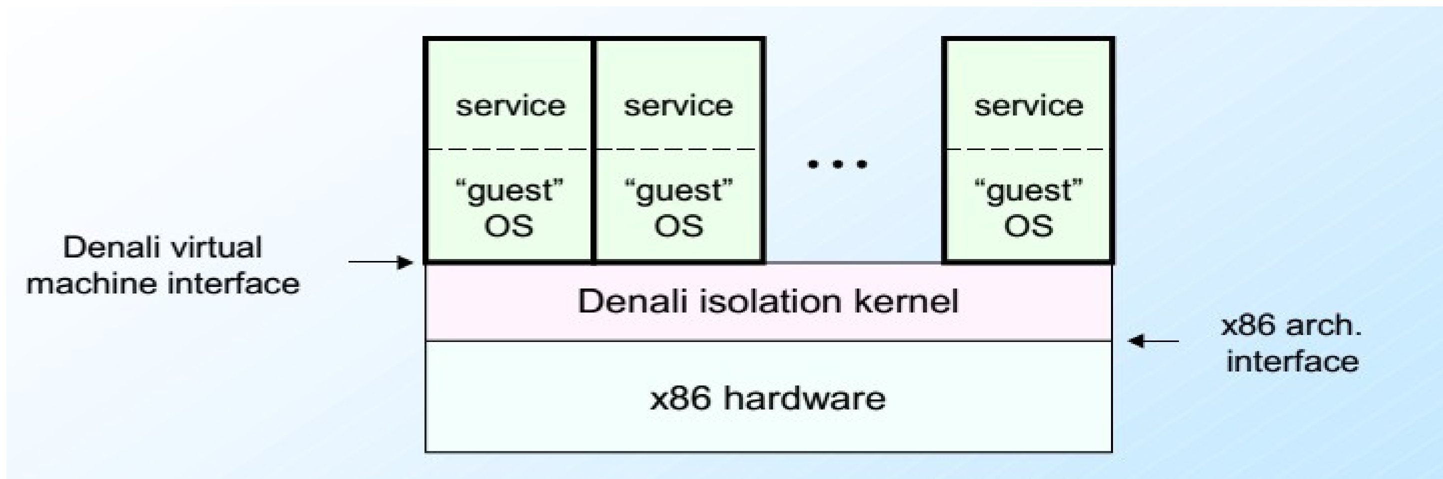

μDenali

- **Isolation kernel serves two roles**
  - virtualization: exposes the Denali virtual interface
  - resource management: multiplexes physical resources across virtual machines
- **Fairly standard mechanisms**
  - per-VM kernel thread stack, timer driven context switching, paging regions striped across disks
- **Two policies for CPU multiplexing**
  - gatekeeper: enforces admission control, by selecting a subset of active machines to admit into system
  - scheduler: controls context switching among active machines
    - round-robin scheduling

µDenali

- **A virtual I/O device is basically a queuing system**
  - virtual ethernet NIC has two queues
    - incoming (Rx) packet queue
    - outgoing (Tx) packet queue
- **Isolation kernel multiplexes and demultiplexes data from queues**
  - two policy questions:
    - what is the queueing discipline?
    - how many buffers should be allocated to each queue?

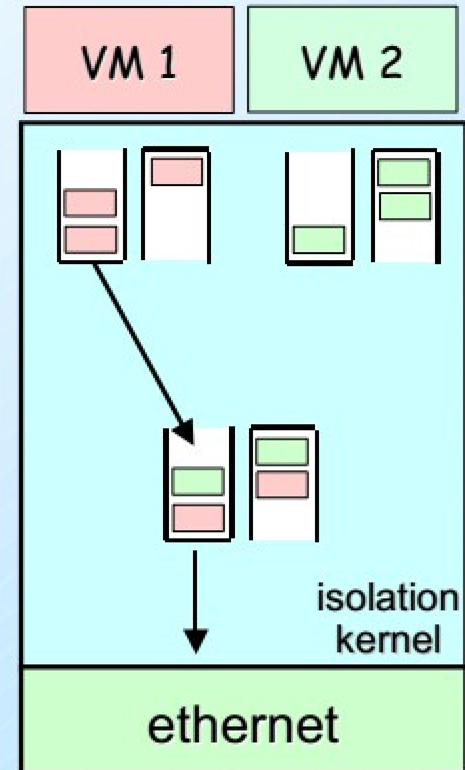

Semantic Awareness

- Account for different guest OS
- provide information that is more detailed
- Parse kernel memory to build a process table map
- Unaware VMI simply see memory as bits

LARES

- Gives each VM an internal "hook"
  - Activate an external monitoring control upon execution
- Monitor can interrupt execution and pass control to a security mechanism
  - The hook is injected into the VM OS
  - Hypervisor write-protects both the hook and the transfers control
  - Triggers at a meaningful system execution point

Semantically Unaware (AntFarm)

- Monitor the VM's memory management unit
- Can construct the virtual-to-physical memory mapping
- Infer information about the machine's processes and OS
- Anticipatory Scheduling

	Process Create	Addr Spc Create	Inferred Create	Process Exit	Addr Spc Exit	Inferred Exit	Context Switch	CS Inferred
Linux 2.4 x86								
Fork Only	1000	1000	1000	1000	1000	1000	3331	3331
Fork + Exec	1000	1000	1000	1000	1000	1000	3332	3332
Vfork + Exec	1000	1000	1000	1000	1000	1000	3937	3937
Compile	815	815	815	815	815	815	4447	4447
Linux 2.6 x86								
Fork Only	1000	1000	1000	1000	1000	1000	3939	3939
Fork+Exec	1000	**2000**	**2000**	1000	**2000**	**2000**	4938	4938
Vfork + Exec	1000	1000	1000	1000	1000	1000	3957	3957
Compile	748	**1191**	**1191**	748	**1191**	**1191**	2550	2550
Windows								
Create	1000	1000	1000	1000	1000	1000	74431	74431
Compile	2602	2602	2602	2602	2602	2602	835248	835248

IntroVirt

- It supports the construction of vulnerability specific predicates
- Attempt to bridge the "semantic gap" between
  - The VMI application
  - The target VM
- Using functionality on the target VM itself to lend context to the acquired data
- Basic mechanism insert assertion + replay VM

IntroVirt: the patch complexity

Application	Reference	Description of bug	Type of bug	# lines in	
				pred	patch
Linux kernel	CAN-2003-0961	integer overflow in do_brk	integer overflow	8	2
OpenSSL	CAN-2002-0656	SSL2 client master key arg buffer overflow	buffer overflow	7	3
squid	CAN-2005-0173	squid_ldap_auth incorrectly handles usernames w/ spaces	malformed input	27	20
Linux kernel	CAN-2004-0109	ISO9660 fs long symlink buffer overflow	buffer overflow	41	17
find	[20]	TOCTTOU race condition	race condition	63	N/A
bind	CAN-2005-0033	buffer overflow in q_usedns	buffer overflow	16	2
emacs	CAN-2005-0100	format string vulnerability in movemail utility	format string	9	1
gv	CAN-2002-0838	unsafe call to sscanf	buffer overflow	4	2
imapd	CAN-2005-0198	incorrect logic in CRAM-MD5 authentication	logic error	6	1
Linux kernel	CVE-2003-0985	mremap zero-area VMA remapping vulnerability	missing validation	8	2
Linux kernel	CVE-2004-0077	mremap missing do_munmap return value check	missing validation	15	7
Linux kernel	CAN-2004-0415	file offset pointer race condition	race condition	107	90
osCommerce	CAN-2005-0458	cross-site scripting vulnerability in contact_us.php	malformed input	27	1
phpBB	CAN-2004-1315	code injection via highlight parameter	malformed input	30	1
smbd	CAN-2003-0201	buffer overflow in call_trans2open	buffer overflow	10	1
squid	CAN-2005-0094	buffer overflow in gopherToHTML	buffer overflow	8	4
util-linux	CVE-2002-0638	chsh/chfn temporary file race condition	race condition	25	1
wu-ftpd	CVE-2000-0573	format string vulnerability in lreply	format string	16	4
wu-ftpd	CAN-2003-0466	off-by-one bug in fb_realpath	off-by-one	11	1
xpdf/cups	CAN-2005-0064	decryption function buffer overflow vulnerability	buffer overflow	7	2

Event Replay

- Ability to replay, or log events on a VM is useful
  - Debugging OSs
  - Replaying compromises
- VM must record enough information to reconstruct interesting portions
- The penalty is to record extra information

Revirt

- An example of a logging VMI
- Serves as the basis for time-traveling VMs that allow replay from any previous VM state

ReVirt

Workload	Runtime with logging (normalized to UMLinux *without* logging)	Log growth rate	Replay runtime (normalized to UMLinux *with* logging)
POV-Ray	1.00	0.04 GB/day	1.01
kernel-build	1.08	0.08 GB/day	1.02
NFS kernel-build	1.07	1.2 GB/day	1.03
SPECweb99	1.04	1.4 GB/day	0.88
daily use	≈ 1	0.2 GB/day	0.03

IaaS, Overlay and Security

- Trust in the network of VMs that is mapped onto the cloud may be increased by inserting into the networks some VMs that monitors the self of the processes running on the VMs and the data these VMs exchanges
- The coverage of the checks on process and data can be increased by minimizing the number of processes that each VM runs i.e. by increasing the number of VMs
- This has obvious advantages in term of safety and of performance and overcomes the classical distinction between host and network IDS e.g. to protect a process Pcrit running a critical application we can
  - Map it onto a distinct VM, VMcrit
  - Introduce a further VM to protect Pcrit self
  - Monitoring the communication to/from VMcrit even from processes that where mapped onto the same physical node

VINCI: Virtual Interacting Network CommunIty

- A software architecture that exploits virtualization to share in a secure way a cloud system.
- It decomposes users into communities: a set of users, their applications, a set of services and of shared resources.
- Users with distinct privileges and applications with distinct trust levels belong to distinct communities = Each community is paired with a level that defines the security requirements and the trust in the community
- Each community is supported by a virtual community network = VCN
  - a structured
  - highly parallel

  overlay network that interconnects VMs built by instantiating one of a predefined set of VM templates.

VM templates

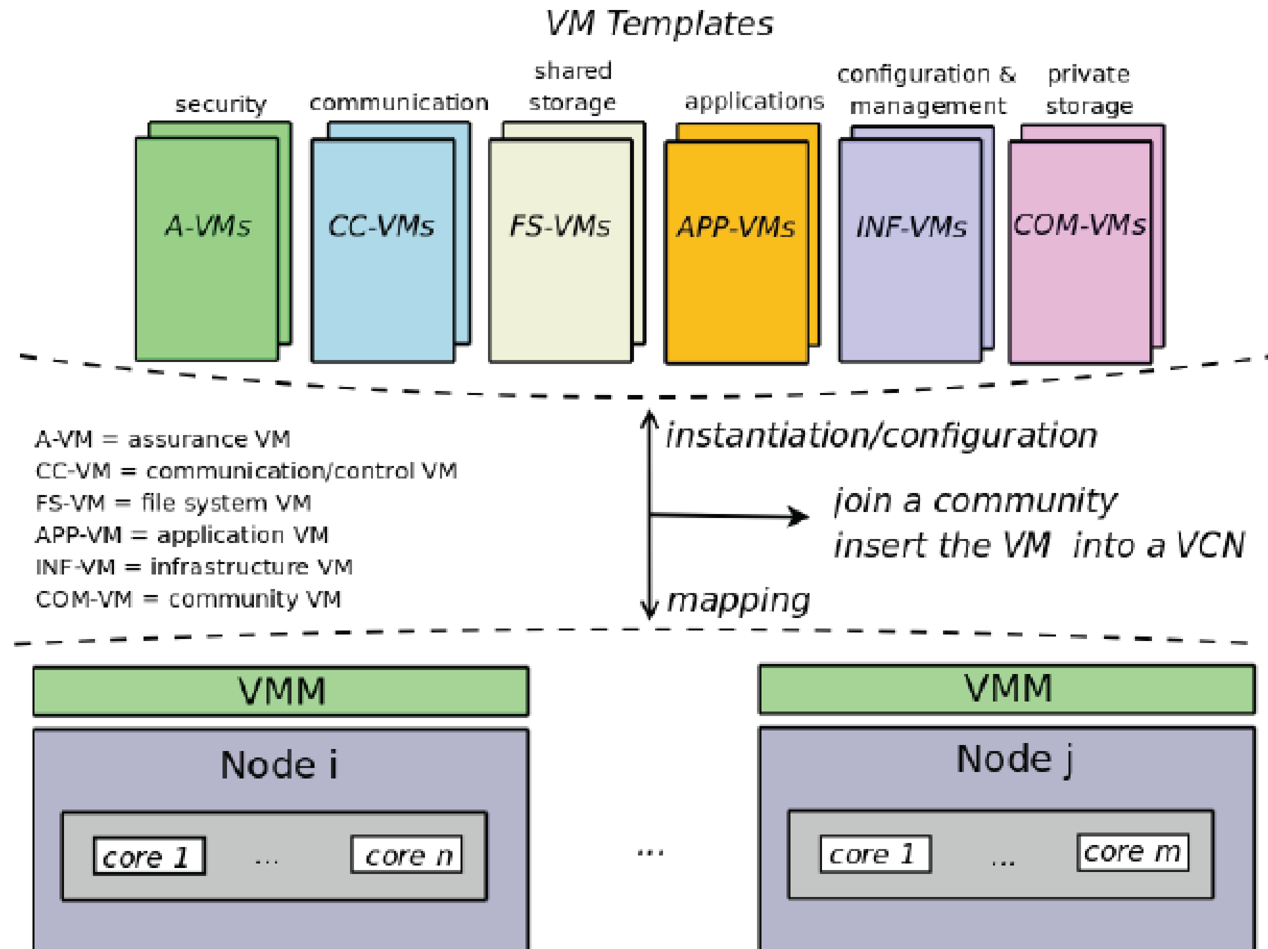

Interaction

- Direct Interaction within VMs in the same community
- Community VM manages resources shared among the same community
- Firewall VM protect the communications among VMs of distinct communities e.g. communication routed through nodes shared with low level communities are encrypted
- File System VM manages resources shared among distinct community by applying
  - a security policy based upon the level of the various community (e.g. Bell La Padula policy)
  - Tainting to protect data
- Some VMs are introduced to manage the infrastructure (VM mapping, allocation, migration)

Tainting

APPROACH

- Track OS-level information flow provenance by assigning a unique identifier (color) to each potential malware entry point
- Color individual processes/data based on their interaction with potential entry points or other previously colored processes/data
- Color-based identification of malware contaminations
- Color-based reduction of log data to be analyzed
- Highlight event anomalies via abnormal color interactions present in logs
- Leverage virtual machine technology for tamper resistance of log coloring

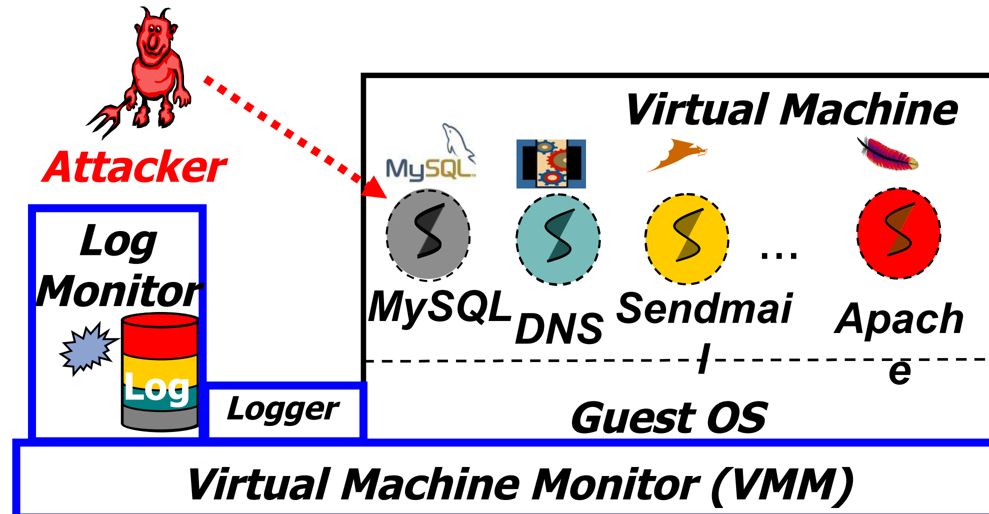

Dynamic taint analysis

1. Taint analysis should be applied anytime a malicious user input can be the vector of an attack. Very important even in the case of web applications.
2. Mark input data as “tainted”
3. Monitor program execution to track how tainted attributes propagate
4. Check when tainted data is used in dangerous ways

Dynamic taint analysis

TaintCheck performs dynamic taint analysis on a program by running the program in its own emulation environment.

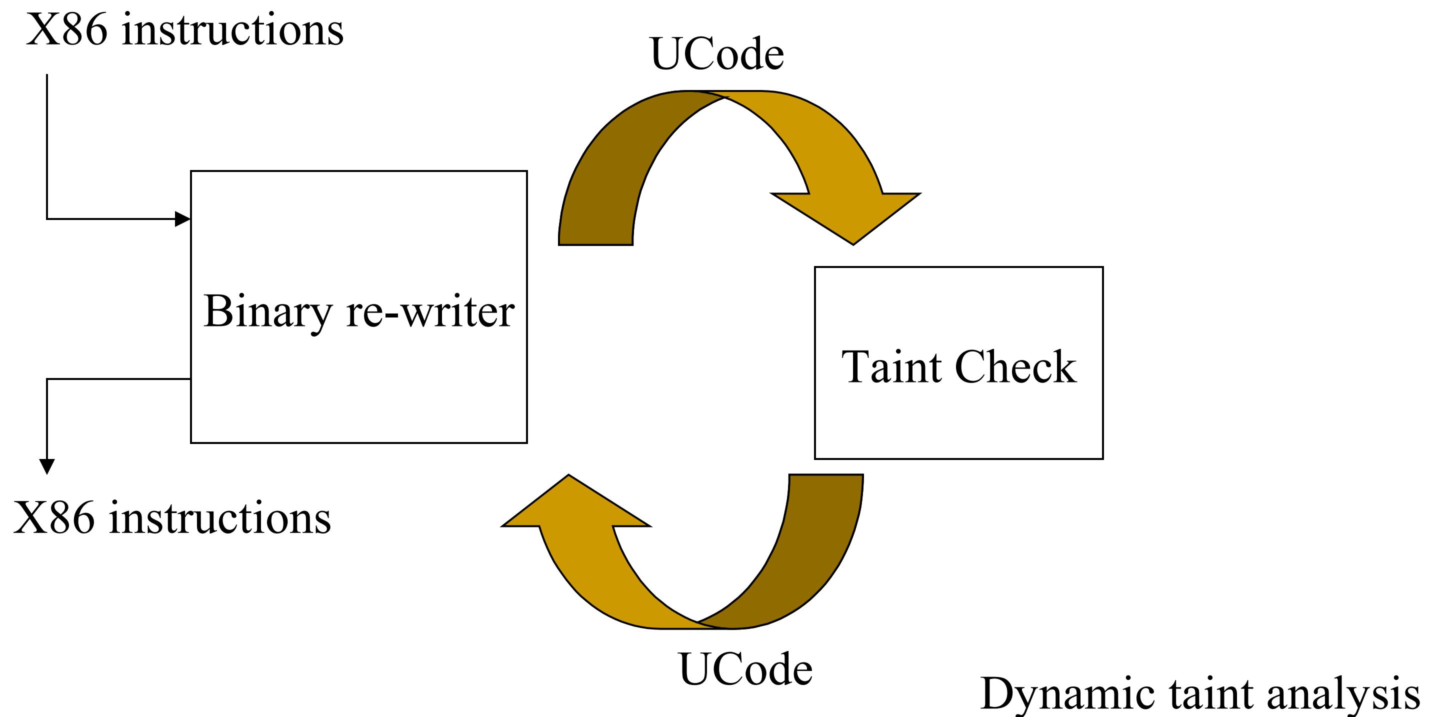

Dynamic taint analysis

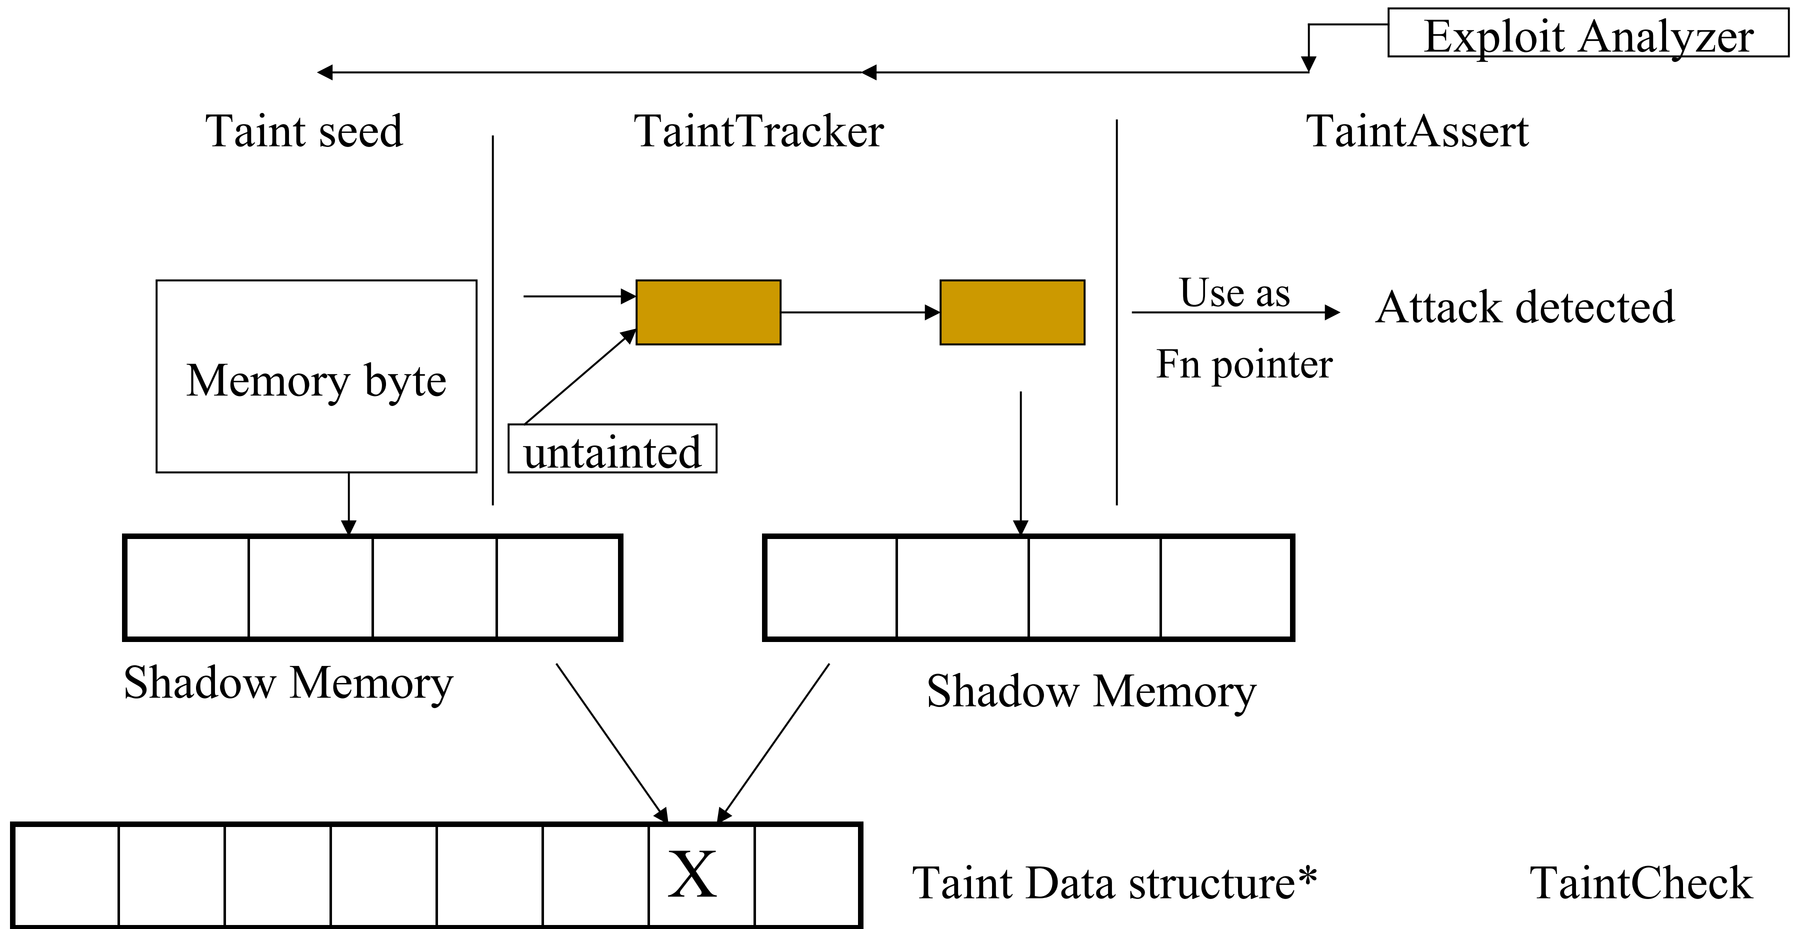

*TDS holds the system call number, a snapshot of the current stack, and a copy of the data that was written

Dynamic taint analysis

TaintSeed

- It marks any data from untrusted sources as “tainted”
    - Each byte of memory has a four-byte shadow memory that stores a pointer to a Taint data structure if that location is tainted, or a NULL pointer if it is not.

Memory is mapped to TDS

Dynamic taint analysis

TaintTracker

- It tracks each instruction that manipulates data in order to determine whether the result is tainted.
  - When the result of an instruction is tainted by one of the operands, TaintTracker sets the shadow memory of the result to point to the same Taint data structure as the tainted operand.

Memory is mapped to TDS

Result is mapped to TDS

Dynamic taint analysis

TaintAssert

- It checks whether tainted data is used in ways that its policy defines as illegitimate.

Dynamic taint analysis

Exploit Analyzer

- The Exploit Analyzer can provide useful information about how the exploit happened, and what the exploit attempts to do.

Dynamic taint analysis

Types of attacks detected by TaintCheck are

- Overwrite attack
  - jump targets (such as return addresses, function pointers, and function pointer offsets), whether altered to point to existing code (existing code attack) or injected code (code injection attack).
- Format string attacks
  - an attacker provides a malicious format string to trick the program into leaking data or into writing an attacker-chosen value to an attacker-chosen memory address.
    - E.g.. use of %s and %x format tokens to print data from the stack or possibly other locations in memory.

Dynamic taint analysis

Why to use TaintCheck ?

- **Does not require source code or specially compiled binaries.**
- **Reliably detects most overwrite attacks.**
- **Has no known false positives.**
- **Enables automatic semantic analysis based signature generation.**

Evaluation

False Negatives

- A false negative occurs if an attacker can cause sensitive data to take on a value without that data becoming tainted.
  - E.g. if (x == 0)y = 0; else if (x == 1) y = 1; ...
- If values are copied from hard-coded literals, rather than arithmetically derived from the input.
  - IIS translates ASCII input into Unicode via a table
- If TaintCheck is configured to trust inputs that should not be trusted.
  - data from the network could be first written to a file on disk, and then read back into memory.

Evaluation

False Positives

- TaintCheck detects that tainted data is being used in an illegitimate way even when there is no attack taking place.
    - It indicates, there are vulnerabilities in the program
        - E.g. A program uses tainted data as a format string, but makes sure it does not use it in a malicious way.

Evaluation

Synthetic

- To detect
    - Overwritten return addresses
    - Overwritten function pointer
    - Format string vulnerability

Actual exploits

- ATPhttpd exploit (buffer overflow)
- Cfingerd exploit (format string vulnerability)
- Wu-ftpd exploit (format string vulnerability)

Evaluation

Program	Overwrite Method	Overwrite Target	Detected
ATPhttpd	buffer overflow	return address	✔
synthetic	buffer overflow	function pointer	✔
synthetic	buffer overflow	format string	✔
synthetic	format string	none (info leak)	✔
cfingerd	`syslog` format string	GOT entry	✔
wu-ftpd	`vsnprintf` format string	return address	✔

Evaluation

Performance

- CPU bound
  - a 2.00 GHz Pentium 4, and 512 MB of RAM, running RedHat 8.0. was used to compress bzip2(15mb)
    - » Normal runtime 8.2s
    - » Valgrind nullgrind skin runtime25.6s (3.1 times longer)
    - » Memcheck runtime 109s (13.3 times longer)
    - » TaintCheck runtime 305s (37.2 times longer)
- Short-lived
- Common case

Evaluation

Performance

- CPU bound
- Short-lived
  - Basic blocks are cached and hence the penalty is acceptable over long lived programs. For short lived programs it is still significantly large
    - » Normal runtime for Cfingerd was0.0222s
    - » Valgrind nullgrind skin runtime took 13 times longer
    - » Memcheck runtime took 32 times longer
    - » TaintCheck runtime took 13 times longer
- Common case

Evaluation

Performance

- CPU bound
- Short-lived
- Common case
  - For network services the latency experienced is due to network and/or disk I/O and the TaintCheck performance penalty should not be noticeable

Application

It is not practical to implement TaintCheck as a standalone due to the performance overhead

- TaintCheck enabled honeypots could use TaintCheck to monitor all of its network services
  - TaintCheck will verify the exploit and provide additional information about the detected attack
- TaintCheck with OS randomization
  - identify which request contained an attack and generate signature for the attack or blocking future requests from the user.
- TaintCheck in a distributed environment

Evaluation

Performance

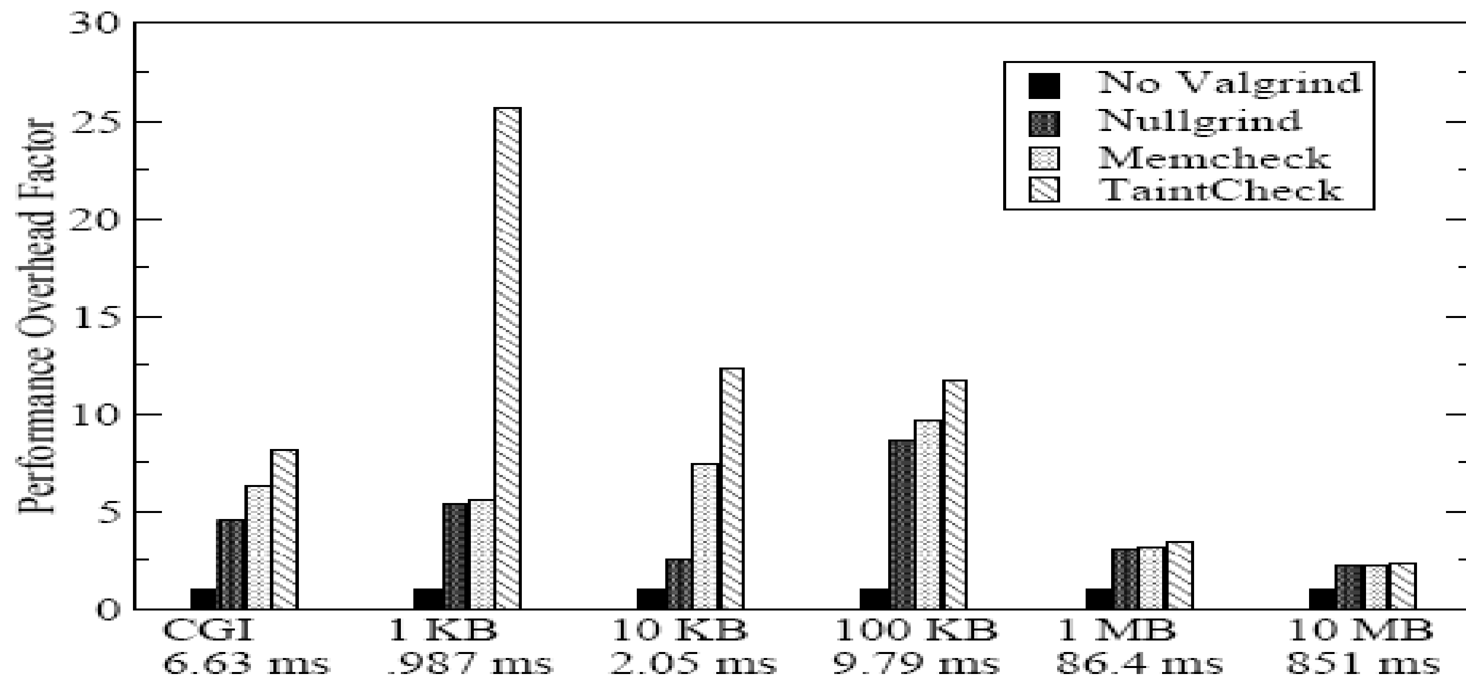

Application

Automatic semantic analysis based signature generation

- as it monitors how each byte of each attack payload is used by the vulnerable program at the processor-instruction level.