

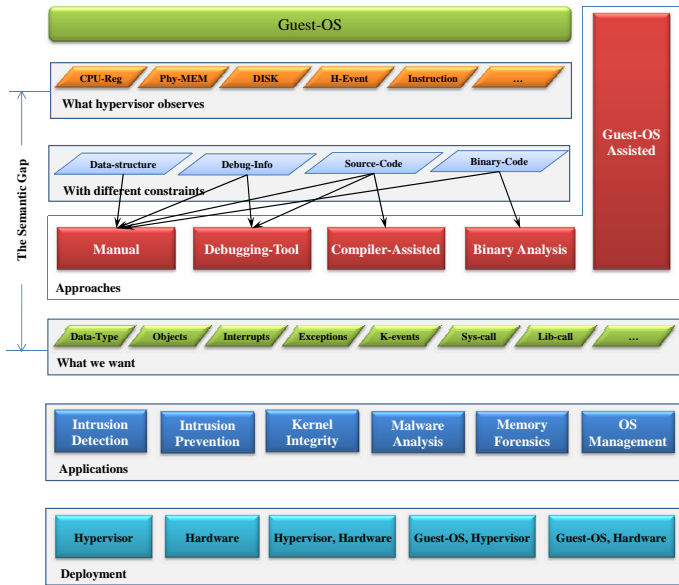
# All You Ever Wanted to Know About Virtual Machine Introspection: Applications

**Zhiqiang Lin**

Department of Computer Sciences  
The University of Texas at Dallas

August 24<sup>th</sup>, 2015

# The Road Map



# Outline

- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment

- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment

# Security Applications

## Security

### 1 Detection

- Kernel level
- User level
- Code modification detection
- Data modification detection
- Forensic analysis (including malware analysis)

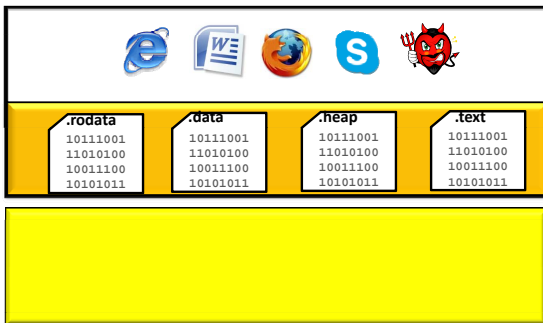
### 2 Prevention

- Kernel level
- User level
- Code protection
- Data protection

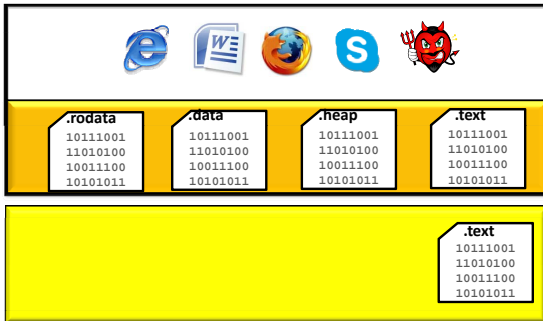
### 3 Recovery

- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment

# Detection: Kernel Code and Read-only Data

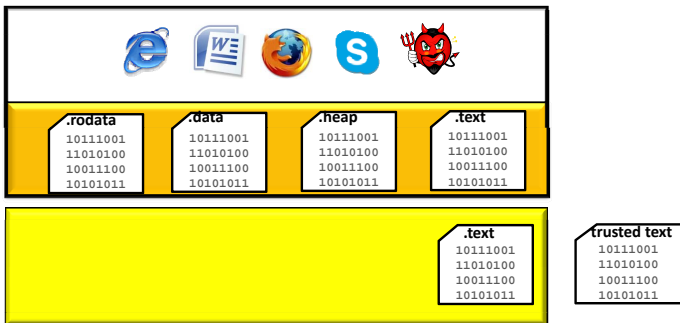


# Detection: Kernel Code and Read-only Data

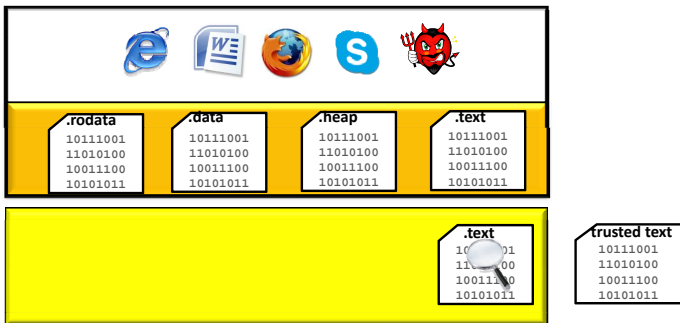




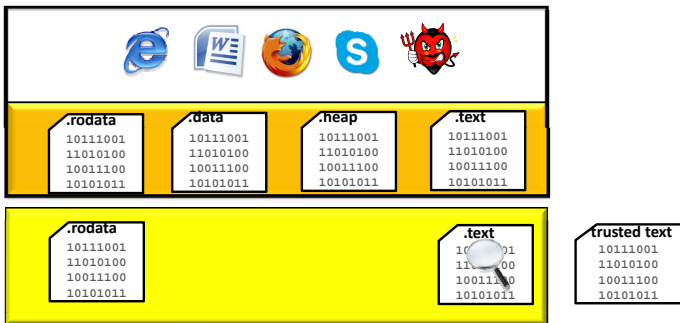
# Detection: Kernel Code and Read-only Data



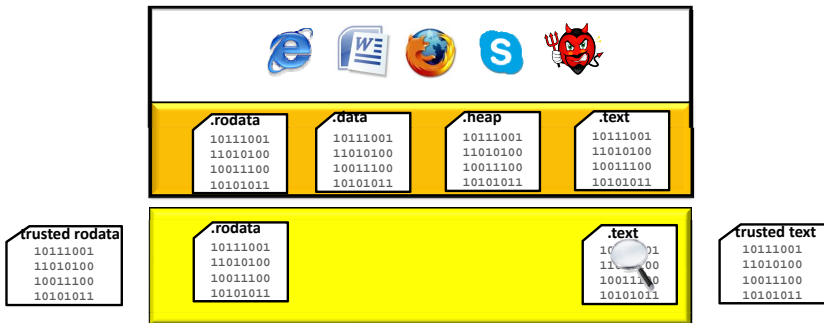
# Detection: Kernel Code and Read-only Data



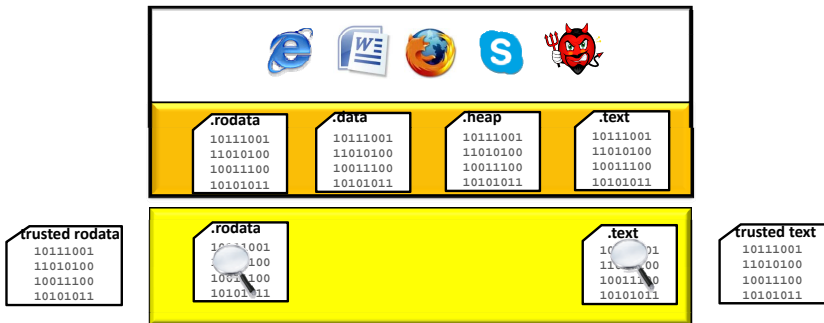
# Detection: Kernel Code and Read-only Data



# Detection: Kernel Code and Read-only Data



# Detection: Kernel Code and Read-only Data



# On Kernel Heap Data: Data Structure Traversal



**.heap**

10111001

11010100

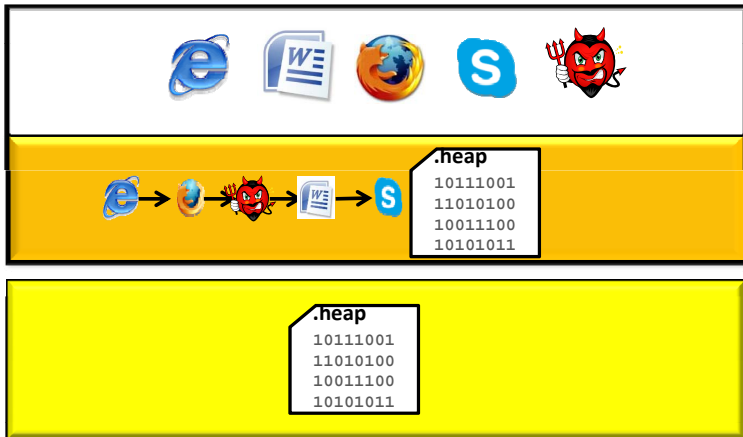
10011100

10101011

# On Kernel Heap Data: Data Structure Traversal

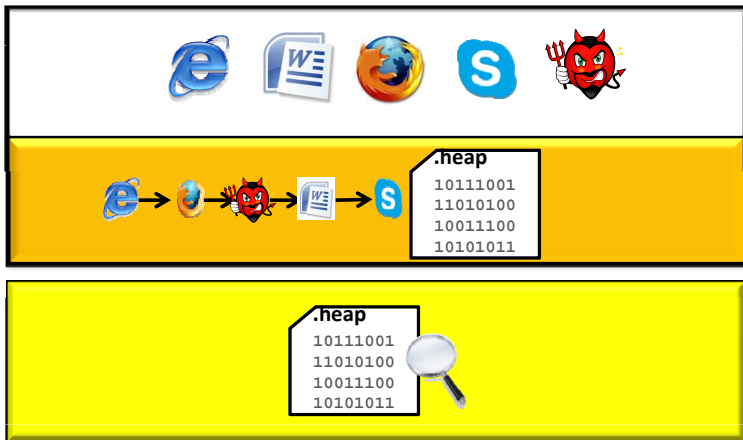


# On Kernel Heap Data: Data Structure Traversal

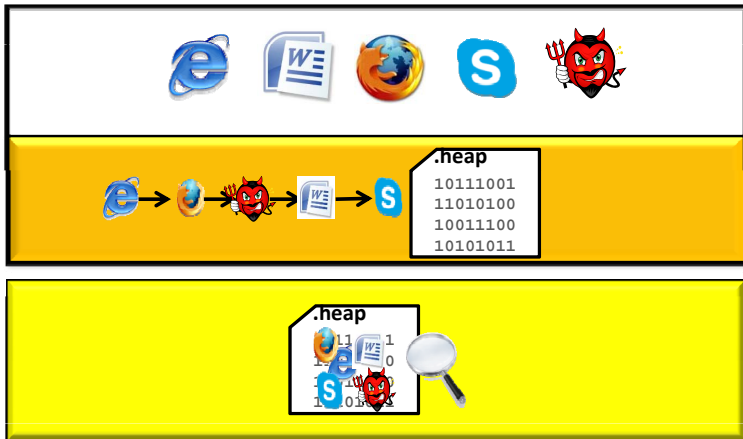




# On Kernel Heap Data: Data Structure Traversal



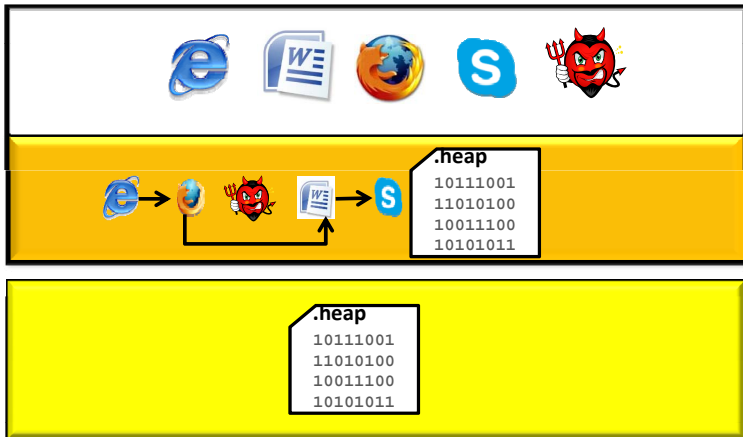
# On Kernel Heap Data: Data Structure Traversal



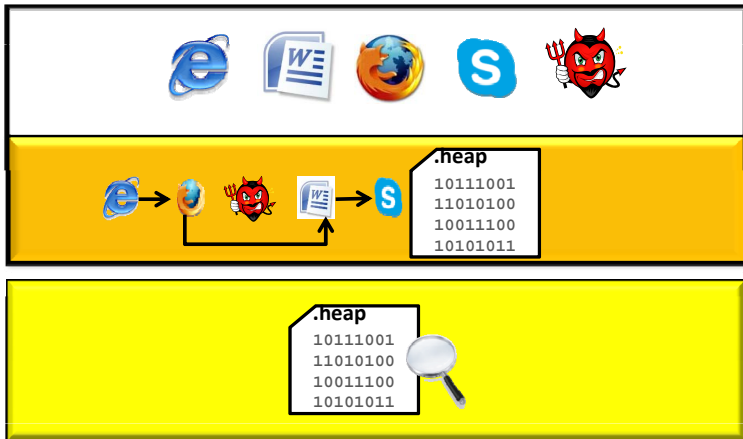
# On Kernel Heap Data: Data Instance Scanning



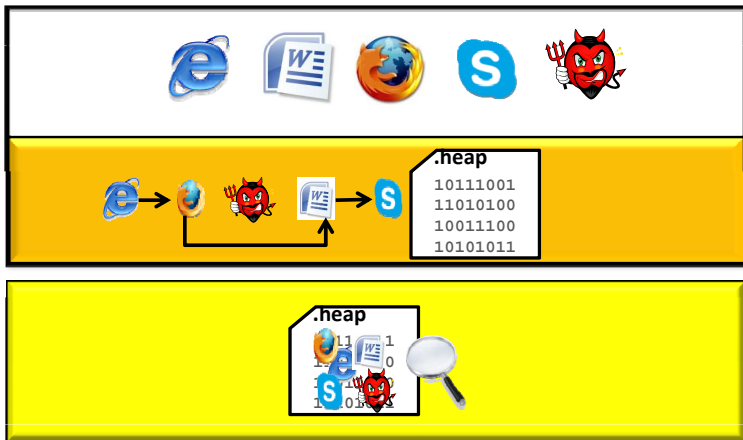
# On Kernel Heap Data: Data Instance Scanning



# On Kernel Heap Data: Data Instance Scanning



# On Kernel Heap Data: Data Instance Scanning



# On Kernel Heap Data: Using Invariant

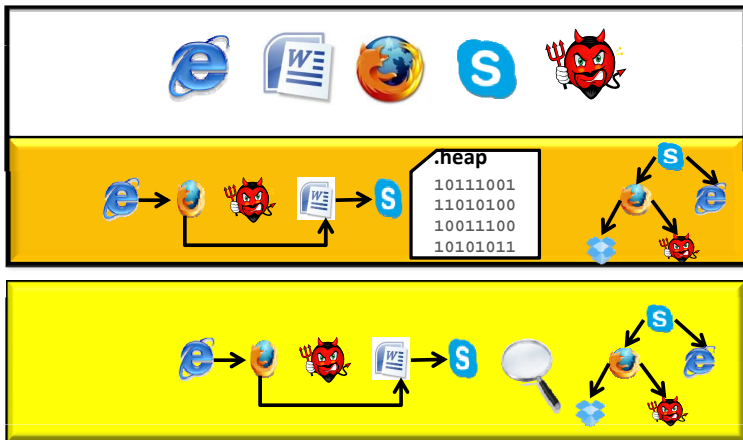


# On Kernel Heap Data: Using Invariant

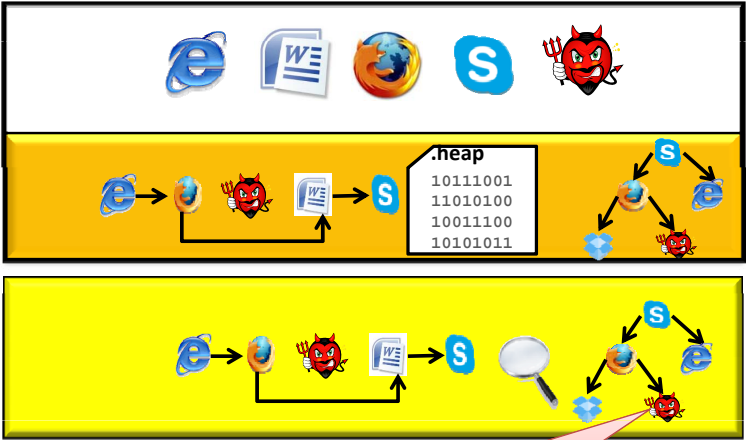




# On Kernel Heap Data: Using Invariant




# On Kernel Heap Data: Using Invariant




Invariant: Set of processes in Scheduler tree = Process list

# Using Data Out Grafting/Redirection [CCS'11,SP'12]

Secure VM (SVM)




Guest VM (GVM)




**Hypervisor**

# Using Data Out Grafting/Redirection [CCS'11,SP'12]

### Secure VM (SVM)




### Guest VM (GVM)



**Hypervisor**

# Using Data Out Grafting/Redirection [CCS'11,SP'12]

Secure VM (SVM)



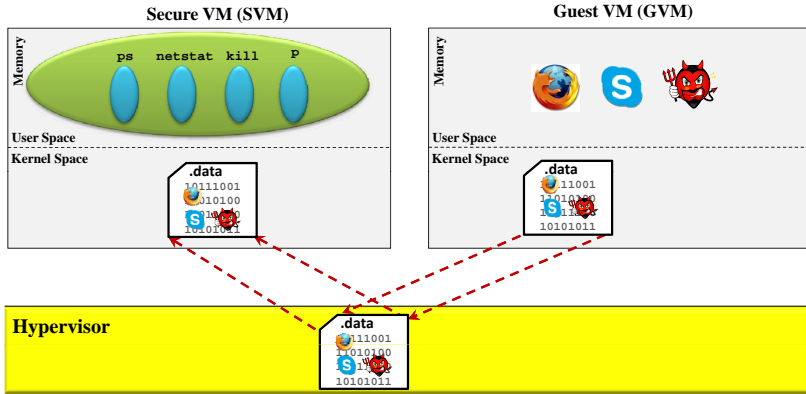
Guest VM (GVM)



Hypervisor




# Using Data Out Grafting/Redirection [CCS'11,SP'12]



# Using Code Implanting [DSN'10, ATC'14, DSN'14]

Secure VM (SVM)




Guest VM (GVM)



**Hypervisor**

# Using Code Implanting [DSN'10, ATC'14, DSN'14]

Secure VM (SVM)



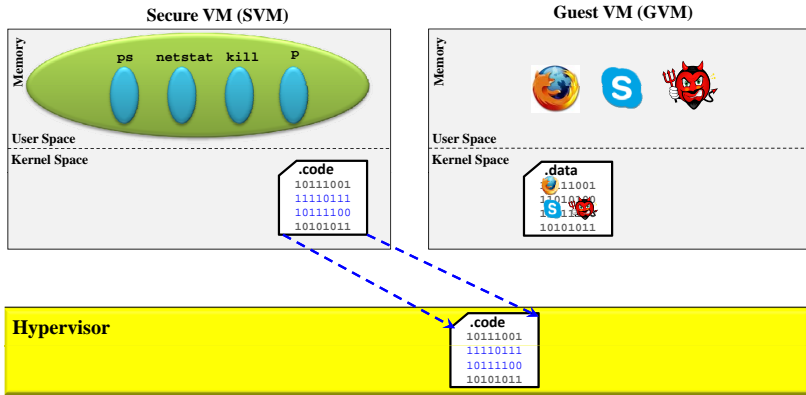
Guest VM (GVM)



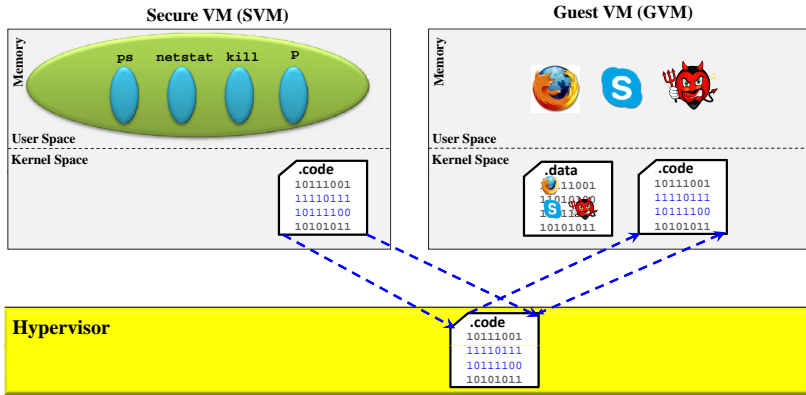
**Hypervisor**



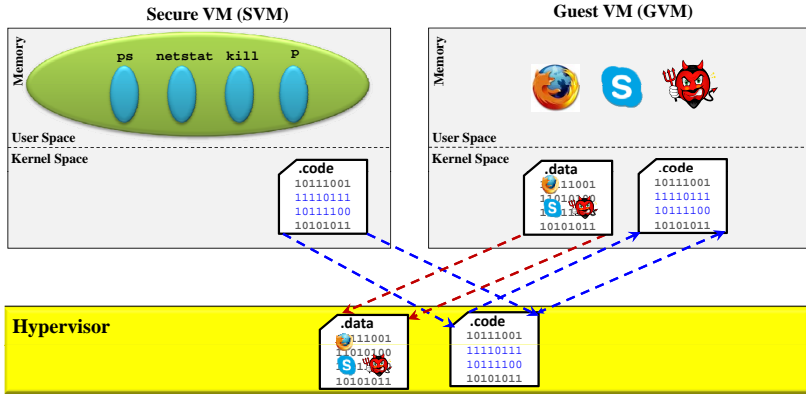
# Using Code Implanting [DSN'10, ATC'14, DSN'14]



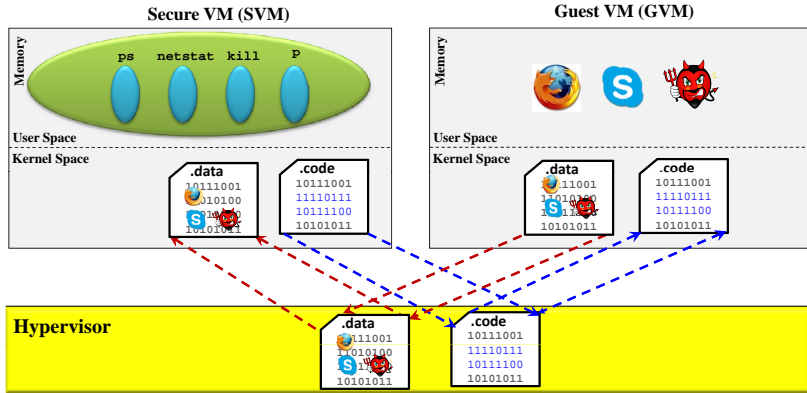
# Using Code Implanting [DSN'10, ATC'14, DSN'14]



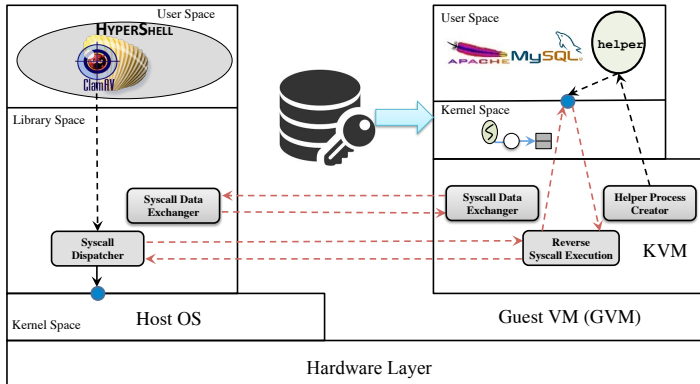
# Using Code Implanting [DSN'10, ATC'14, DSN'14]



# Using Code Implanting [DSN'10, ATC'14, DSN'14]



# Disk Introspection: FDE disk virus scanning [ATC'14]



# Disk Introspection: FDE disk virus scanning [ATC'14]



1. Encrypted by dm-crypt
2. 101,415 files
3. 1336.09 megabytes in size

# Disk Introspection: FDE disk virus scanning [ATC'14]




1. Encrypted by dm-crypt
2. 101,415 files
3. 1336.09 megabytes in size

Clamav successfully detect two viruses!!

- 1 Security Applications
  - Detection
  - **Prevention**
  - Recovery
- 2 Non Security Applications
- 3 Deployment

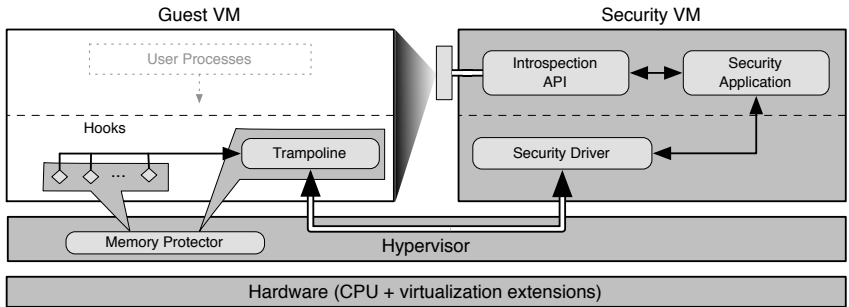


# Continues Monitoring [Payen et al. SP'08]

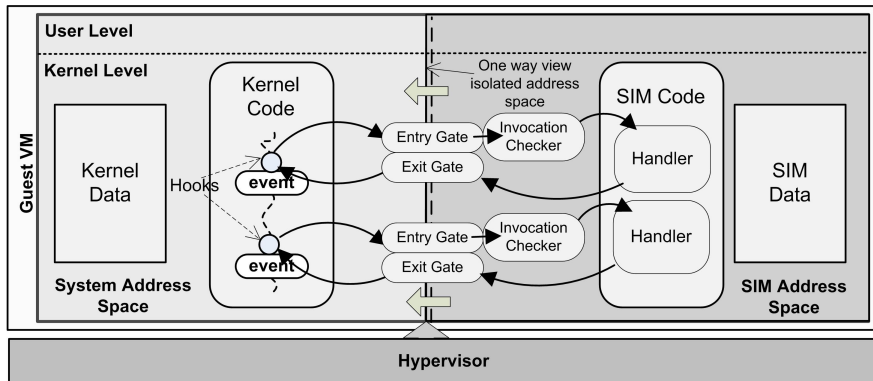


Formal model of secure active monitoring shown with potential attacks.

# Continues Monitoring [Payen et al. SP'08]

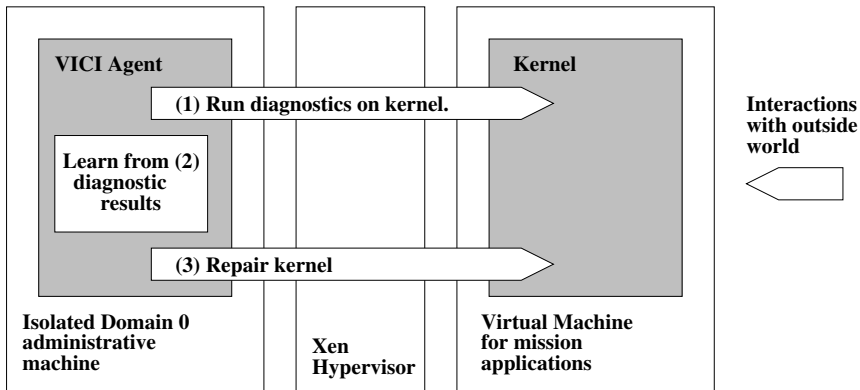


# In-VM Monitoring w/ Hardware [Sharif et al. CCS'09]



- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment

# VMI based Attack Repair [Fraser et al. ACSAC'08]



# Out-of-Box Attack Recovery, Repair

Rootkit	Targeted Function Pointer	Repaired?
adore-2.6	kernel global, heap object	✗
hookswrite	IDT table	✓
int3backdoor	IDT table	✓
kbdv3	syscall table	✓
kbeast-v1	syscall table, tcp4_seq_show	✓
mood-nt-2.3	syscall table	✓
override	syscall table	✓
phalanx-b6	syscall table, tcp4_seq_show	✓
rkit-1.01	syscall table	✓
rial	syscall table	✓
suckit-2	IDT table	✓
synapsys-0.4	syscall table	✓

Table : Rootkit Repairing with An Exterior [Fu and Lin, VEE'13] Tool.

- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment


# Non Applications

## Non Security


- 1 Virtual machine management
- 2 Process management
- 3 High performance computing
- 4 Autonomous computing
- 5 ...



# Out-of-VM Management: Writable VMI [ATC'14]

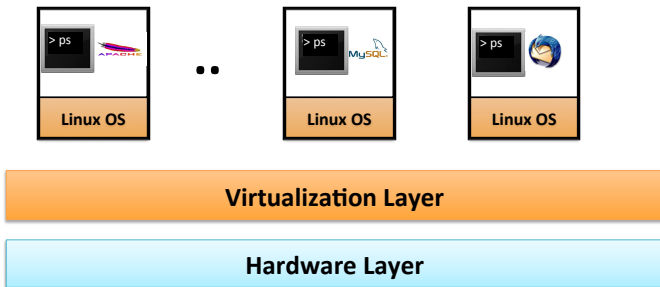


# Out-of-VM Management: Writable VMI [ATC'14]




- ### Advantages
- Only install the management utilities at hypervisor layer.
  - Automated, uniformed, and centralized management.

# In-VM Management: Existing Approaches



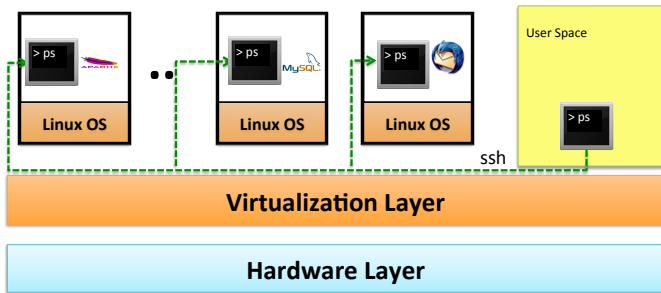
# In-VM Management: Existing Approaches



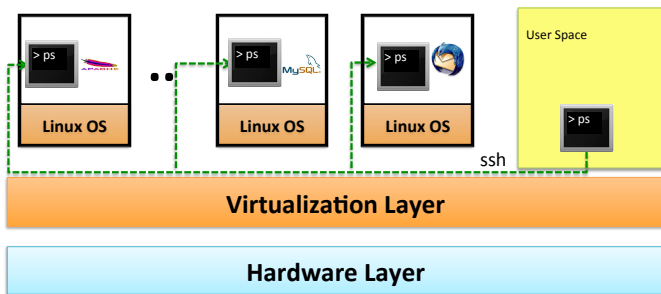
## Disadvantages

- Scattered, distributed
- Install, update, and execute in each VM

# In-VM Management: Existing Approaches



# In-VM Management: Existing Approaches



## Disadvantages

- Requiring the (admin) login password.
- Requiring install the management utilities in each VM.

# Performance Impact: HyperShell [ATC'14]

Process	S	B(ms)	D(ms)	T(X)										
ps	✗	1.33	5.42	4.08	date	✗	0.11	0.12	1.09	mkdir	✓	0.10	0.19	1.90
pidstat	✗	1.95	7.56	3.88	w	✗	0.95	6.62	6.97	mkfifo	✓	0.10	0.19	1.90
nice	✓	0.07	0.11	1.57	hostname	✓	0.04	0.06	1.50	mknod	✓	0.10	0.19	1.90
getpid	✓	0.01	0.02	2.00	groups	✓	0.21	0.62	2.95	mv	✓	0.15	0.31	2.07
mpstat	✗	0.29	0.66	2.28	hostid	✓	0.16	0.56	3.50	rm	✓	0.08	0.15	1.88
pstree	✗	0.69	6.03	8.74	locale	✓	0.09	0.17	1.89	od	✓	0.12	0.35	2.92
chrt	✓	0.11	0.16	1.45	getconf	✓	0.09	0.34	3.78	cat	✓	0.07	0.18	2.57
renice	✓	0.11	0.18	1.64	<b>System Utils</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	link	✓	0.07	0.13	1.86
top	✗	504.92	510.85	1.01	uptime	✗	0.07	0.47	6.71	comm	✓	0.08	0.22	2.75
nproc	✓	0.07	0.26	3.71	sysctl	✓	8.5	42.72	5.03	shred	✗	0.72	0.92	1.28
sleep	✓	1.27	1.28	1.01	arch	✓	0.07	0.11	1.57	truncate	✓	0.07	0.26	3.71
pgrep	✓	0.89	4.72	5.30	dmesg	✓	0.38	0.51	1.34	head	✓	0.07	0.15	2.14
pkill	✓	0.87	4.33	4.98	lscpu	✓	0.26	1.21	4.65	vdir	✓	0.63	3.95	6.27
snice	✓	0.17	0.65	3.82	mcookie	✗	0.29	0.49	1.69	nl	✓	0.08	0.17	2.13
echo	✓	0.07	0.09	1.29	<b>Disk/Devices</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	tail	✓	0.08	0.20	2.50
pwdx	✓	0.05	0.07	1.40	blkid	✓	0.14	0.61	4.36	namei	✓	0.07	0.13	1.86
pmap	✓	0.16	0.36	2.25	badblocks	✓	0.35	0.44	1.26	whereis	✓	2.05	4.86	2.37
kill	✓	0.01	0.04	4.00	lspci	✓	31.40	36.52	1.16	stat	✓	0.27	0.78	2.89
killall	✓	0.62	3.03	4.89	iostat	✓	0.45	1.04	2.31	readlink	✓	0.07	0.12	1.71
<b>Memory</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	du	✓	0.11	0.53	4.82	unlink	✓	0.07	0.13	1.86
free	✗	0.04	0.08	2.00	df	✓	0.16	0.35	2.19	cut	✓	0.08	0.17	2.13
vmstat	✗	0.19	0.33	1.74	<b>Filesystem</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	dir	✓	0.07	0.20	2.86
slabtop	✗	0.22	0.36	1.64	sync	✓	8.07	6.53	0.81	mktemp	✓	0.09	0.18	2.00
<b>Modules</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	getcap	✓	0.04	0.08	2.00	rmdir	✓	0.07	0.13	1.86
rmmod	✓	0.51	3.14	6.16	lsuf	✓	3.31	6.12	1.85	ptx	✓	0.12	0.45	3.75
modinfo	✓	0.48	1.54	3.21	pwd	✓	0.07	0.11	1.57	chcon	✓	0.06	0.12	2.00
lsmod	✓	0.10	0.17	1.70	<b>Files</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	ifconfig	✗	0.32	1.15	3.59
<b>Environment</b>	<b>S</b>	<b>B(ms)</b>	<b>D(ms)</b>	<b>T(X)</b>	chgrp	✓	0.19	0.47	2.47	ip	✓	0.10	0.20	2.00
who	✓	0.14	0.72	5.14	chmod	✓	0.07	0.14	2.00	route	✓	138.65	150.32	1.08
env	✓	0.07	0.11	1.57	chown	✓	0.19	0.47	2.47	ipmaddr	✓	0.13	0.34	2.62
printenv	✓	0.07	0.1	1.43	cp	✓	0.11	0.27	2.45	iptunnel	✓	0.09	0.29	3.22
whoami	✓	0.19	0.45	2.37	uniq	✓	0.09	0.35	3.89	nameif	✓	0.10	0.21	2.10
stty	✓	0.11	0.46	4.18	file	✓	0.87	1.72	1.98	netstat	✗	0.25	0.37	1.48
users	✓	0.09	0.53	5.89	find	✓	0.20	0.58	2.90	arp	✓	0.14	0.24	1.71
uname	✓	0.09	0.11	1.22	grep	✓	0.35	2.14	6.11	ping	✗	15.02	18.2	1.21
id	✓	0.26	0.85	3.27	ln	✓	0.08	0.14	1.75	Avg.	-	7.27	8.45	2.73
					ls	✓	0.14	0.27	1.93					

- 1 Security Applications
  - Detection
  - Prevention
  - Recovery
- 2 Non Security Applications
- 3 Deployment



# Deployment

- 1 Bare metal (e.g., Xen, vSphere, Hyper-V)
- 2 Hosted, Native Hypervisor (e.g., KVM)
- 3 Hosted, Emulation Hypervisor (e.g., QEMU)
- 4 Extra Hardware (e.g., PCI device in Copilot)



# Deployment Comparison

<b>Metric</b>	<b>Definition</b>
Flexibility	How many constraints are imposed on the monitor
Security	How well the deployment type provides for security coverage
Invisibility	How difficult the presence of the monitor is to detect from within the VM
Speed	How much system slowdown occurs compared to no monitor running
Space	How much storage capability the deployment type possesses

**Table :** Definitions of the metrics used to compare out-of-VM monitor deployment types.

# Deployment Comparison

Approaches	Flexibility	Security	Invisibility	Speed	Space
Bare Metal	●	◐	◑	●	●
Hosted, Native Hypervisor	●	◐	◑	●	●
Hosted, Emulation Hypervisor	●	◐	◑	○	●
Extra Hardware	○	●	●	●	○

**Table :** Comparison between different out-of-VM monitor deployment types. Note that symbol ○ denotes a low degree for that comparison item, ◐ denotes a medium degree, and ● denotes a high degree.