Toward Resilient Cloud Environment:

Case for Virtual Machine Introspection Using Hardware Architectural Invariants

Z. Kalbarczyk

C. Pham, C. Di Martino, R. Iyer

Coordinated Science Laboratory

Department of Electrical and Computer Engineering

University of Illinois at Urbana-Champaign

What Happens in an **Internet Minute?**

And **Future Growth** is **Staggering**

SOURCE: http://www.intel.com/content/www/us/en/communications/internet-minute-infographic.html

Magnitude of the Problems:
Five-Minute Snapshot of In-and-Out Traffic within NCSA

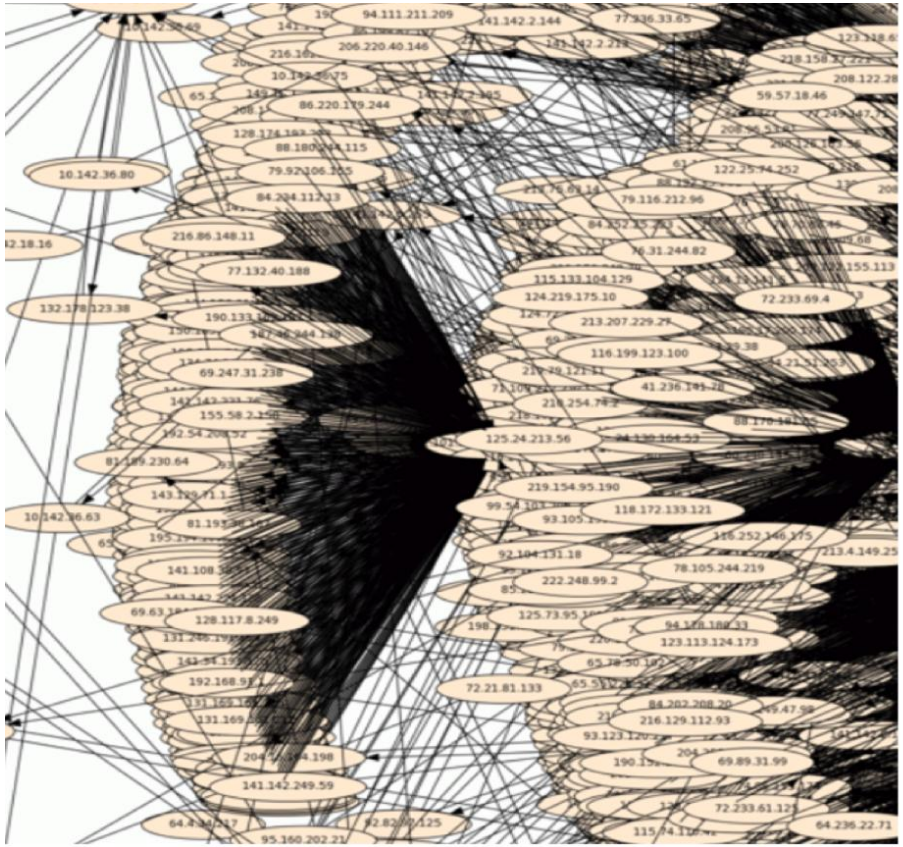

(a)

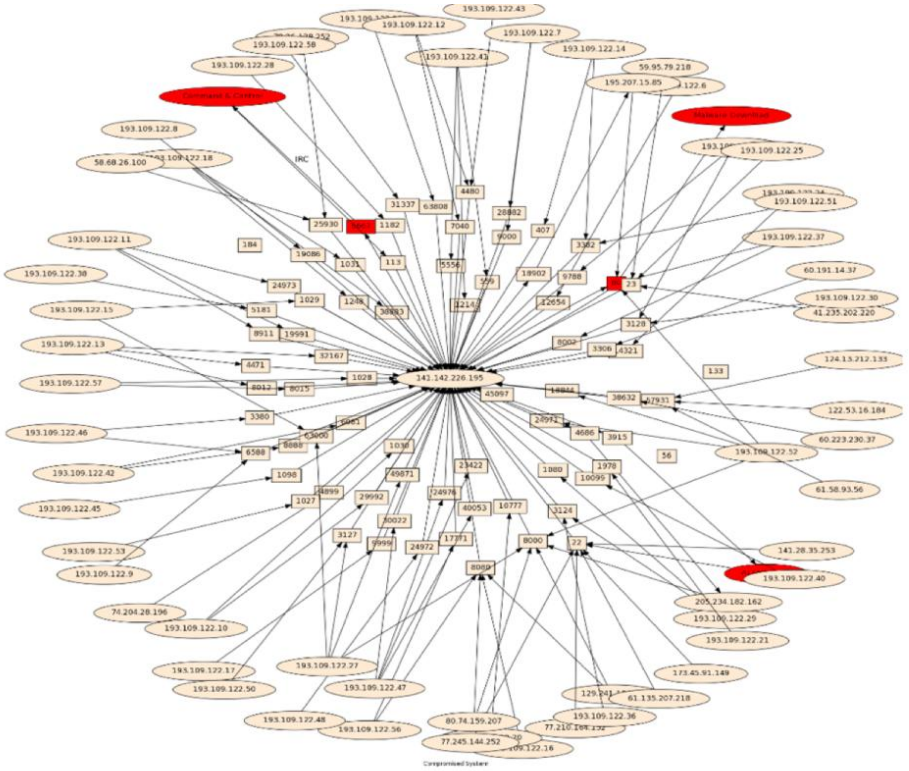

(b)

Cloud Computing - Growing Interest vs. Security Problems

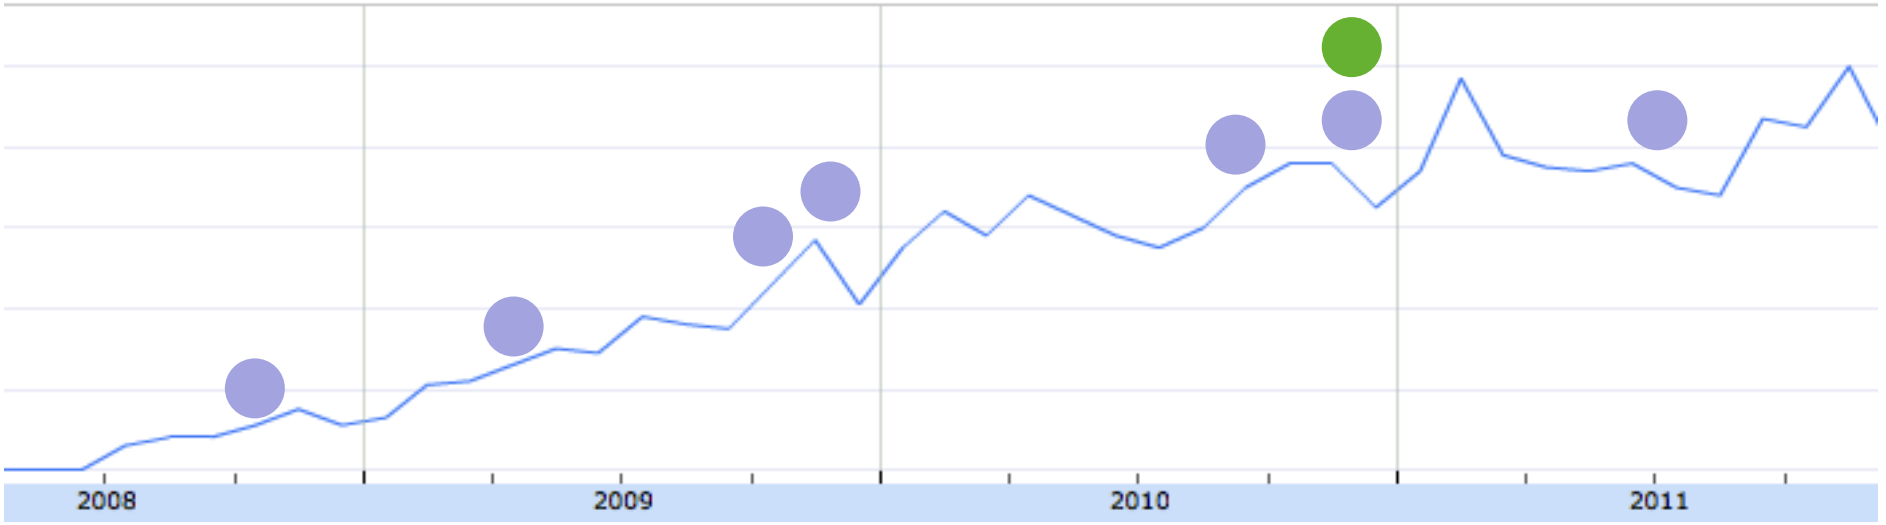

Jul'08 - Spammers set up mail spamming instances in the Amazon's EC2 cloud.

Apr'09 - Texas datacenters operations are suspended for FBI investigation.

Nov'09 - Side channel attack of Amazon's EC2 service.

Dec'09 - Zeus crime-ware using Amazon's EC2 as command and control server.

Sep'10 - Google Engineer Stalked Teens, Spied on Chats

Dec'10 - Microsoft BPOS cloud service hit with data breach

June'11 - Dropbox: Authentication Bug Left Cloud Storage Accounts Wide Open

Dec'10 - Anonymous hacker group failed to take down Amazon

Cloud Computing - Growing Interest vs. Growing Number of Outages

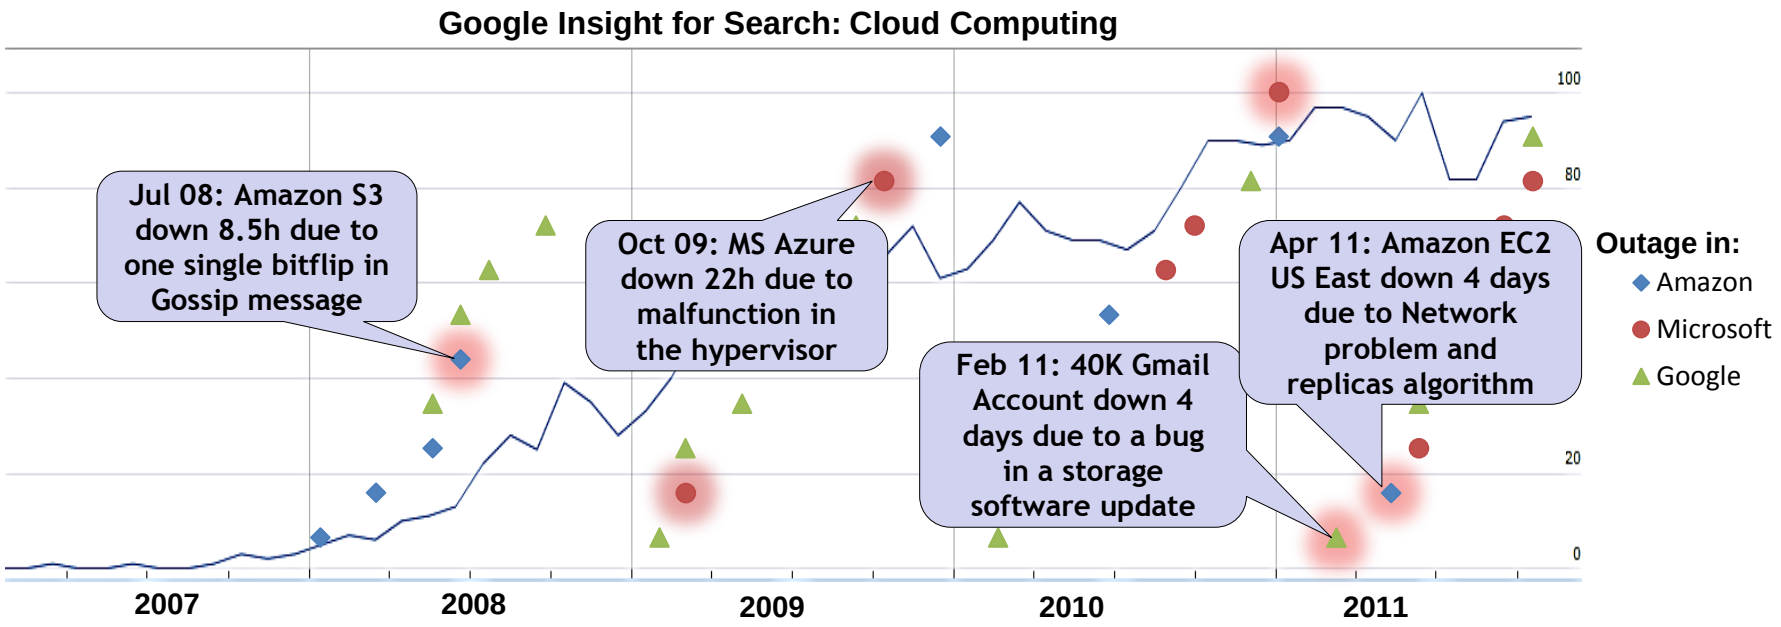

- Providing a higher level of availability and security is one of the biggest challenges of Cloud computing

Challenges in VM Monitoring

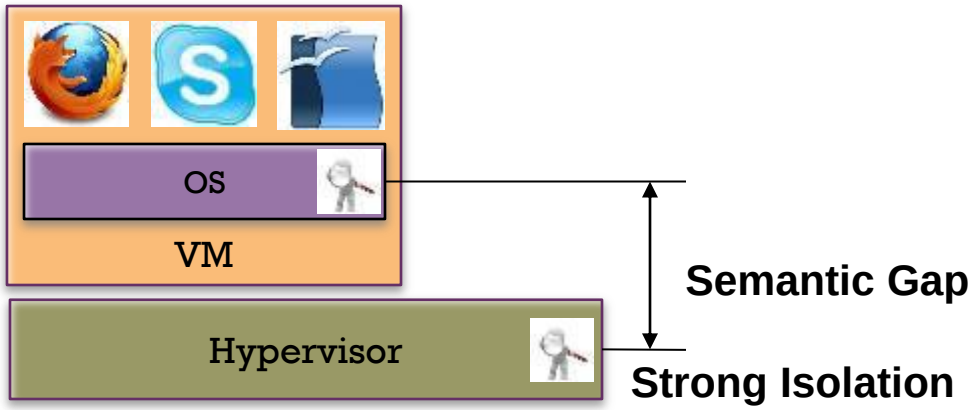

- Challenge: Semantic Gap
- Cur
  - U de the V
  - Limitations:
    - Require effort to understand the guest OS
    - Monitoring tools need to be updated as the guest OS updated
    - Share the same view with attacker: can be manipulated

Our Solutions:

VM Introspection based on

the **Architectural Invariants** of VM environment

What Do We Monitor?

- Guest system's architectural state
  - VM Events, General Purpose and Control Registers
- Guest system's virtual devices
  - Network interfaces, hard disks, memory
- **Advantages:**
  - Non-intrusive to the guest system
  - Hypervisor independent
  - Guest system independent

Reliability and Security Checkers

- **Guest OS hang detection**
  - Infinitive time between two consecutive context switches
- **Hidden rootkit process detection**
  - The number of running processes displayed by the guest system (Task Manager, PS, TOP) is smaller than the number reported by our monitoring tool
- Hypervisor hang detection
  - Infinitive time between VM Exit and VM Entry events
- Guest OS boot sequence integrity
- Process termination detection
- Measure system utilization
- ….

VMIAI Integration with KVM Architecture

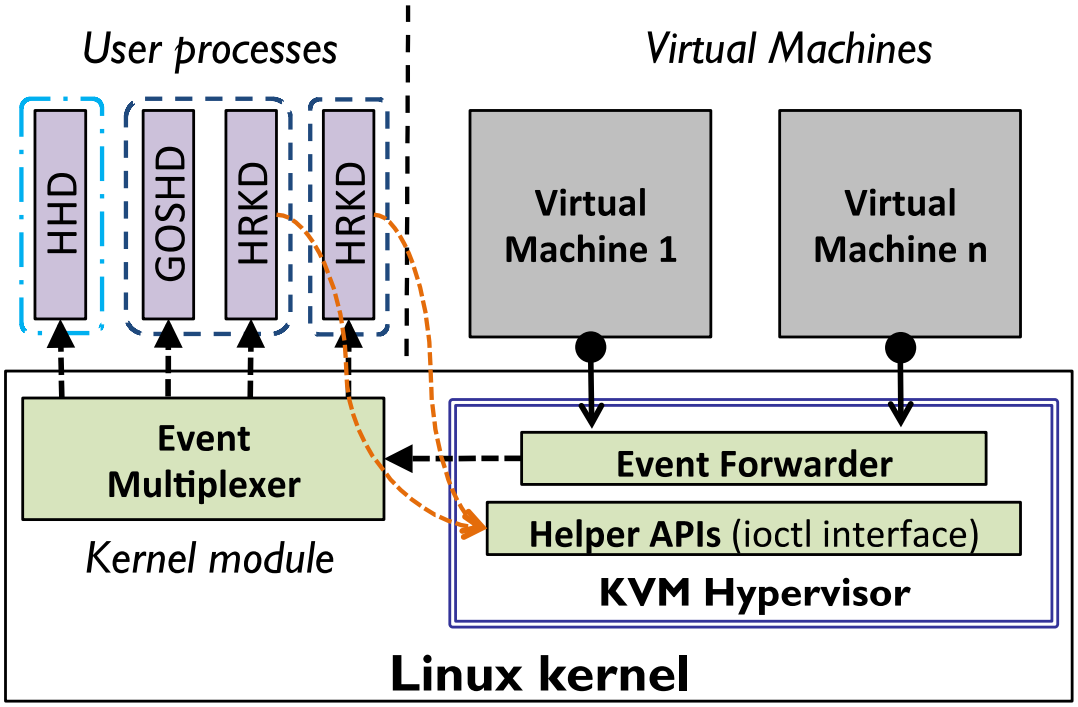

- Example detection modules:
  - Hypervisor hang detection (HHD)
  - Guest OS hang detection (GOSHD)
  - Hidden Rootkit detection (HRKD)

Hidden Rootkit Process Detection

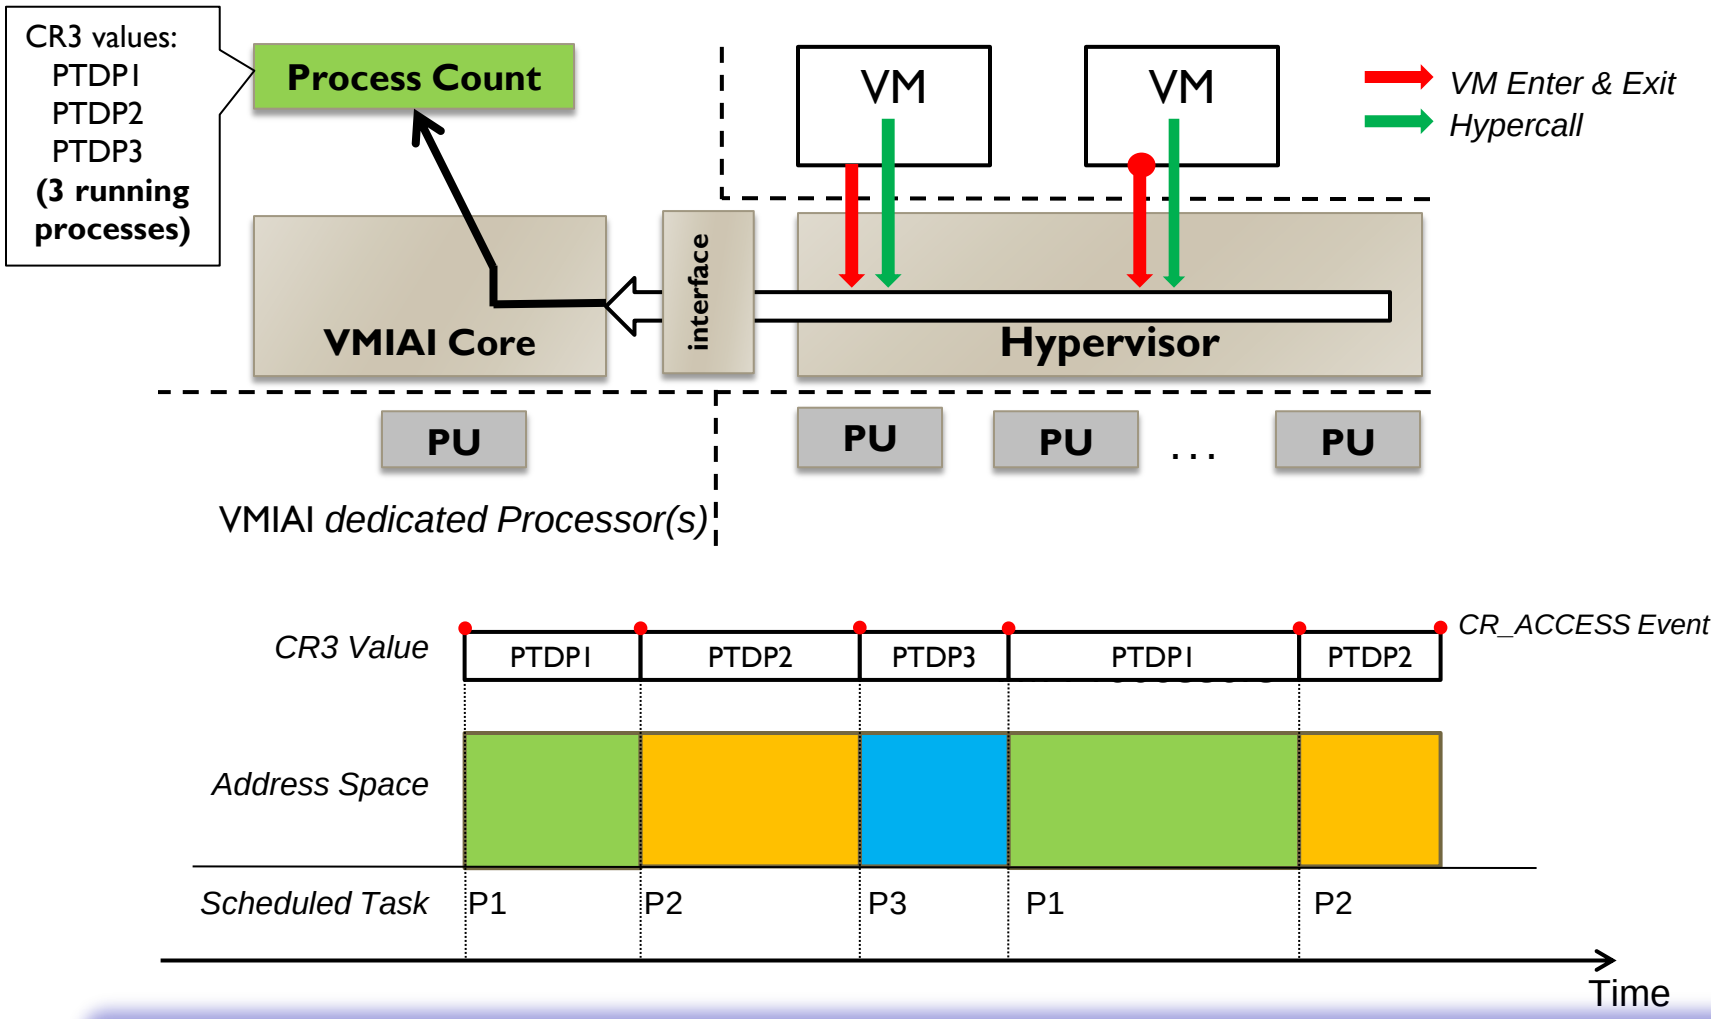

Detect hidden process (rootkit): Process cannot hide from VMIAI 's view

Evaluation of HRKD: Detection Coverage

- HRKD evaluated against real world rootkits on Windows and Linux
- All rootkits successfully detected

Rootkit	*Target OS*	*Hiding techniques*
FU	Windows XP, Vista	DKOM
HideProc	Windows XP, Vista	..
AFX	Windows XP	Hijack system calls
HideToolz	Windows Vista, 7	Hijack system calls
HE4Hook	Windows XP	Hijack system calls
BH	Windows XP	Hijack system calls
Enyelkm 1.2	Linux kernel 2.6	…
SucKIT	Linux kernel 2.6	Kmem, dkom
PhalanX	Linux kernel 2.6	DKOM

- Detection capability not affected by implementation or hiding techniques of the rootkits.
- HRKD can detect future hidden rootkits regardless of their newly invented hiding mechanism

Conclusions

- Propose Virtual Machine Introspection framework based on Hardware Architectural Invariants to provide security and reliability monitoring for guest VMs and hypervisor
- Introduced two monitoring techniques
  - Hidden Rootkit Detection (HRKD)
    - detects real-world rootkits, including a new type of hidden rootkit that is specifically designed to defeat the current VMI techniques
  - Guest OS Hang Detection (GOSHD)
    - detect hangs occurring in the guest OS kernel with short detection latencies
- Proposed monitoring techniques cause less than 3% and >0.5% performance loss to Disk IO intensive and CPU intensive workload, respectively