Virtualization Security State of the Union

Ruxcon Melbourne 2010

David Jorm

djorm@redhat.com

Contents

- Virtualization & multi-tenancy
- New attack surfaces
- Vulnerability examples
- Defence attempts
- Market impact

The Olden Days

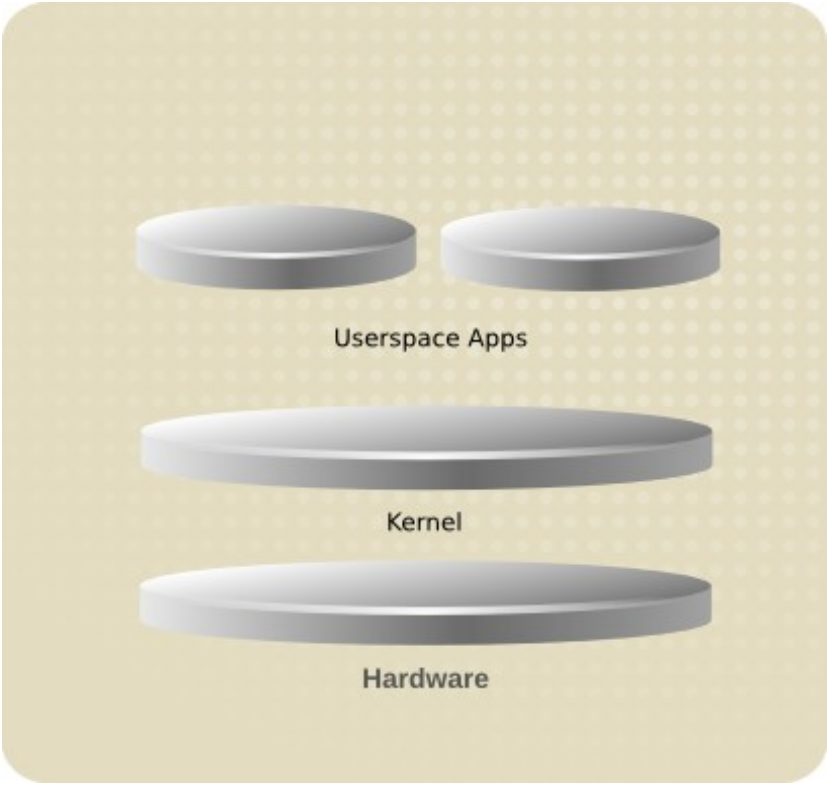

- Attack surfaces limited to userspace, kernel and networking
- Attacking one userspace app could affect others, e.g. suid root binaries

Virtualization

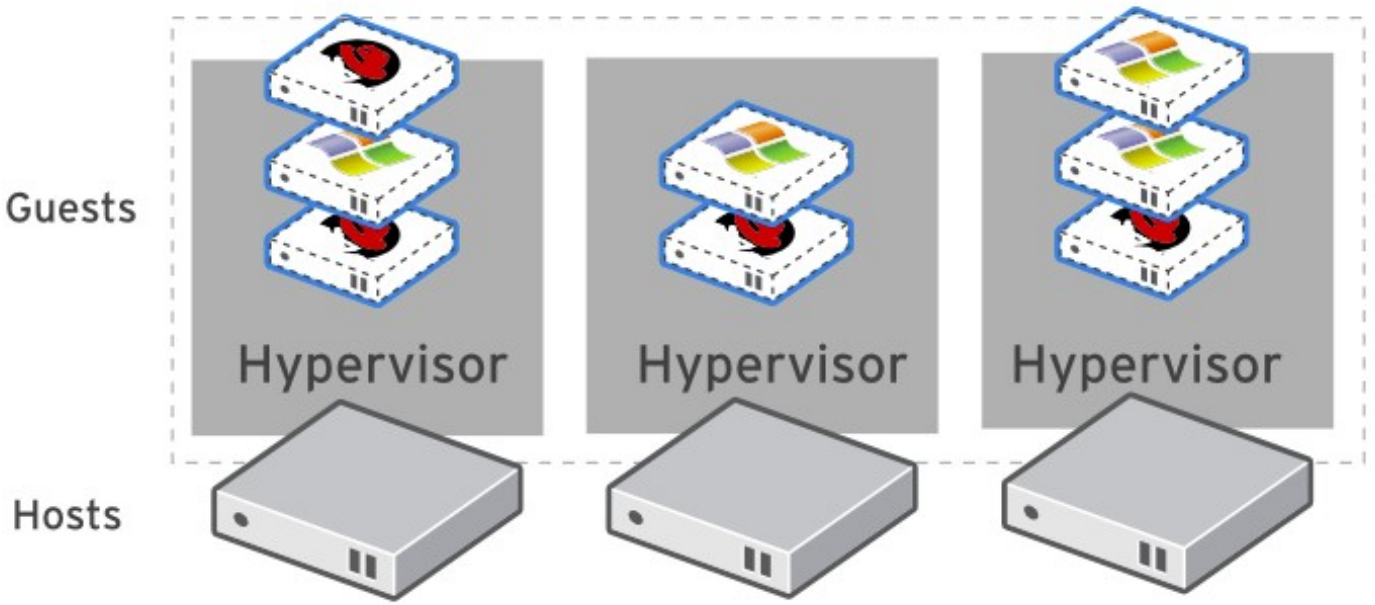

- One host runs multiple guests
- Guests share resources – CPU, disk, memory
- Guests have different security contexts

Guest Portability

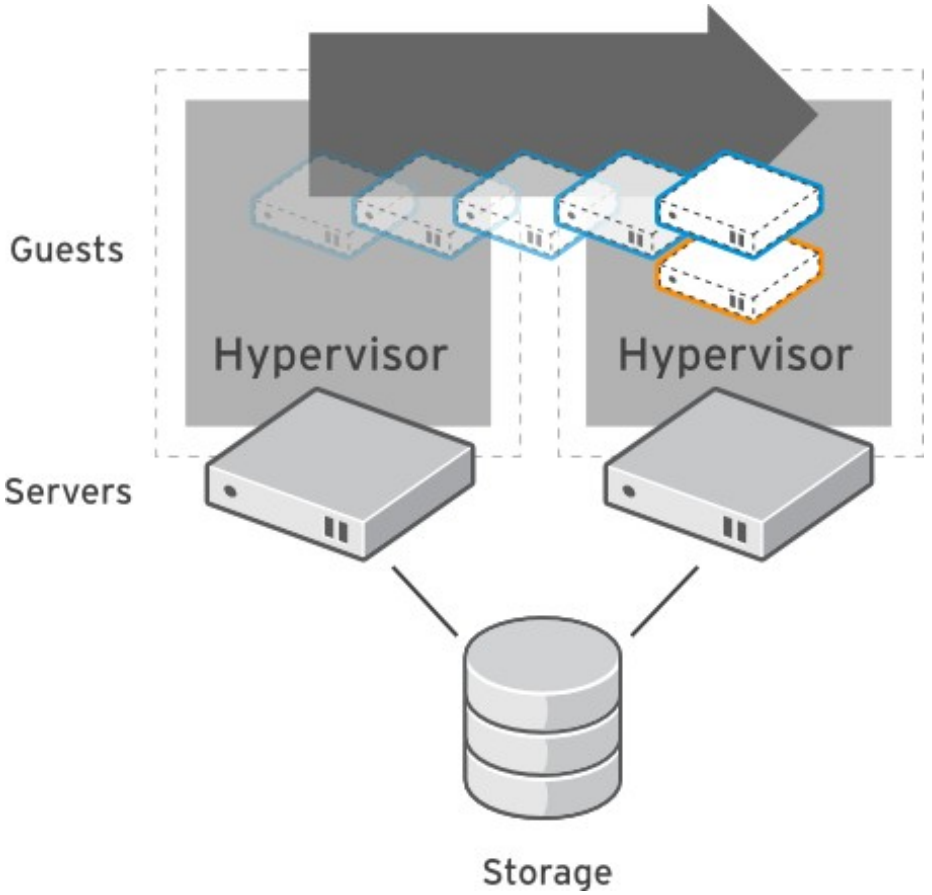

- Live migration in enterprise virtualization platforms de-couples guests from host machine.
- A given guest could be running on any host

Cloud

- Guests can run on any host, outside your environment
- Guests share physical resources with other guests from outside your network boundary
- Physical access controls no longer apply – what is inside and outside your network is now purely virtual

Attack Surfaces

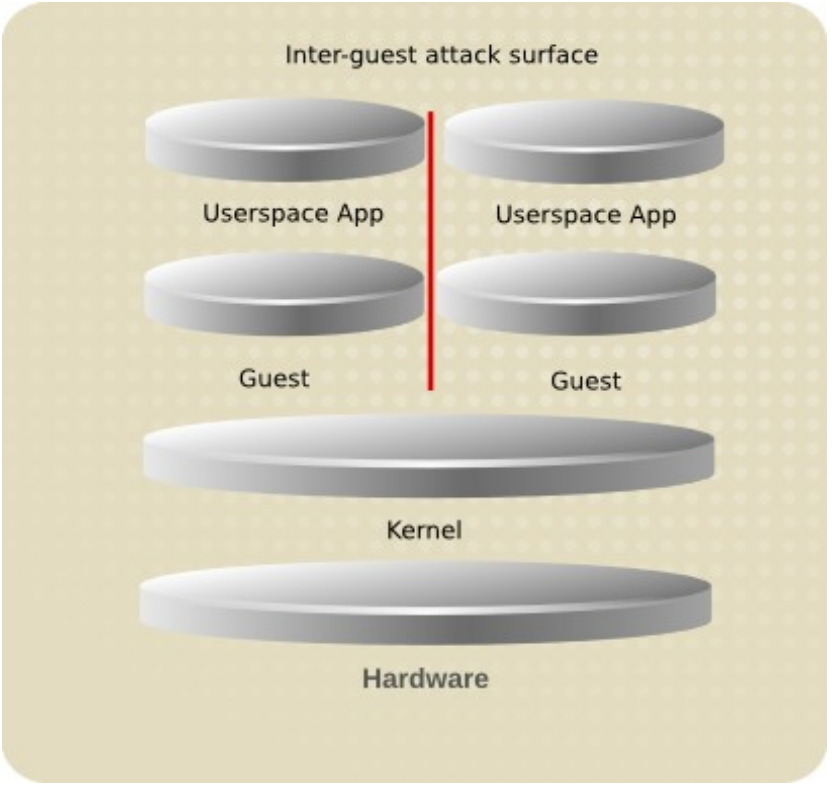

- New inter-guest attack surface, in addition to existing attack surfaces
- Blah blah theory surface, where is the ownage?

Vmware workstation 5.5 attack (2007)

- Shared clipboard feature – share copy/paste buffer between VMs
- Altering setting did not persist without a VM restart or trigger action
- Malicious guest could then read paste buffer from any other machine – passwords?
- On a workstation this is a less serious issue, but demonstrative

http://www.securiteam.com/securitynews/5GP021FKKO.html

Vmware workstation shared folder CVE-2007-1744

- Shared folders from the host are visible to guests
- Exploitation allows a guest to read or write arbitrary files on the host system
- Since other guests' files are on the host, this impacts the inter-guest attack surface

Xen buffer overflow of PV FB CVE-2008-1943

- Buffer overflow of paravirt frame buffer
- Allows guests to crash the host and execute arbitrary code using a crafted shared framebuffer description
- Dom0 thinks it is just painting framebuffer to screen, but it contains metadata. Not validating the metadata leads to vulnerability

Vmware ESX CVE-2009-3733

"A directory traversal vulnerability allows for remote retrieval of any file from the host system. In order to send a malicious request, the attacker will need to have access to the network on which the host resides"

- A guest is on the same network as the host – either via bridging or passthru
- Another guest's image exists as files on, or accessible by, the host system
- Potential for inter-guest impact – no weaponized exploit available to prove

RHEV not zeroing disk images CVE-2010-2224

- Under certain conditions, VM disk images are not zeroed out when deleted
- New VMs will be potentially allocated virtual disks using this un-initialized space
- Malicious VMs could get all forensic and read the pre-existing data

Oracle VM remote command exec CVE-2010-3583

- Discovered by Onapsis, advisory released 2nd Nov
- Oracle VM Agent runs as root, exposes functions through XML-RPC
- Undisclosed vulnerability allows for execution of arbitrary OS commands via XML-RPC
- Oracle VM is Xen based – with command injection, could manipulate other guests using Xen tools

Hypervisor detection via samepage merging

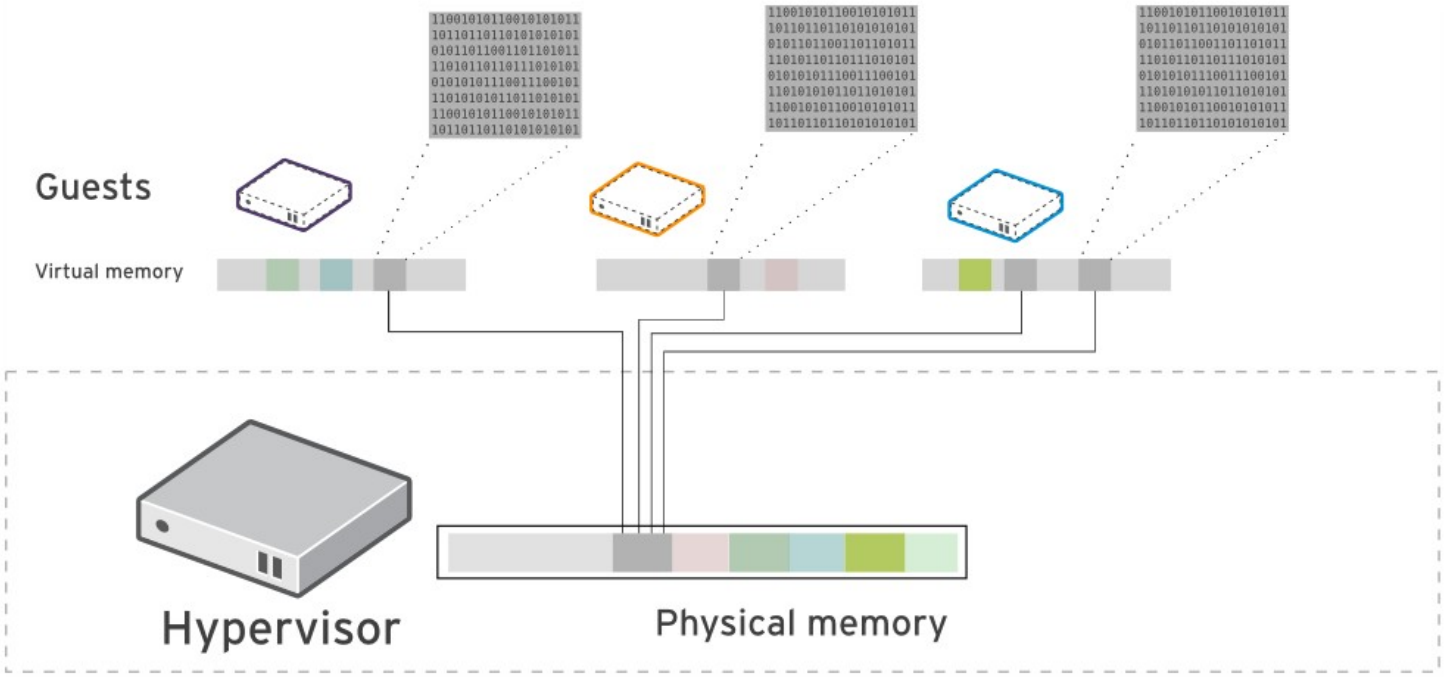

- Samepage merging allows for memory over-commit, using copy on write

Hypervisor detection via samepage merging

- By detecting patterns of memory access time, clear signals for bare metal and virtualization can be found
- Within virtualized environments, the samepage merging technique used will have its own signature
- Allows for detection of hypervisors with hardware acceleration, such as KVM. Applications for malware evasion, since VMs used to analyse them.

Ref: “VMM detection through samepage merging”, Daniel Fernandez

Defence – Network attack

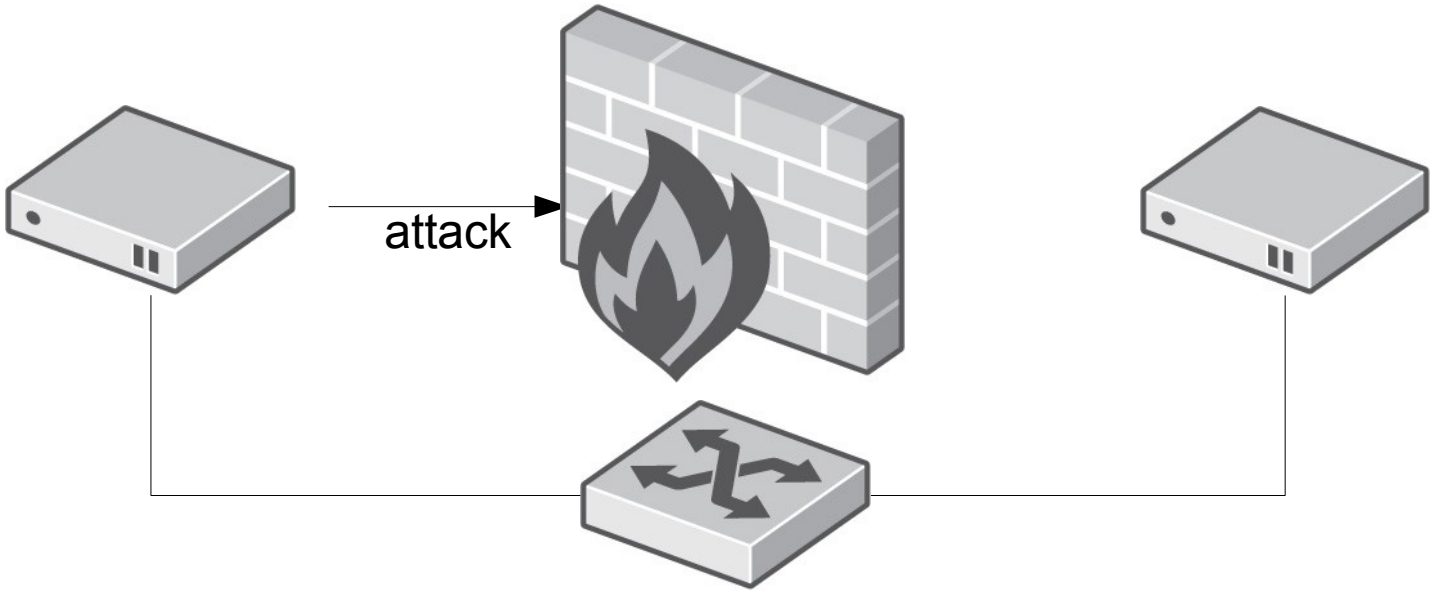

- Firewall used as infrastructure component to limit attack surface for network environments

Defence – Inter-guest attack

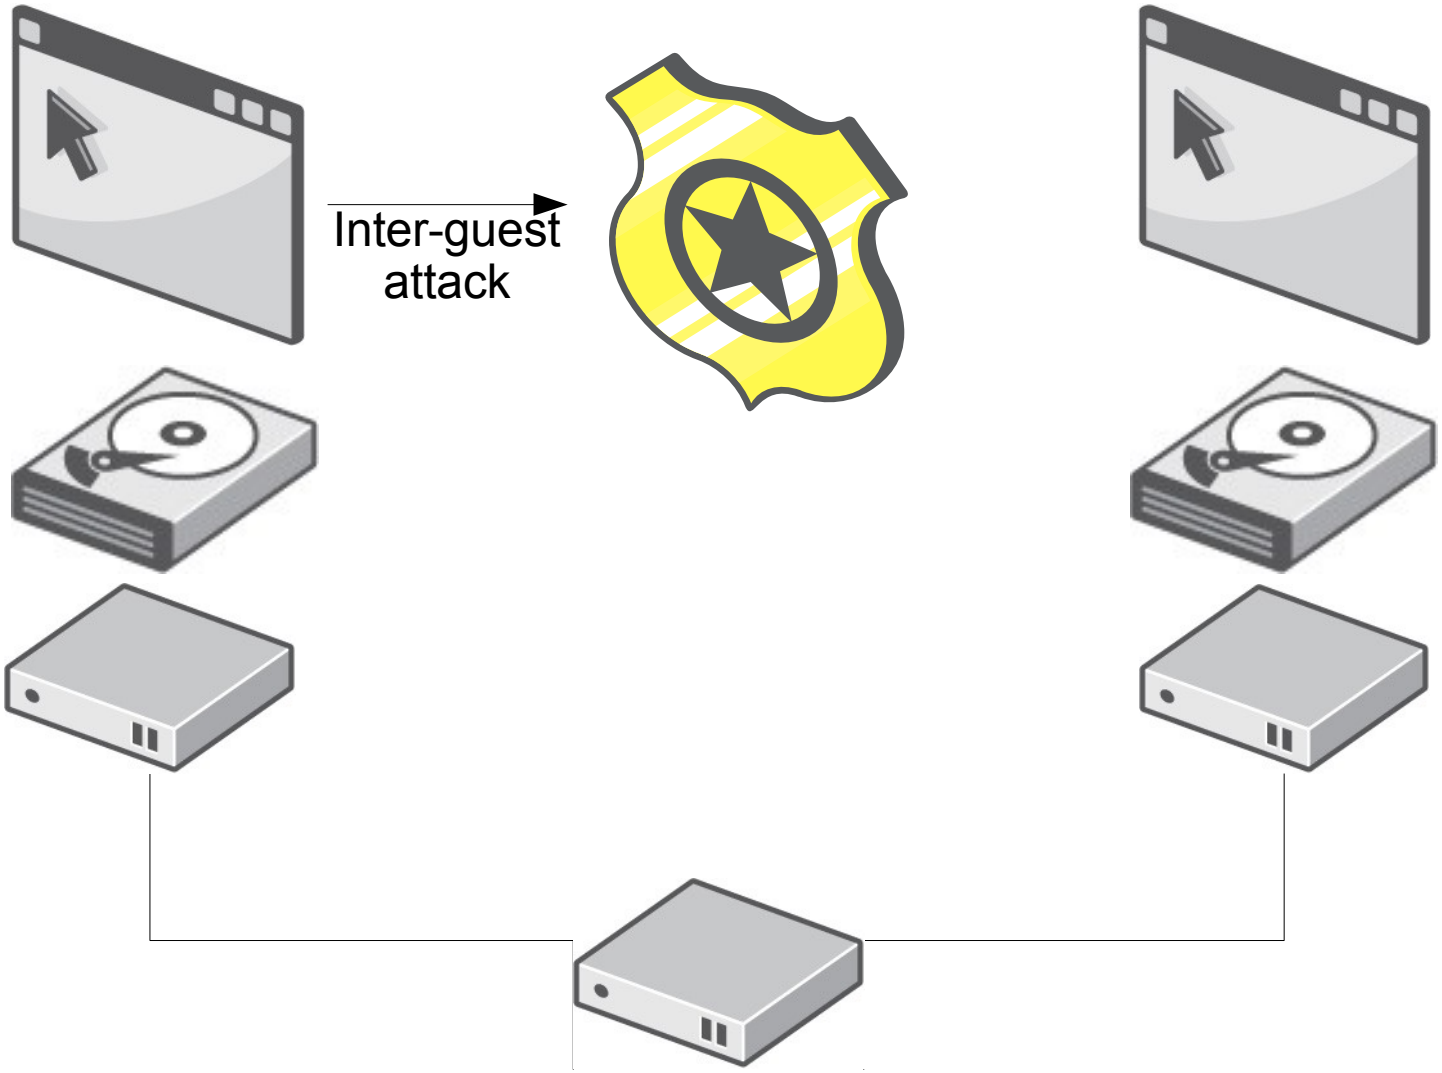

Defence - Linux/KVM

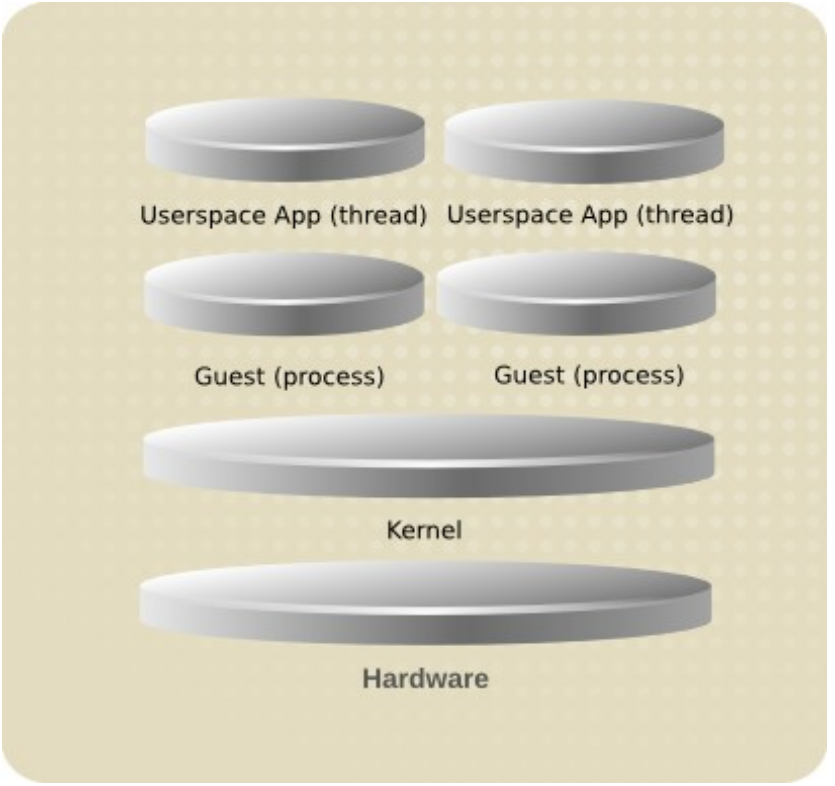

- On KVM, guests are regular processes. Disks are regular files.
- SELinux provides MAC for processes and files

SELinux

- MAC implementation
- Process, files and devices are labeled
- Rules govern how process labels interact with other process/file labels
- Kernel enforces these rules

sVirt

- Isolate guests using MAC policy, contain hypervisor breaches
- Implemented as a pluggable security framework for libvirt, which is a virtualization management abstraction layer
- Various MAC implementations (SELinux, SMACK) can be used with various hypervisors

VMware

“ESX is meant to be used in production environments in which the guest virtual machines can potentially be exposed to malicious users and network traffic. Strong isolation and strict separation of management greatly reduce any risk of harmful activity going beyond the boundaries of the virtual machine.”

Market impact

“More than half of U.S. organizations are adopting cloud services, but only 47 percent of IT professionals believe that those services are properly secured before they're deployed. “

“68 percent of those surveyed thought that cloud computing was too risky”

- Stefanie Hoffman, CRN

Get that paper

Questions?