SecureSwitch: BIOS-Assisted Isolation and Switch between Trusted and Untrusted Commodity OSes

Kun Sun, Jiang Wang, Fengwei Zhang, Angelos Stavrou
Center for Secure Information Systems
George Mason University

GEORGE MASON UNIVERSITY

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

Introduction

- Segregate personal communication and entertainment from business and government operations
  - In untrusted environment: Web browser, online gaming, social web portal
  - In trusted environment: Online banking, shopping, business email
- Trusted computing environment
  - Separate/isolate the trusted environment from the untrusted environment
  - Prevent data leakage even if the host has been previously infected with malware in the untrusted environment.
  - Enables secure transactions on demand with low context switching latency

GEORGE MASON UNIVERSITY

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

Lampson Red/Green System Separation

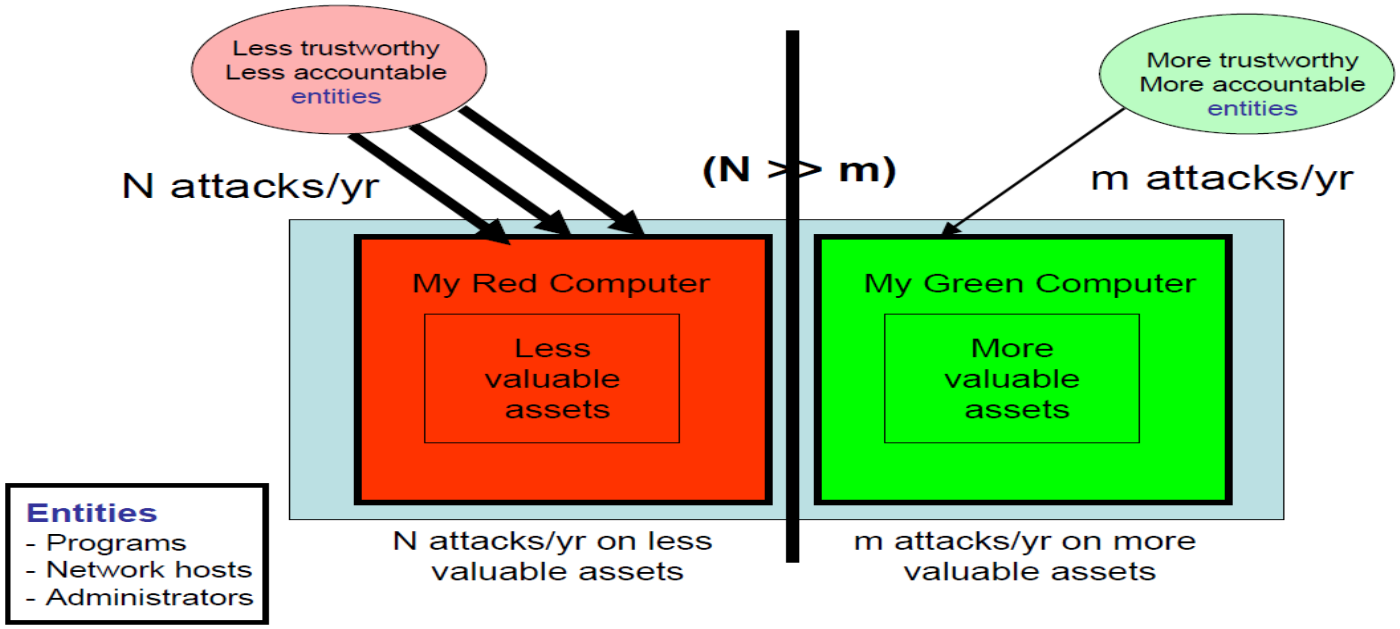

Red/Green system: Policy + Isolation + Accountability +Freedom

* Butler Lampson, Accountability and Freedom Slides, Microsoft, September 2005

- Our work focuses on isolation between two environments
  - Not on how the users decide what applications go into each OS
  - Not on how to give the user control over data exchanges between the two environments

GEORGE MASON UNIVERSITY

Software-based Isolation Solution

	VMM-based	OS-based	Browser-based
Isolation Level	OS level	User / Process level	Applet level
Examples	Xen, VMware, QEMU, UML	FreeBSD Jail, Linux OpenVZ, Solaris Container	Adobe Flash, Java applets, Silverlight
Security concerns	VMM vulnerabilities* Covert Channel	VMM vulnerabilities OS vulnerabilities Privilege escalation	VMM vulnerabilities OS vulnerabilities Browser vulnerabilities

* From 1999 to 2009, 373 vulnerabilities affecting virtualization solutions.
--- "IBM X-Force 2010 Mid-year trend and risk report"

GEORGE MASON UNIVERSITY

Hardware-based Isolation Solution

	Multiple Computers	Multi-boot	VT-x / SVM (DRTM)
Isolation Level	Whole physical computer	OS level;	Instruction level;
Examples		Bootloader: LILO, Grub	Flicker [1]; TrustVisor [2];
Problems	Cost, inflexible	Long switching time	Software compatibility

Our work provides an BIOS-assistant OS level isolation

- without using any mutable software layer (e.g., hypervisor)
- no changes of the OS source code
- no data leak between two OS environments
- fast switching time, around 6 seconds

GEORGE MASON UNIVERSITY

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

ACPI Sleeping States

- Advanced Configuration and Power Interface (ACPI)
  - OS-directed configuration; Power/thermal management
  - Industrial standard widely supported
- Global System States
  - G0 --- Working (System Operational)
  - G1---Sleeping (CPU stopped)
  - G2 ---Soft Off
  - G3 ---Mechanical off (Physical off switch)
- Sleeping States in G1: S1 – S4
  - S3: also *called Standby, Suspend to RAM*
    - DRAM still maintained
  - S4: also called *Hibernation or Suspend to Disk*
    - DRAM not maintained
- Device Power States: D0 – D3
  - D0 - Fully-On
  - D3 -- Power off to device

BIOS, UEFI, and Coreboot

- Basic Input/Output System (BIOS)
  - Initializing hardware like processor, memory, chipset, hard disk, and other I/O devices.
  - Stored in non-volatile ROM chips.

- Unified Extensible Firmware Interface (UEFI)
  - Define a new software interface between OS and firmware.
  - Ease the development by switching into protected mode in early stage and writing code in C language.
  - Partially open source

- ***Coreboot*** (formerly as LinuxBIOS)
  - Similar functionality as UEFI
  - Open source
  - We use Coreboot V4

GEORGE MASON UNIVERSITY

DIMM Mask and DQS Setting

- A motherboard may have more than one Dual In-line Memory Module (DIMM) slot.
  - We assign one DIMM to one OS.

- BIOS uses “DIMM_MASK” variable to control which DIMMs to be enabled.
  - BIOS sets corresponding “data strobes”(DQS) parameters to enable DDR RAM memory access.
  - Require two sets of DQS setting to support the two DIMMs used by two OSes separately.

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

System Architecture

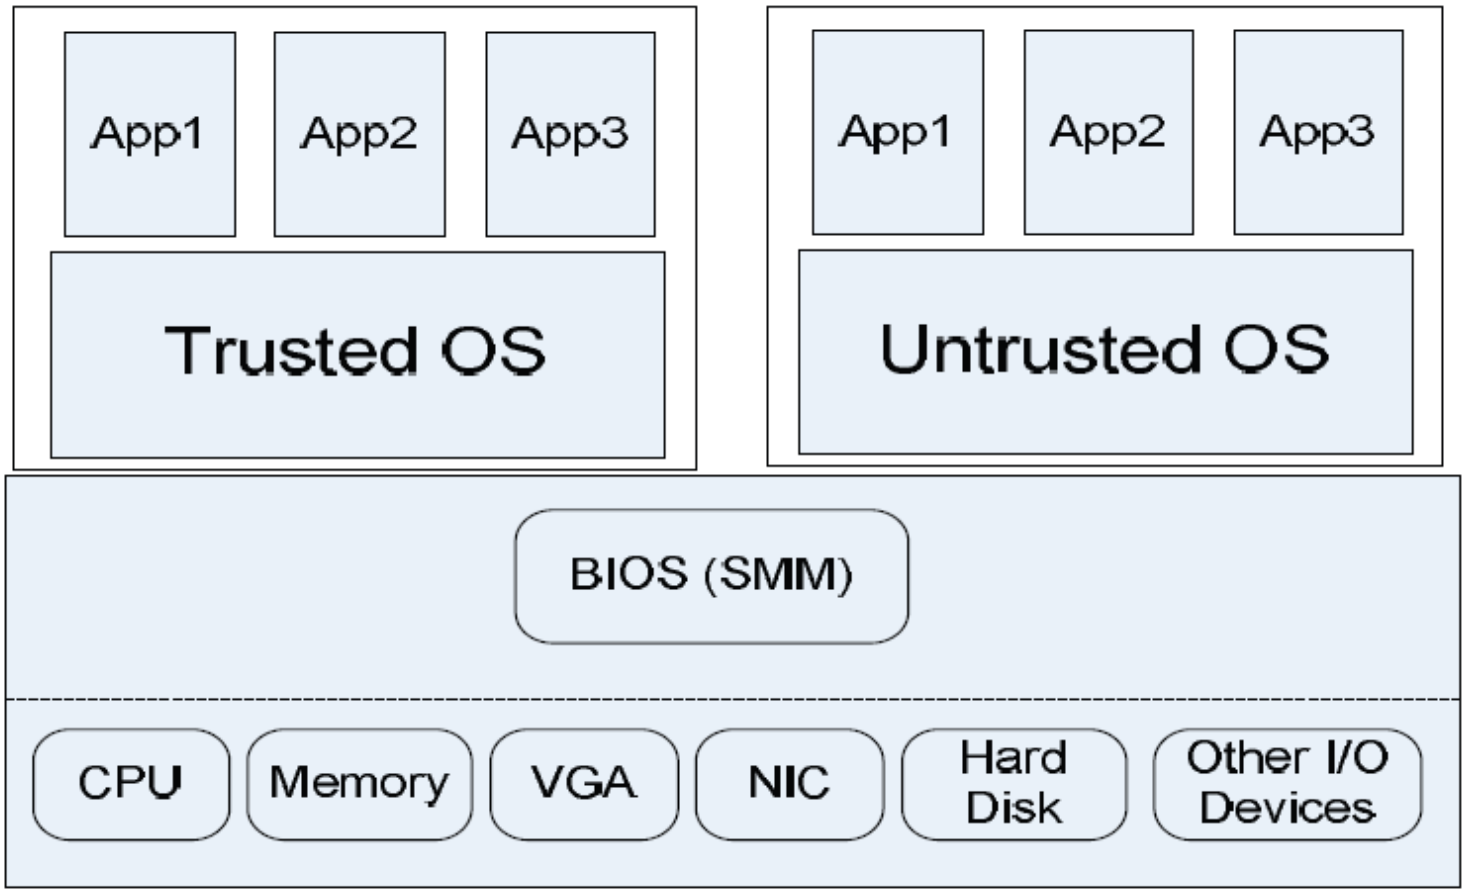

Attack Model

- Assumption
  - BIOS and option ROM on devices can be trusted.
  - No physical access to the protected machine
- Attacks from the untrusted OS
  - *Data exfiltration attacks*: stealing sensitive data from the trusted OS
  - *Spoofing Trusted OS attacks*: deceiving the user into a fake trusted OS to perform sensitive transactions.
  - *Cache-based side channel attacks*: extracting sensitive information
- Out of the scope
  - Denial of Service attacks
  - Network attacks on trusted OS
  - Malicious device firmware

Secure Switching State Machine

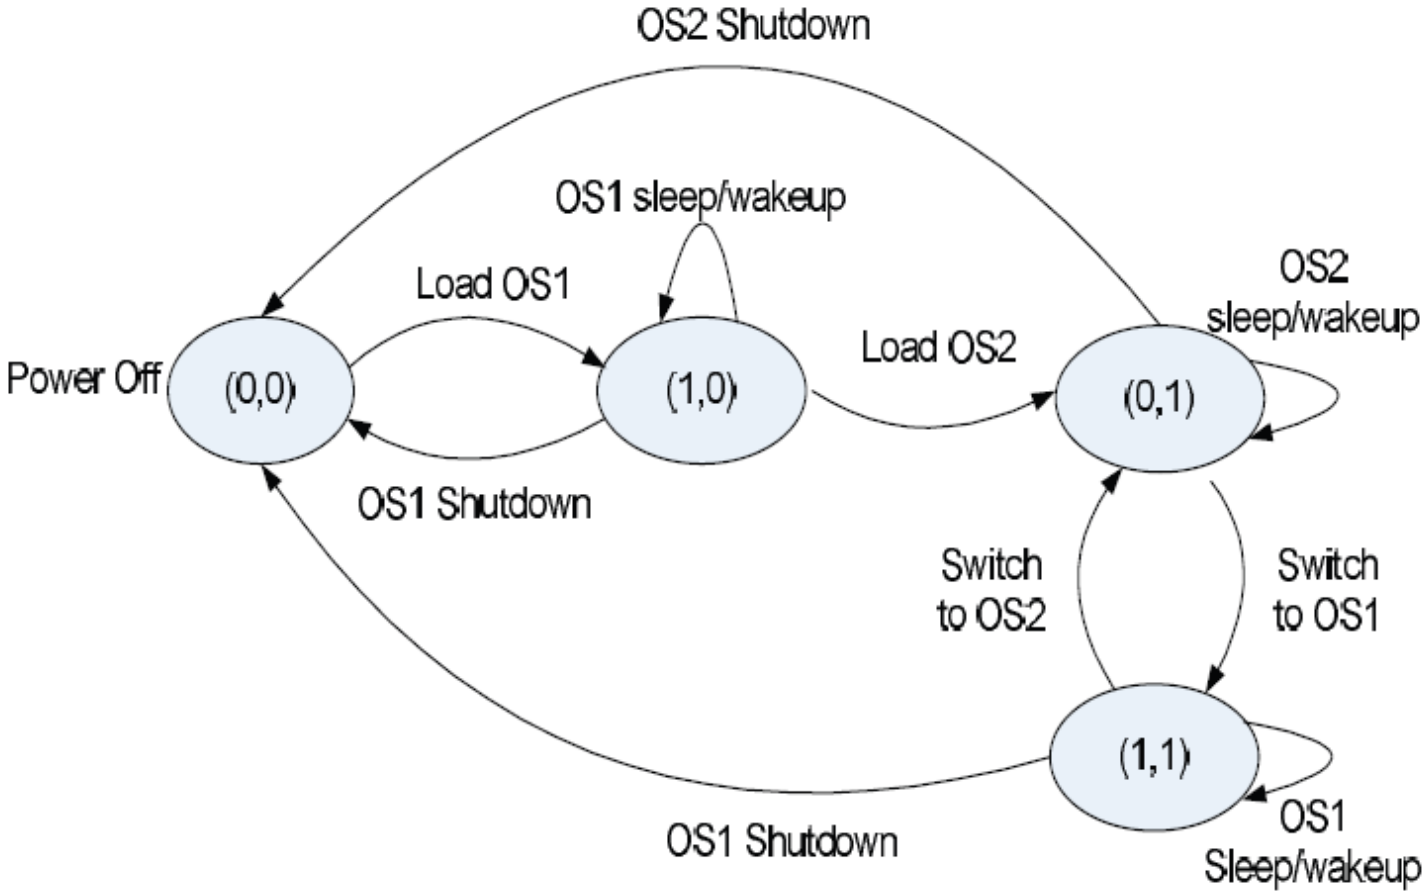

Trusted Path

- A trusted path assures users that they really are working with the operation system they intend to use.
- Prevent spoofing trusted OS attacks
  - Untrusted OS should be truly suspended.
    - Power LED lights up when system is powered on, and blinks in the sleep mode.
  - BIOS should be entered.
    - Press the power button.
  - Protecting system flags/variables
    - OS Flag: which OS should be woken next
    - Boot Flag: if untrusted OS has been loaded
    - Where to save these flags?

Boot Flag: in CMOS
OS Flag: physical jumper (e.g., pins in Parallel port connector)

Secure Switching Process

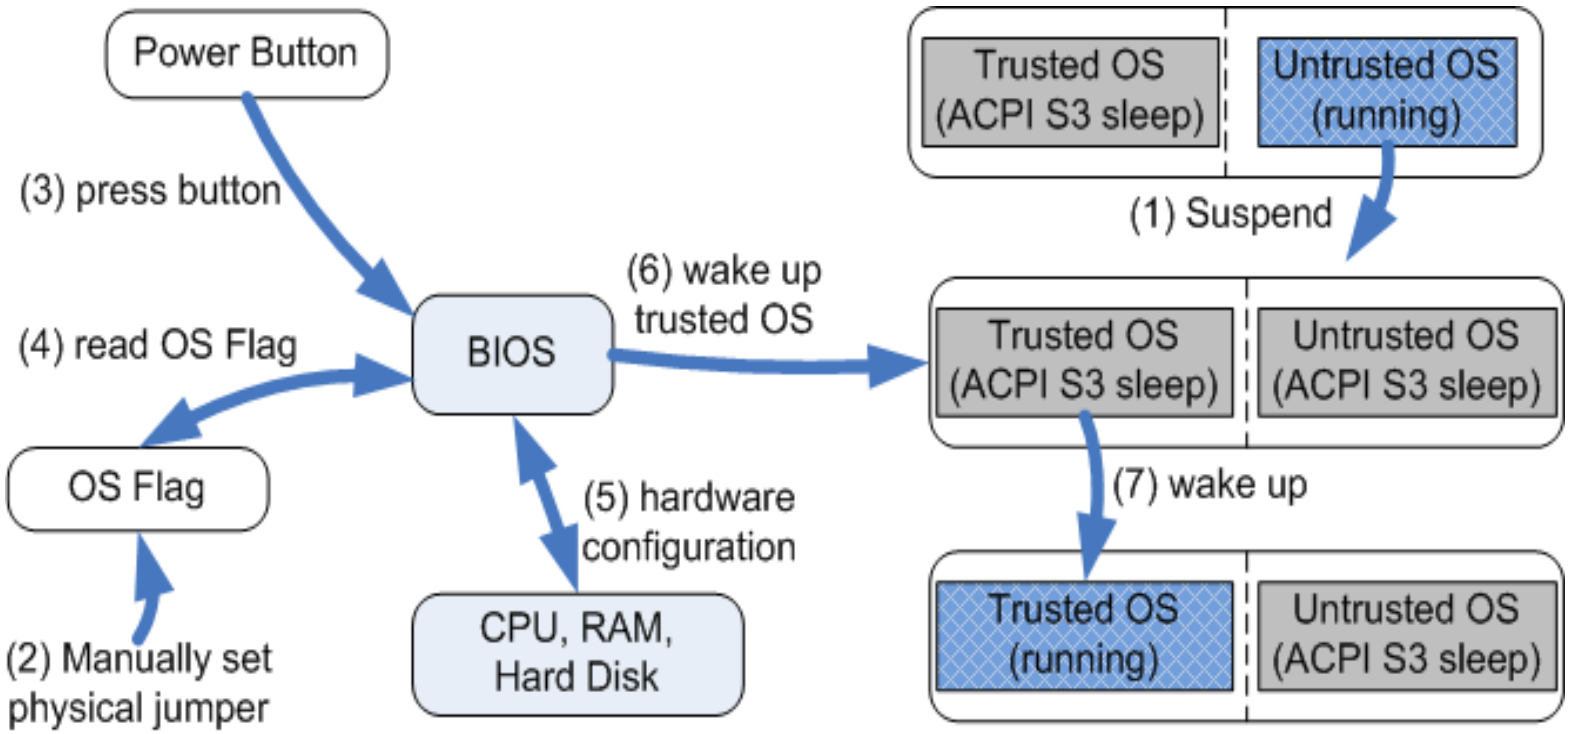

System Isolation

- CPU Isolation: two OSes never run concurrently.
- Memory Isolation: physical-level isolation
- Hard disk isolation: encrypted hard disk + RAM disk
- Other I/O isolation: clean the buffers and states in devices.

- **Isolation Mechanisms:**

	CPU	Memory	Hard Disk	VGA	NIC
OS with ACPI S3	√	√	√	√	√
BIOS		√	√		

GEORGE MASON UNIVERSITY

Physical-level Memory Isolation

- OS environments run in separate Dual In-line Memory Modules (DIMMs).
- BIOS only enables and reports one DIMM for each OS.
  - Two DQS settings for two OSes
  - "DIMM_MASK" controlled by the physical jumper.
  - When the "DIMM_MASK" conflicts with DQS setting, the system crashes
- Only the BIOS can initialize and enable the DIMMs; software cannot initialize or enable DIMMs after the system boots up

Hard Drive Isolation

- Hard disk encryption
  - Two hard disks, one for each OS.
  - Disk lock in ATA specification
- BIOS only enables one hard disk
  - Attacker may change the setting to enable all hard disks
  - We use SMM-based detection to check that the channel enable registers has not been changed
- RAM disk
  - For browser-based application, save a small amount of temporary data in the RAM

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

Prototype

- **Hardware**
  - *Motherboard*: ASUS M2V-MX_SE
  - *CPU*: AMD Sempron 64 LE-1300
  - *DDR2*: Kingston HyperX 1GB
  - *HDD*: Seagate 500GB

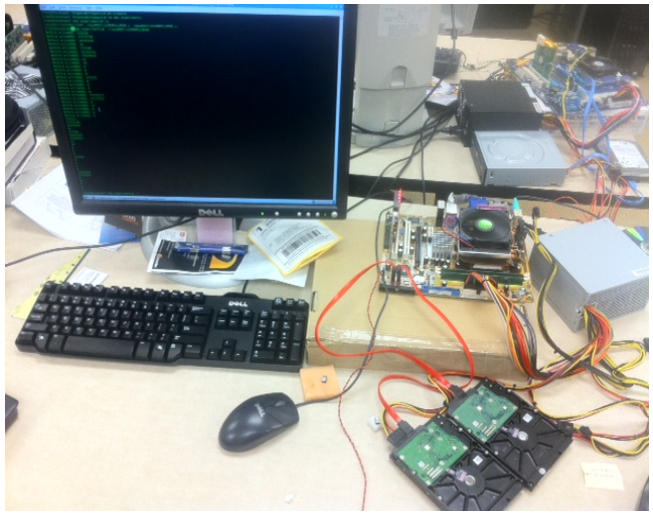

- **Software**
  - *BIOS*: Coreboot + SeaBIOS
  - *Trusted OS*: Linux (Centos 5.5)
  - *Untrusted OS*: Windows XP

GEORGE MASON UNIVERSITY

Performance Analysis

Table 1: Switching Time

Switching Operation	Secure Switch(s)
Windows XP Suspend	4.41
CentOS Wakeup	1.96
Total	6.37
CentOS Suspend	2.24
Windows XP Wakeup	2.79
Total	5.03

GEORGE MASON UNIVERSITY

Linux Suspend Time Breakdown

User Space : 1517.33 ms

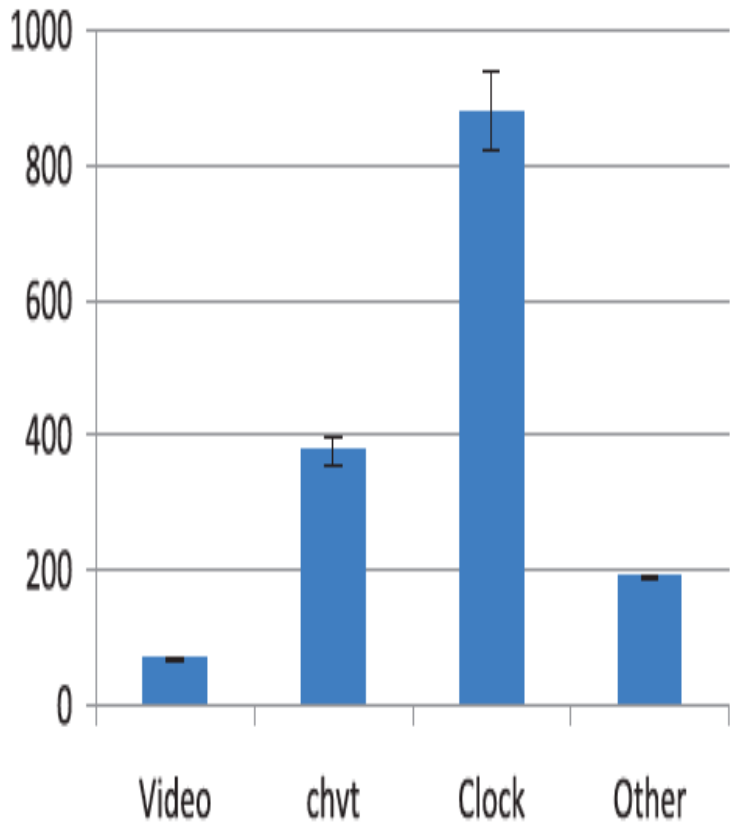

Kernel Space: 1590.14 ms

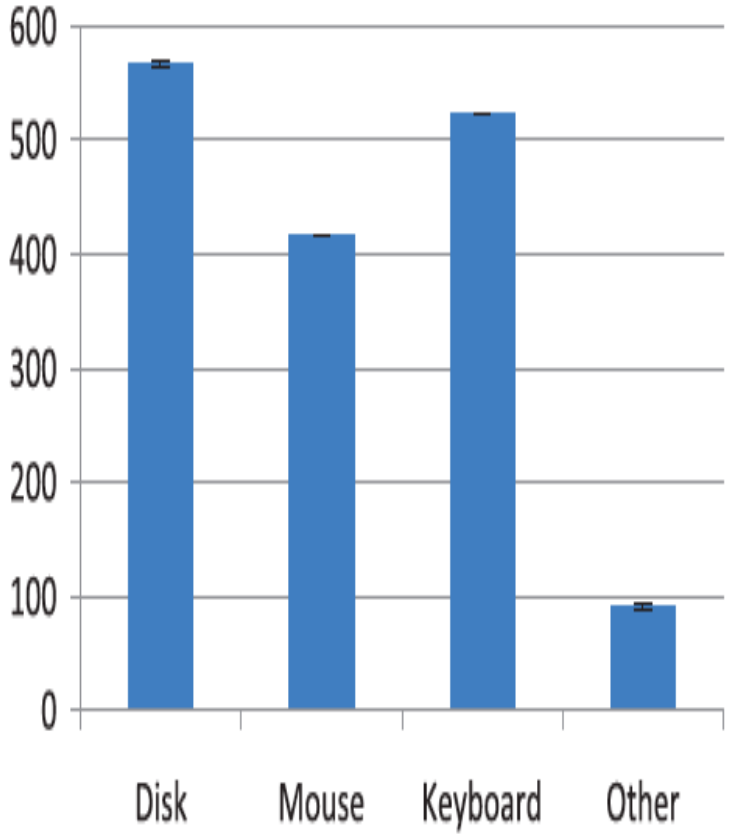

GEORGE MASON UNIVERSITY

Linux Wakeup Time Breakdown

User Space: 621.04 ms

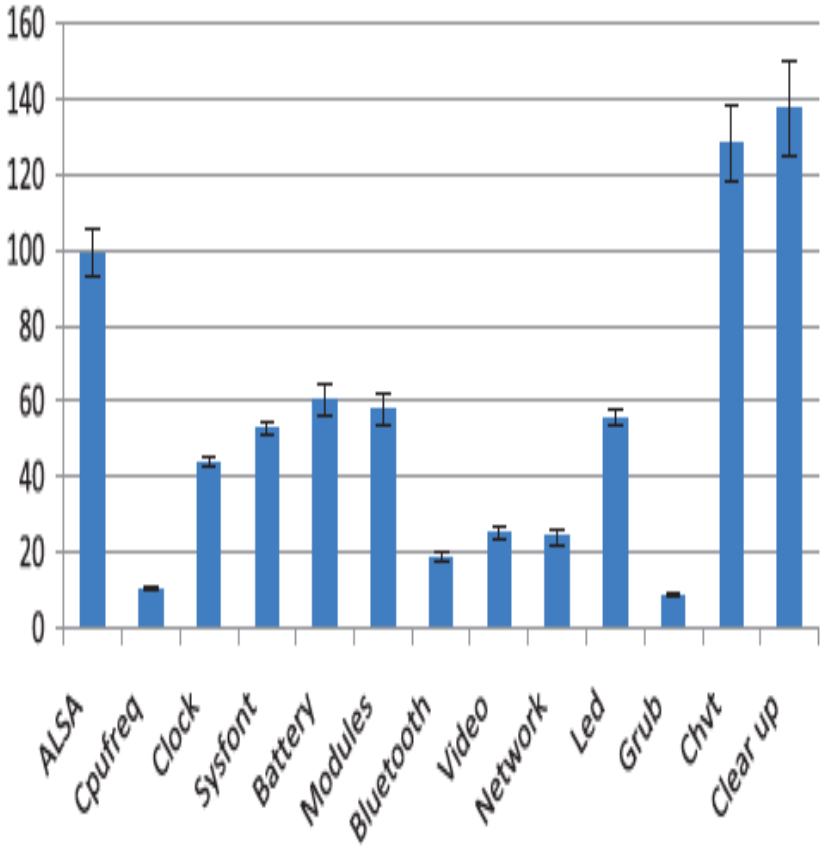

Kernel Space: 1537.22 ms

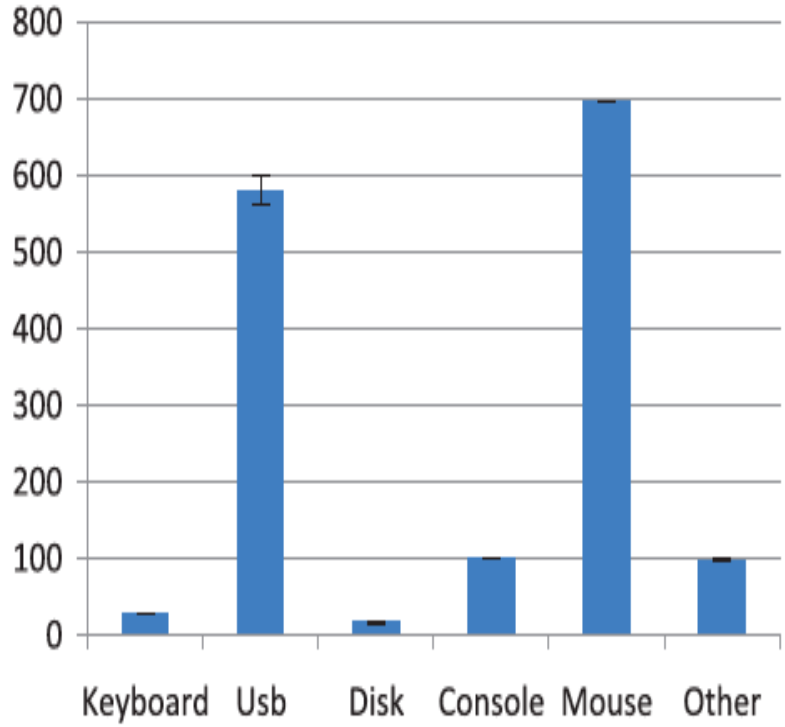

GEORGE MASON UNIVERSITY

Outline

- Introduction
- Related Works
- Background Knowledge
- System Architecture
- Experimental Results
- Discussion
- Summary

GEORGE MASON UNIVERSITY

Comparison with other systems

	SecureSwitch	Lockdown [41]	TrustVisor [29]	Flicker [28]	Xen [14]
Trusted Computing Base	BIOS	Hypervisor+BIOS	Hypervisor	250 LOC	Hypervisor
Switching Time (second)	≈6	40	<0.001	1	<0.001
Hardware Dependency	ACPI	ACPI+TPM+ VT-x/SVM	TPM+ VT-x/SVM	TPM+ VT-x/SVM	VT-x/SVM*
Software Compatibility	High	High	Low	Low	High
Memory Overhead	High	Low	Low	Low	Medium
OS Concurrency	No	No	Yes	No	Yes
Computation Overhead	Low	Medium	Medium	Low	Medium

* Xen requires VT-x/SVM to support full virtualization.

GEORGE MASON UNIVERSITY

Summary

- We develop a BIOS-based secure isolation and switching system to obtain a usable trusted workspace
  - Prevent data leakage
  - Without using hypervisor
  - No changes of OS source code
  - Low switching time

GEORGE MASON UNIVERSITY

Thank you.

- Questions?

GEORGE MASON UNIVERSITY

Reference

- [1]. J. McCune, B. Parno, A. Perrig, M. Reiter, and H. Isozaki. Flicker: An execution infrastructure for TCB minimization. In Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems 2008, pages 315–328. ACM, 2008.
- [2]. J. M. McCune, Y. Li, N. Qu, Z. Zhou, A. Datta, V. Gligor, and A. Perrig. TrustVisor: Efficient TCB reduction and attestation. In Proceedings of the IEEE Symposium on Security and Privacy, 2010.

GEORGE MASON UNIVERSITY