RED HAT :: CHICAGO :: 2009
SUMMIT

Secure Virtualization Using SELinux

Daniel J Walsh
dwalsh@redhat.com

presented by

Virtualization Dream

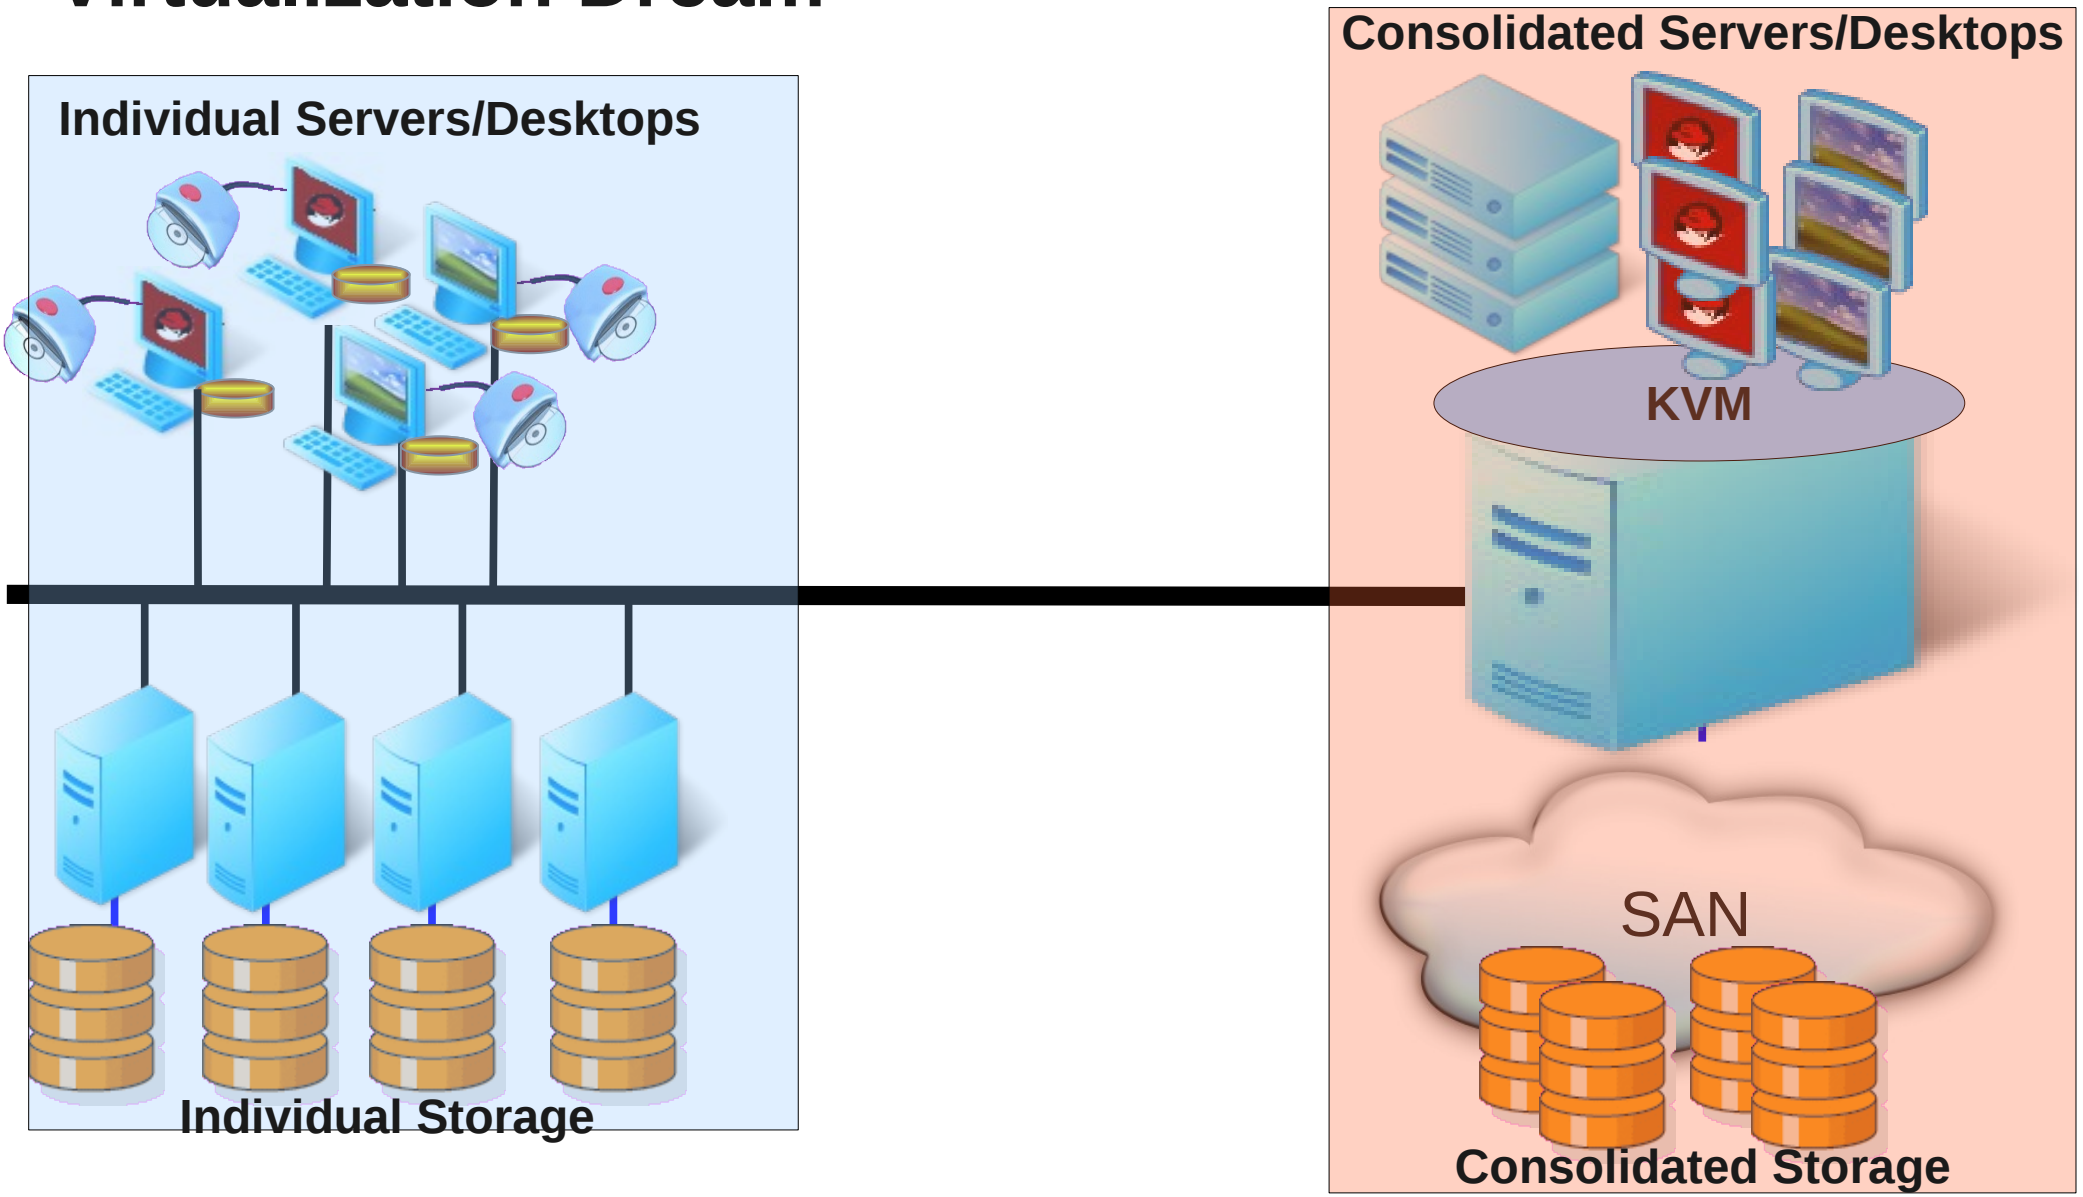

Before Virtualization

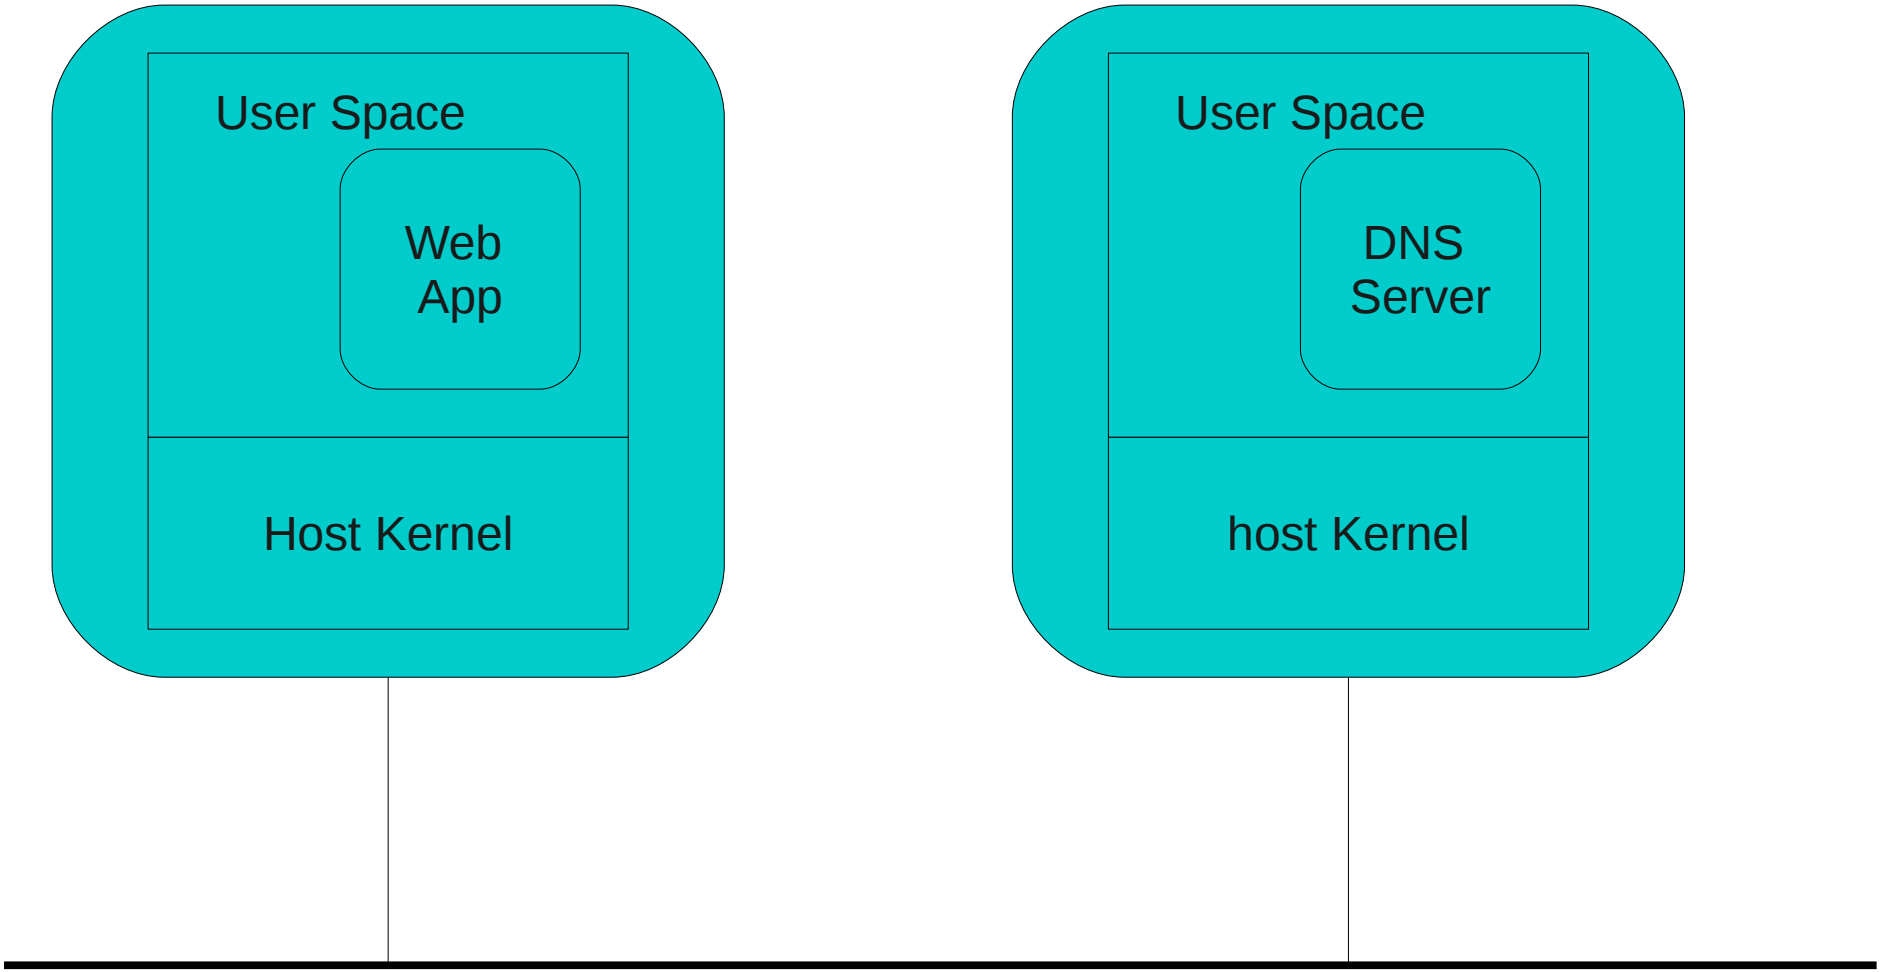

After Virtualization

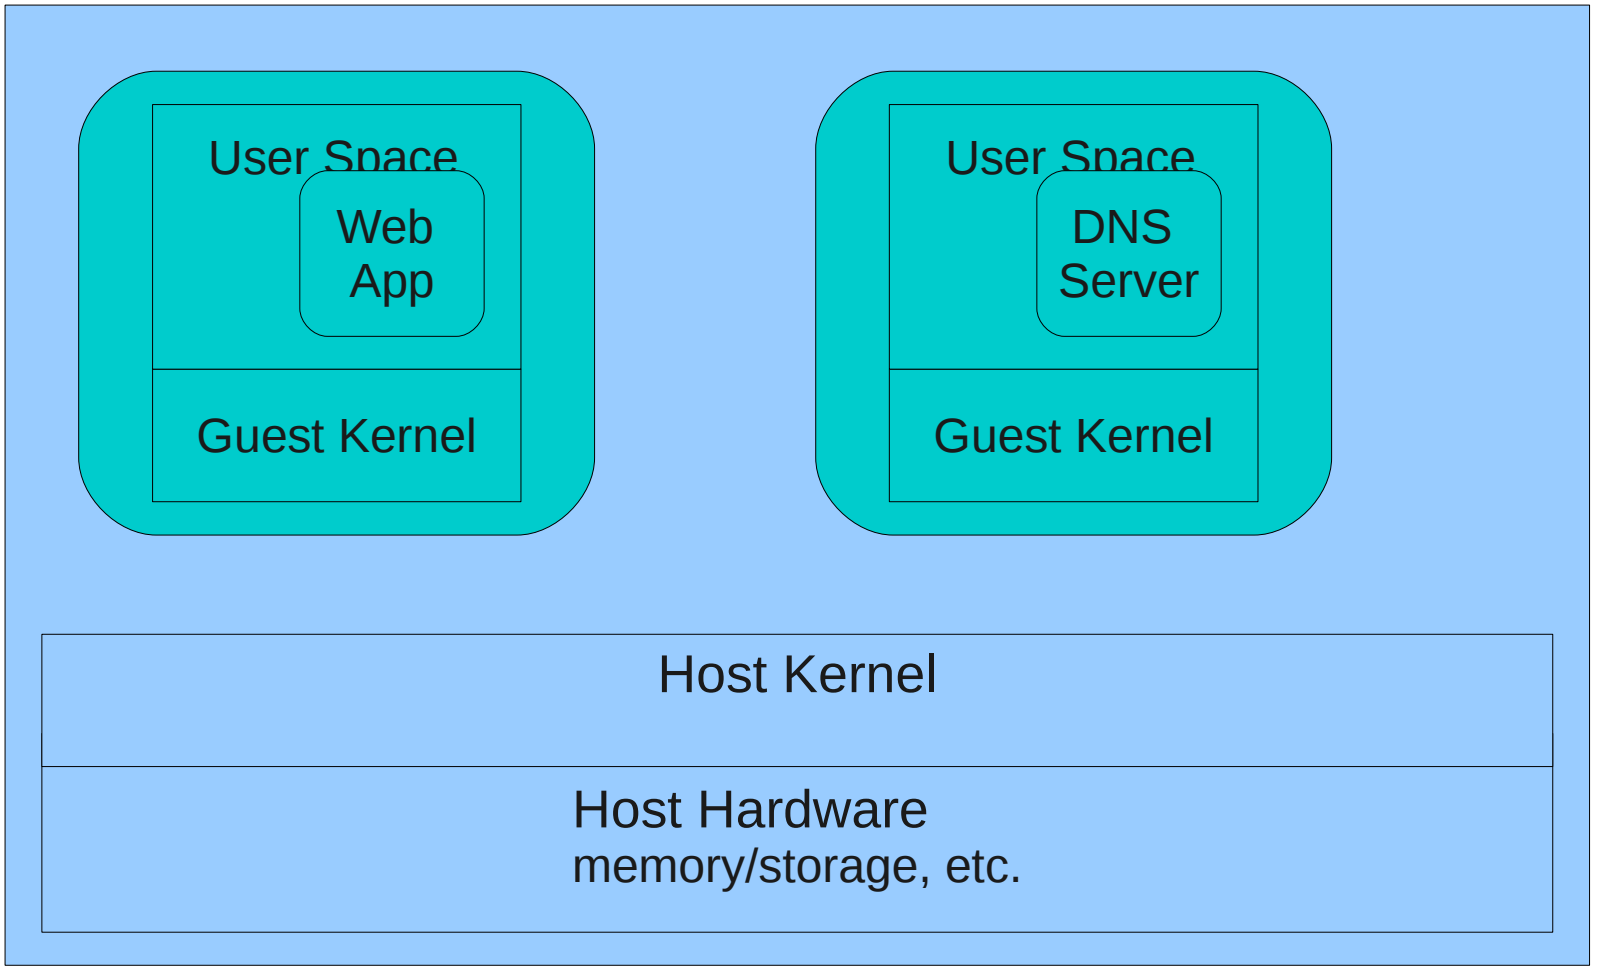

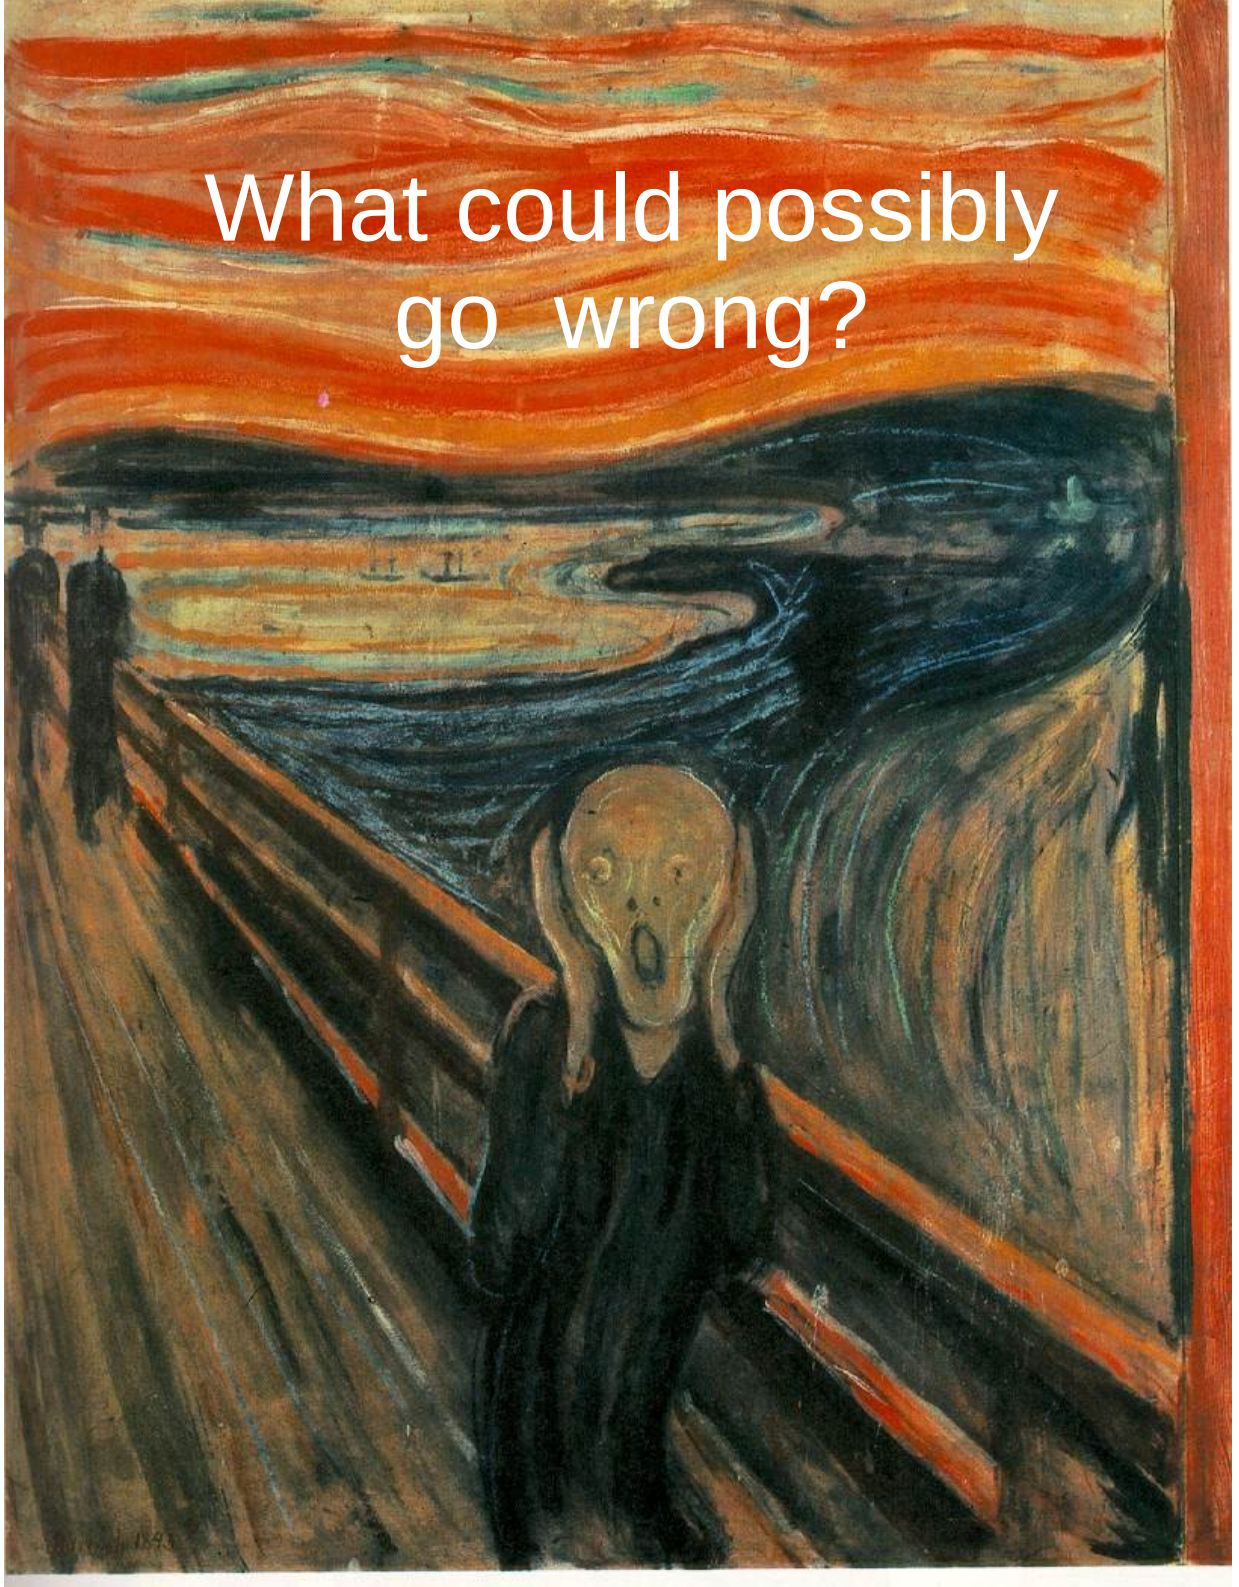

Hypervisor vulnerabilities

- Not theoretical
- Evolving field
- Potentially huge payoffs
- Xen already compromised...

XEN Vulnerability
http://www.hacker-soft.net/Soft/Soft_13289.htm

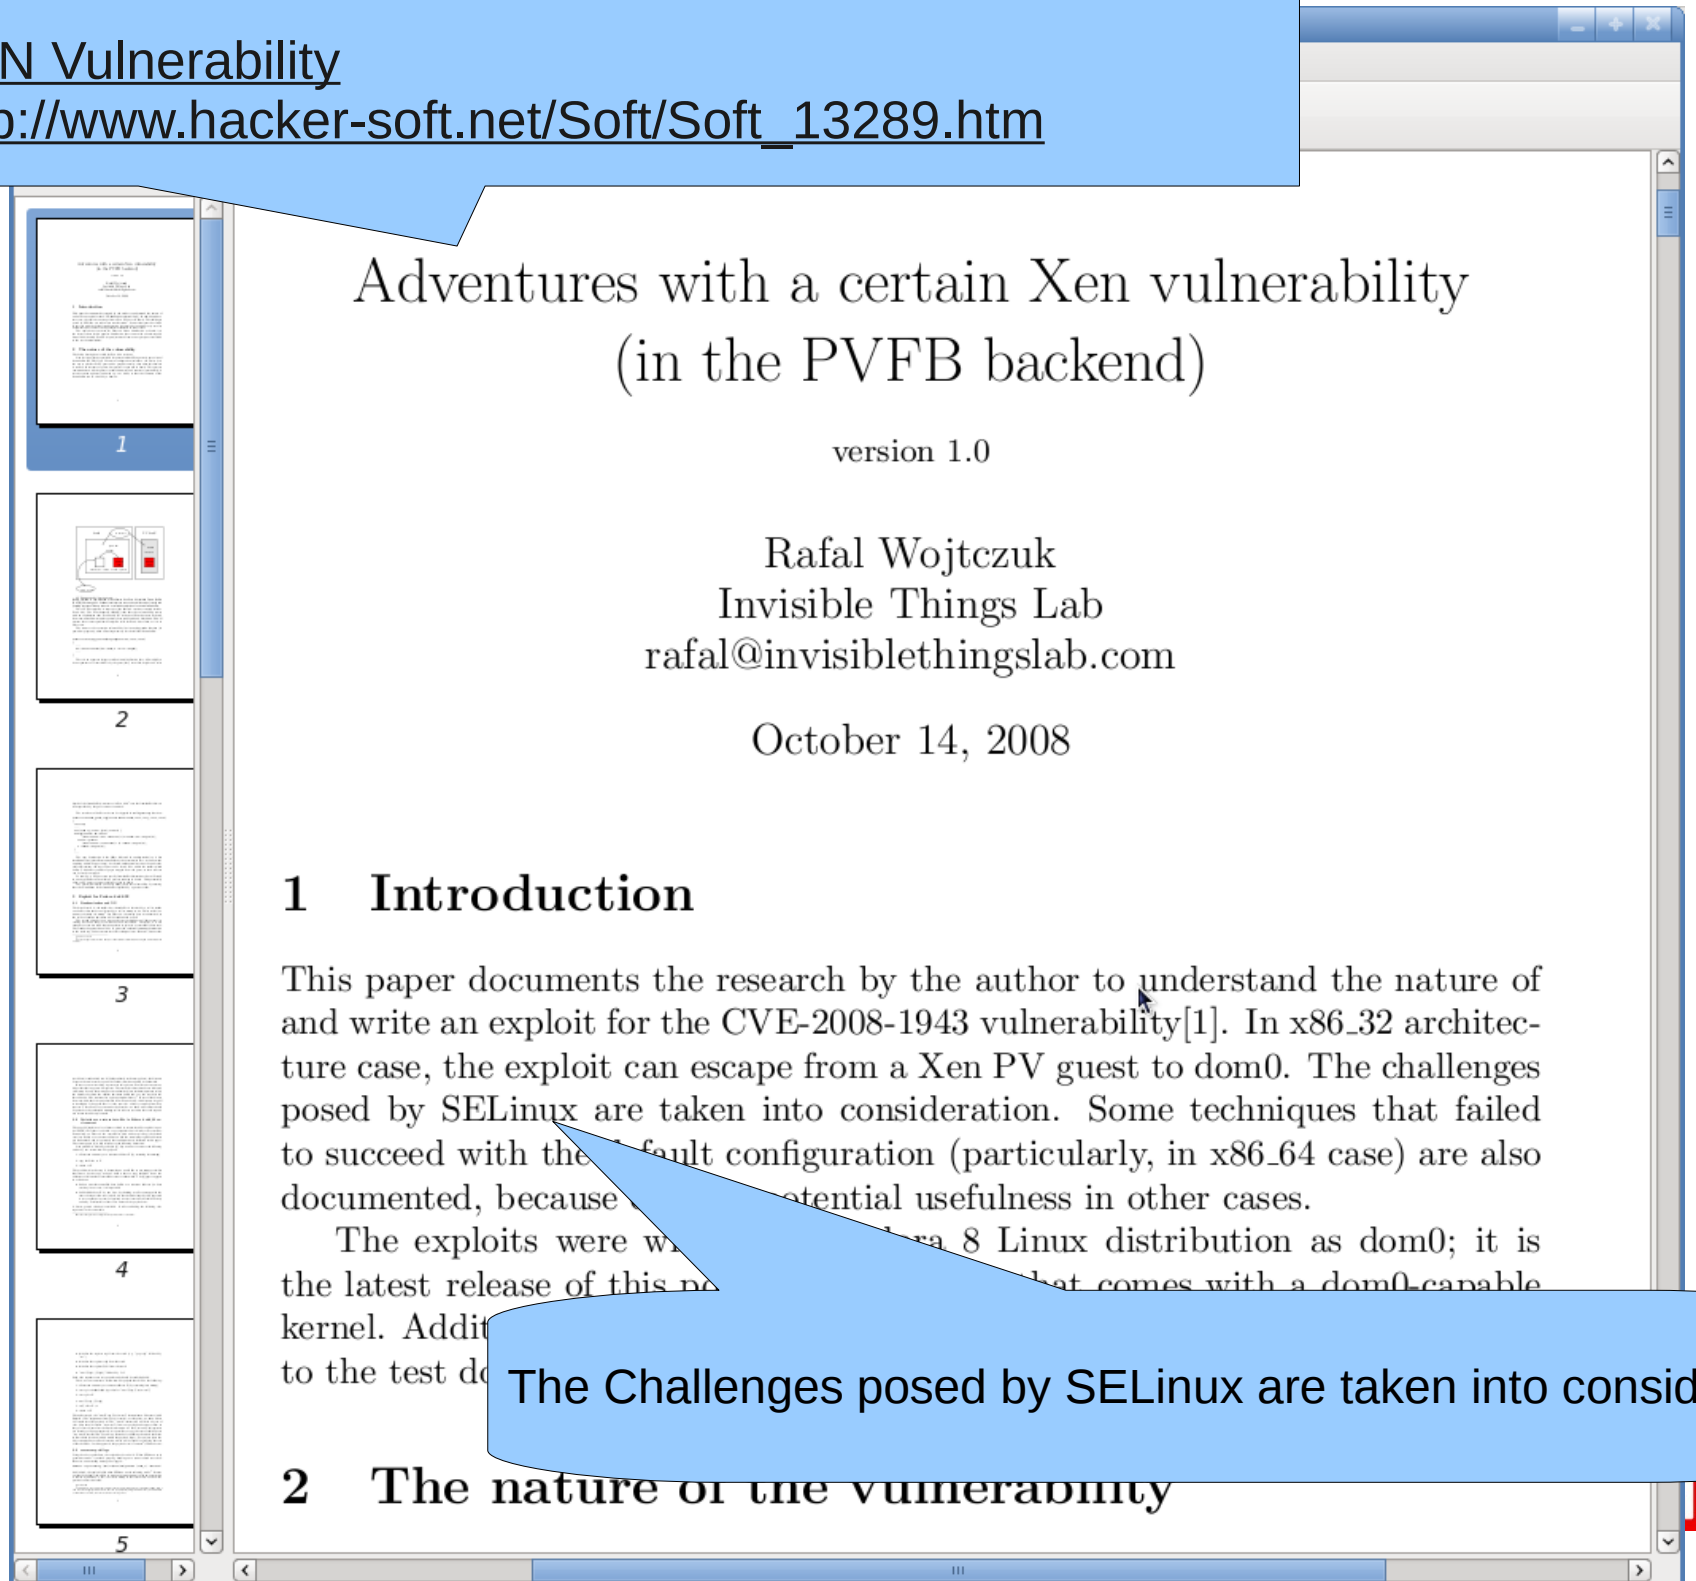

Adventures with a certain Xen vulnerability
(in the PVFB backend)

version 1.0

Rafal Wojtczuk
Invisible Things Lab
rafal@invisiblethingslab.com

October 14, 2008

1 Introduction

This paper documents the research by the author to understand the nature of and write an exploit for the CVE-2008-1943 vulnerability[1]. In x86_32 architecture case, the exploit can escape from a Xen PV guest to dom0. The challenges posed by SELinux are taken into consideration. Some techniques that failed to succeed with the default configuration (particularly, in x86_64 case) are also documented, because of potential usefulness in other cases.

The exploits were written for Fedora 8 Linux distribution as dom0; it is the latest release of this product that comes with a dom0-capable kernel. Additional to the test do

The Challenges posed by SELinux are taken into consideration.

2 The nature of the vulnerability

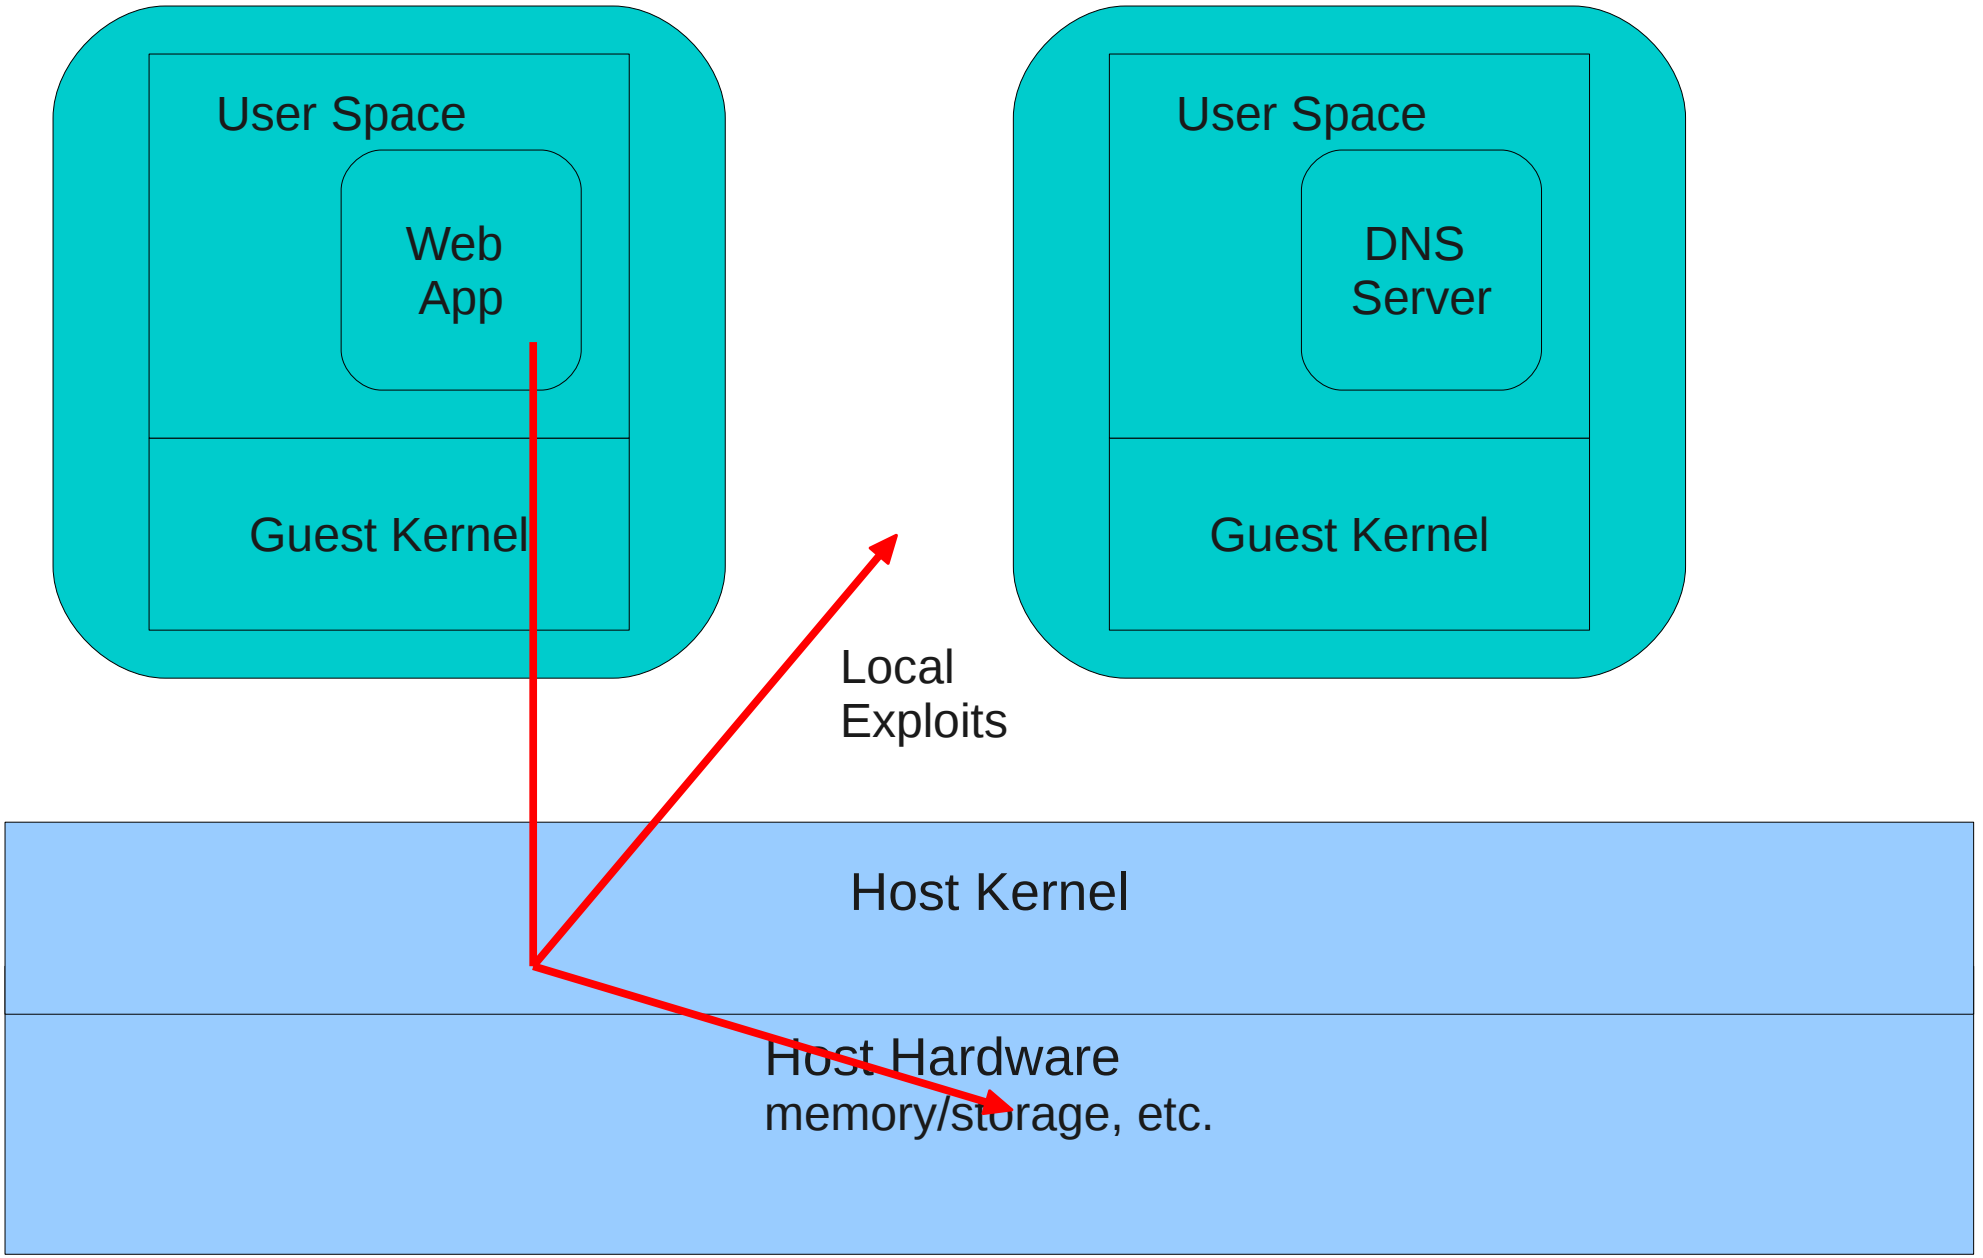
User Space
Web App
Guest Kernel
User Space
DNS Server
Guest Kernel
Local Exploits
Host Kernel
Host Hardware
memory/storage, etc.

Who is the weakest link?

aris™

suse
A NOVELL COMPANY

Host Kernel

Host Hardware
memory/storage, etc.

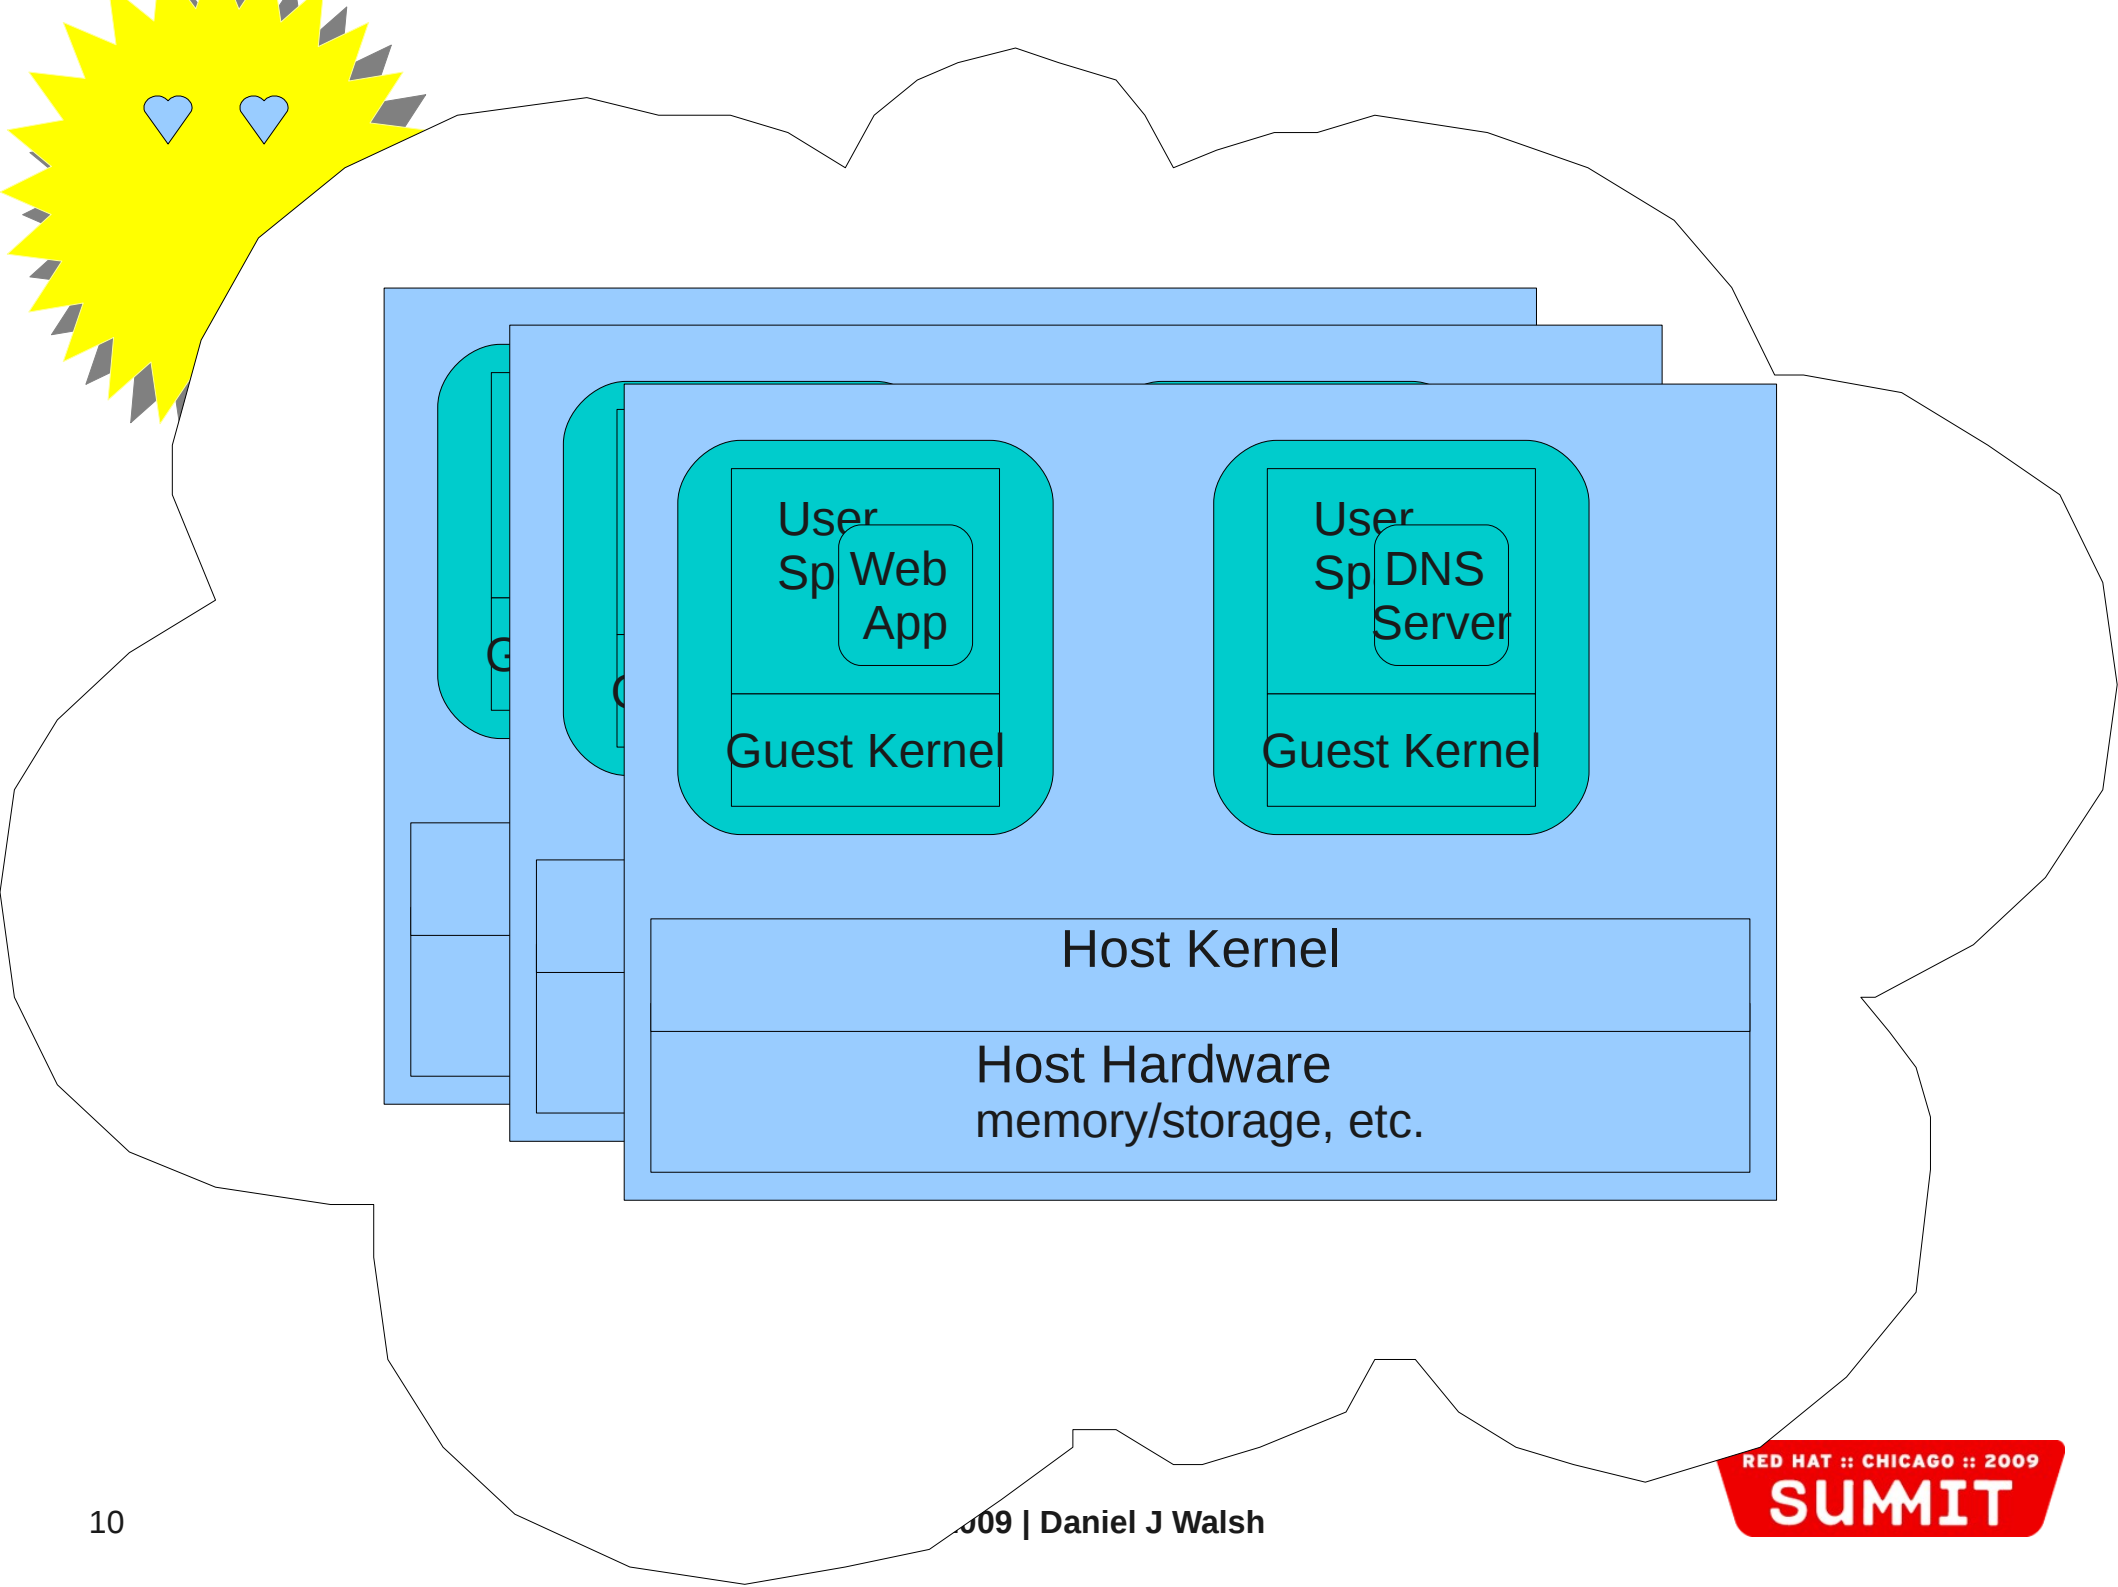
User Sp
Web App
Guest Kernel
User Sp
DNS Server
Guest Kernel
Host Kernel
Host Hardware
memory/storage, etc.

Hansel and Gretl?

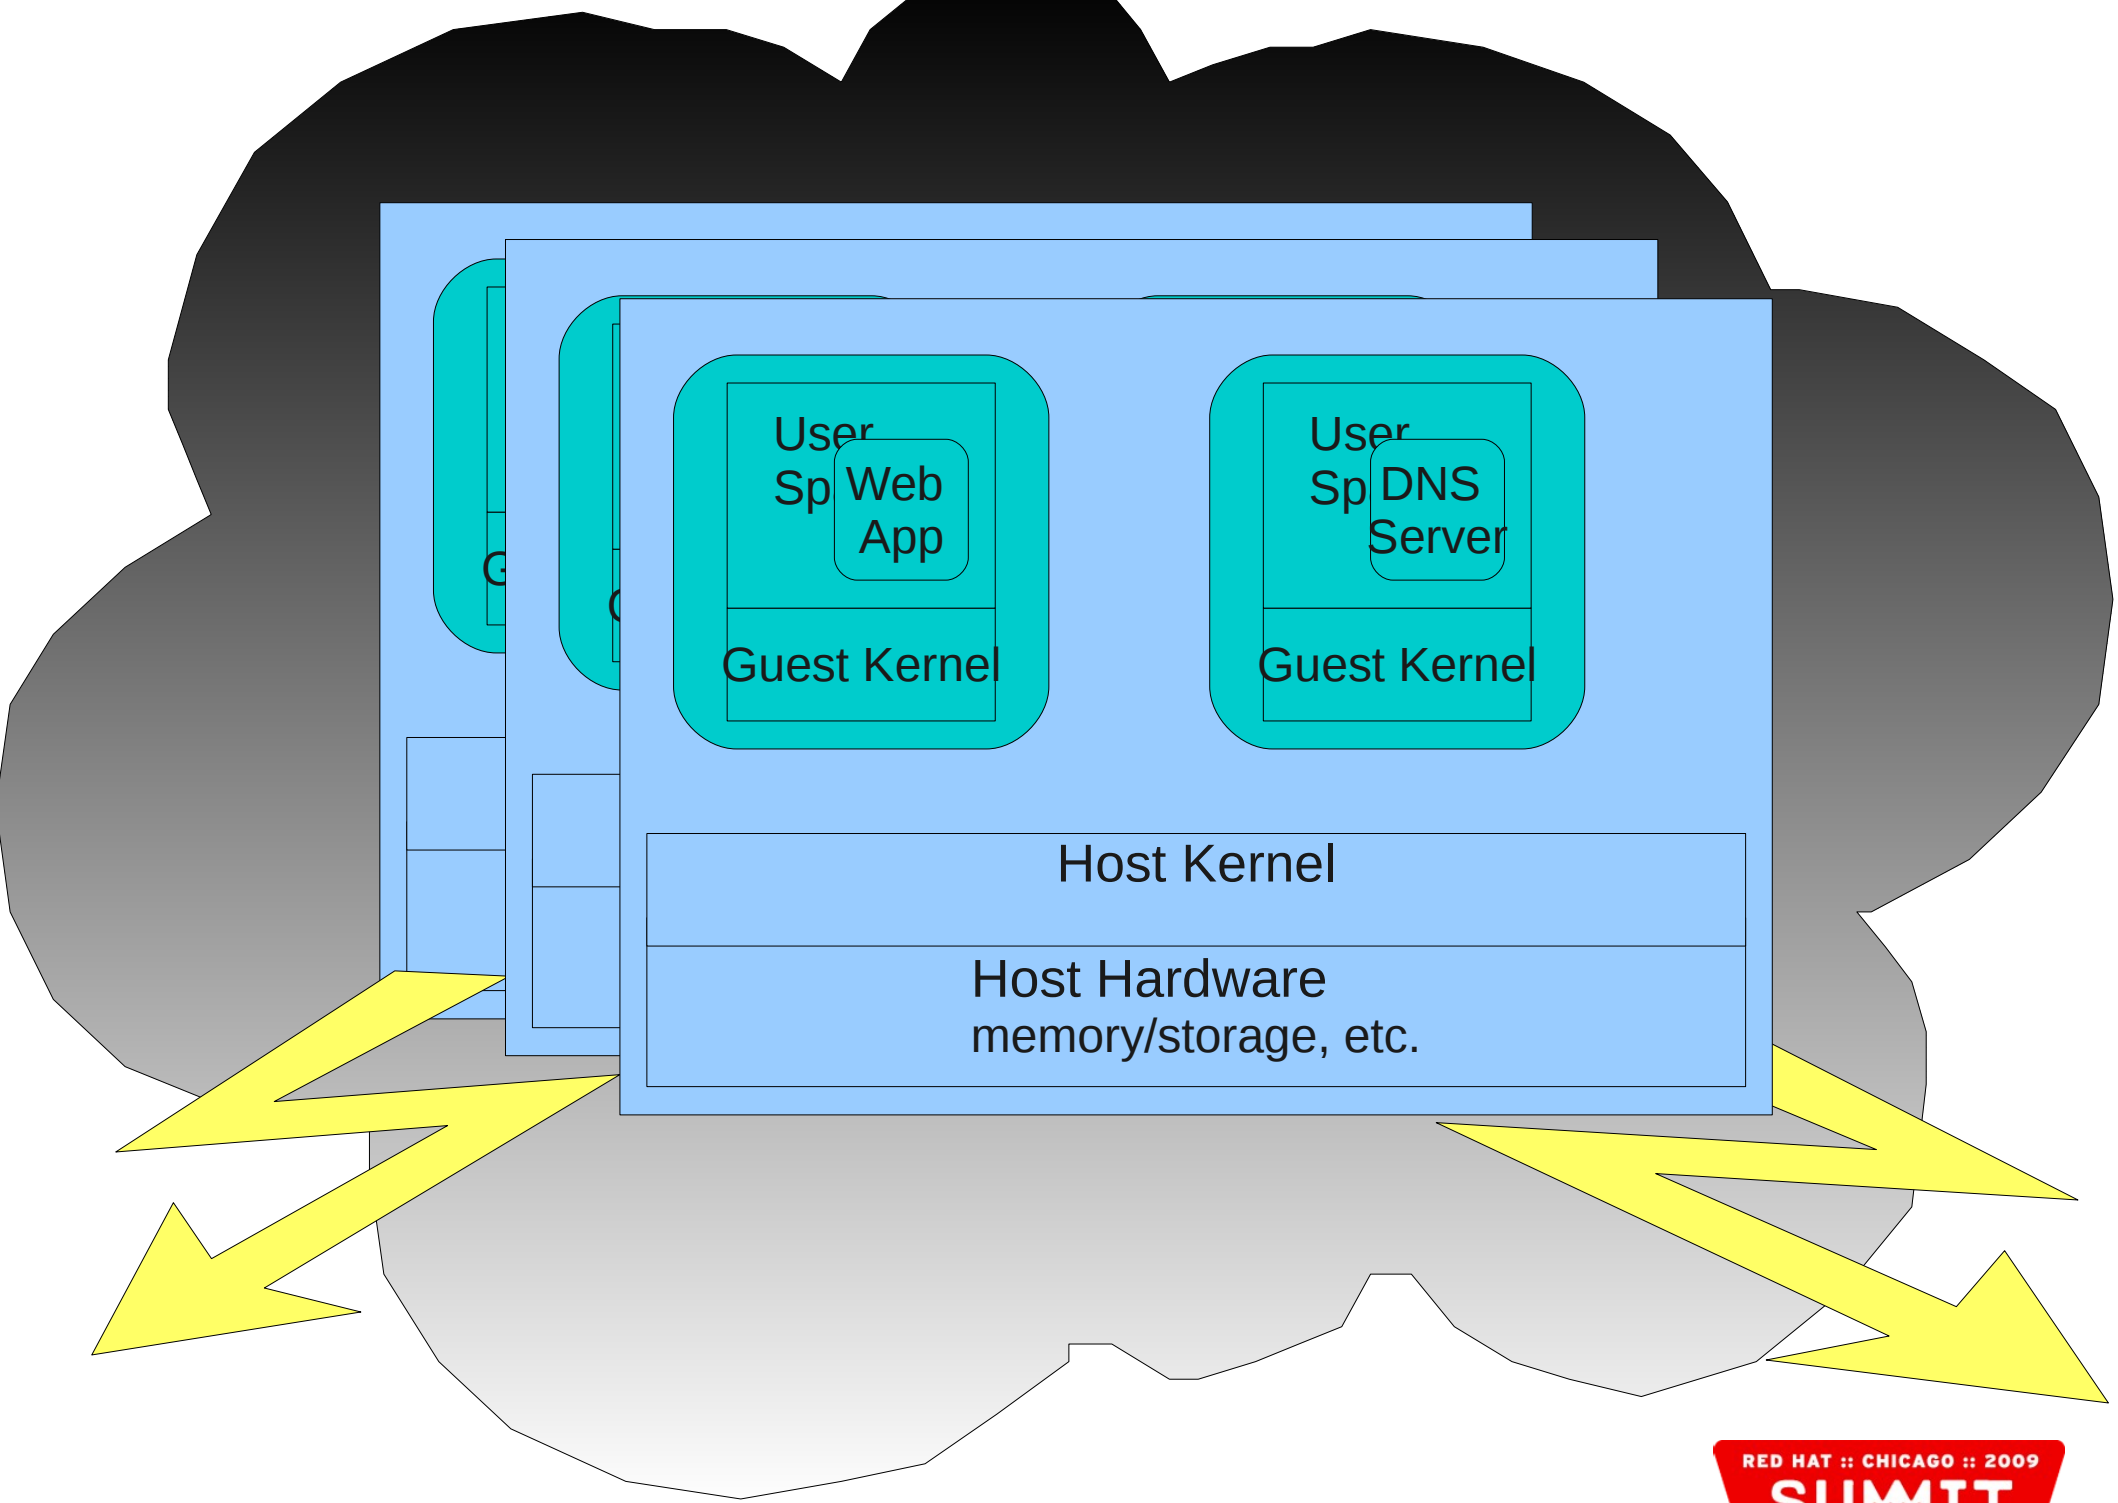

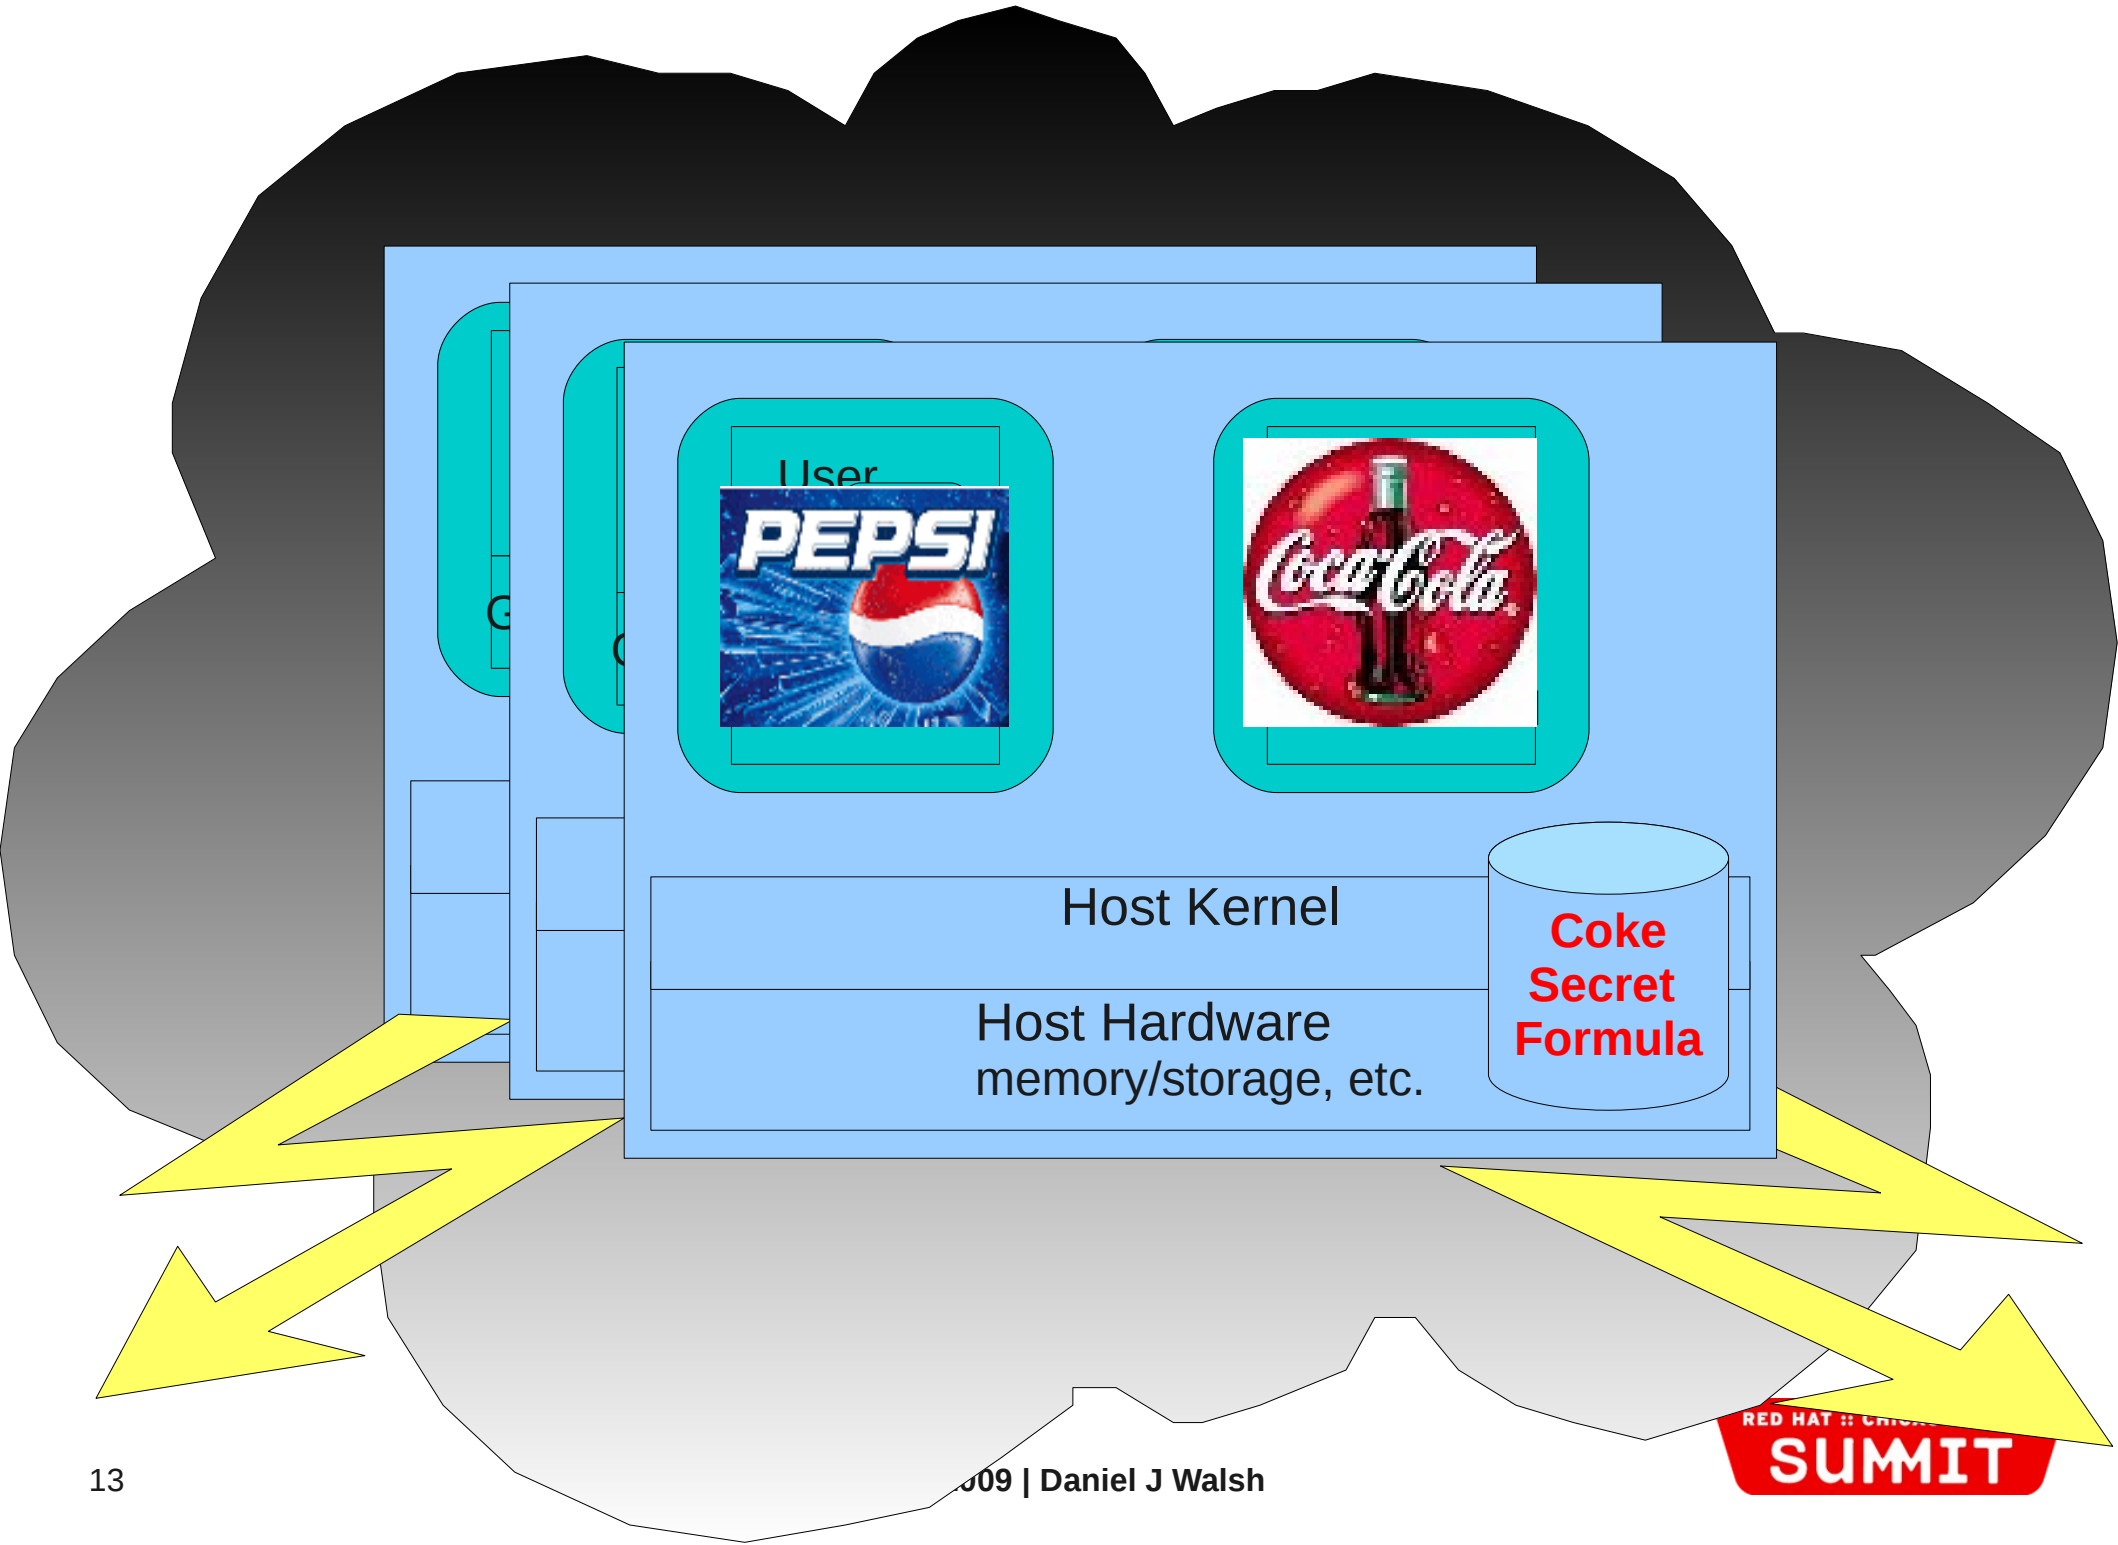
PEPSI
Host Kernel
Host Hardware
memory/storage, etc.
Coke
Secret
Formula

Enter SELinux..

SELinux is all about labeling

- Processes get labels
  - Virtual desktops are processes!!!
- Files/Devices Get Labels
  - Virtual desktop images are stored on files/devices!!!!
- Rules govern how Process Labels Interact with Process/File Labels.
- Kernel Enforces these Rules.

Svirt in a Nutshell

Isolate guests using Mandatory Access Control security policy

Contain Hypervisor Breaches

Libvirt – Dynamic Labeling

- Generates a Random unused MCS label.
  - MCS – Multiple Category Security
- Labels the image file/device - svirt_image_t:MCS1
- Launches the image - svirt_t:MCS1
- Labels R/O Content – virt_content_t:s0
- Labels Shared R/W Content – svirt_t:s0
- Labels image on completion - virt_image_t:s0

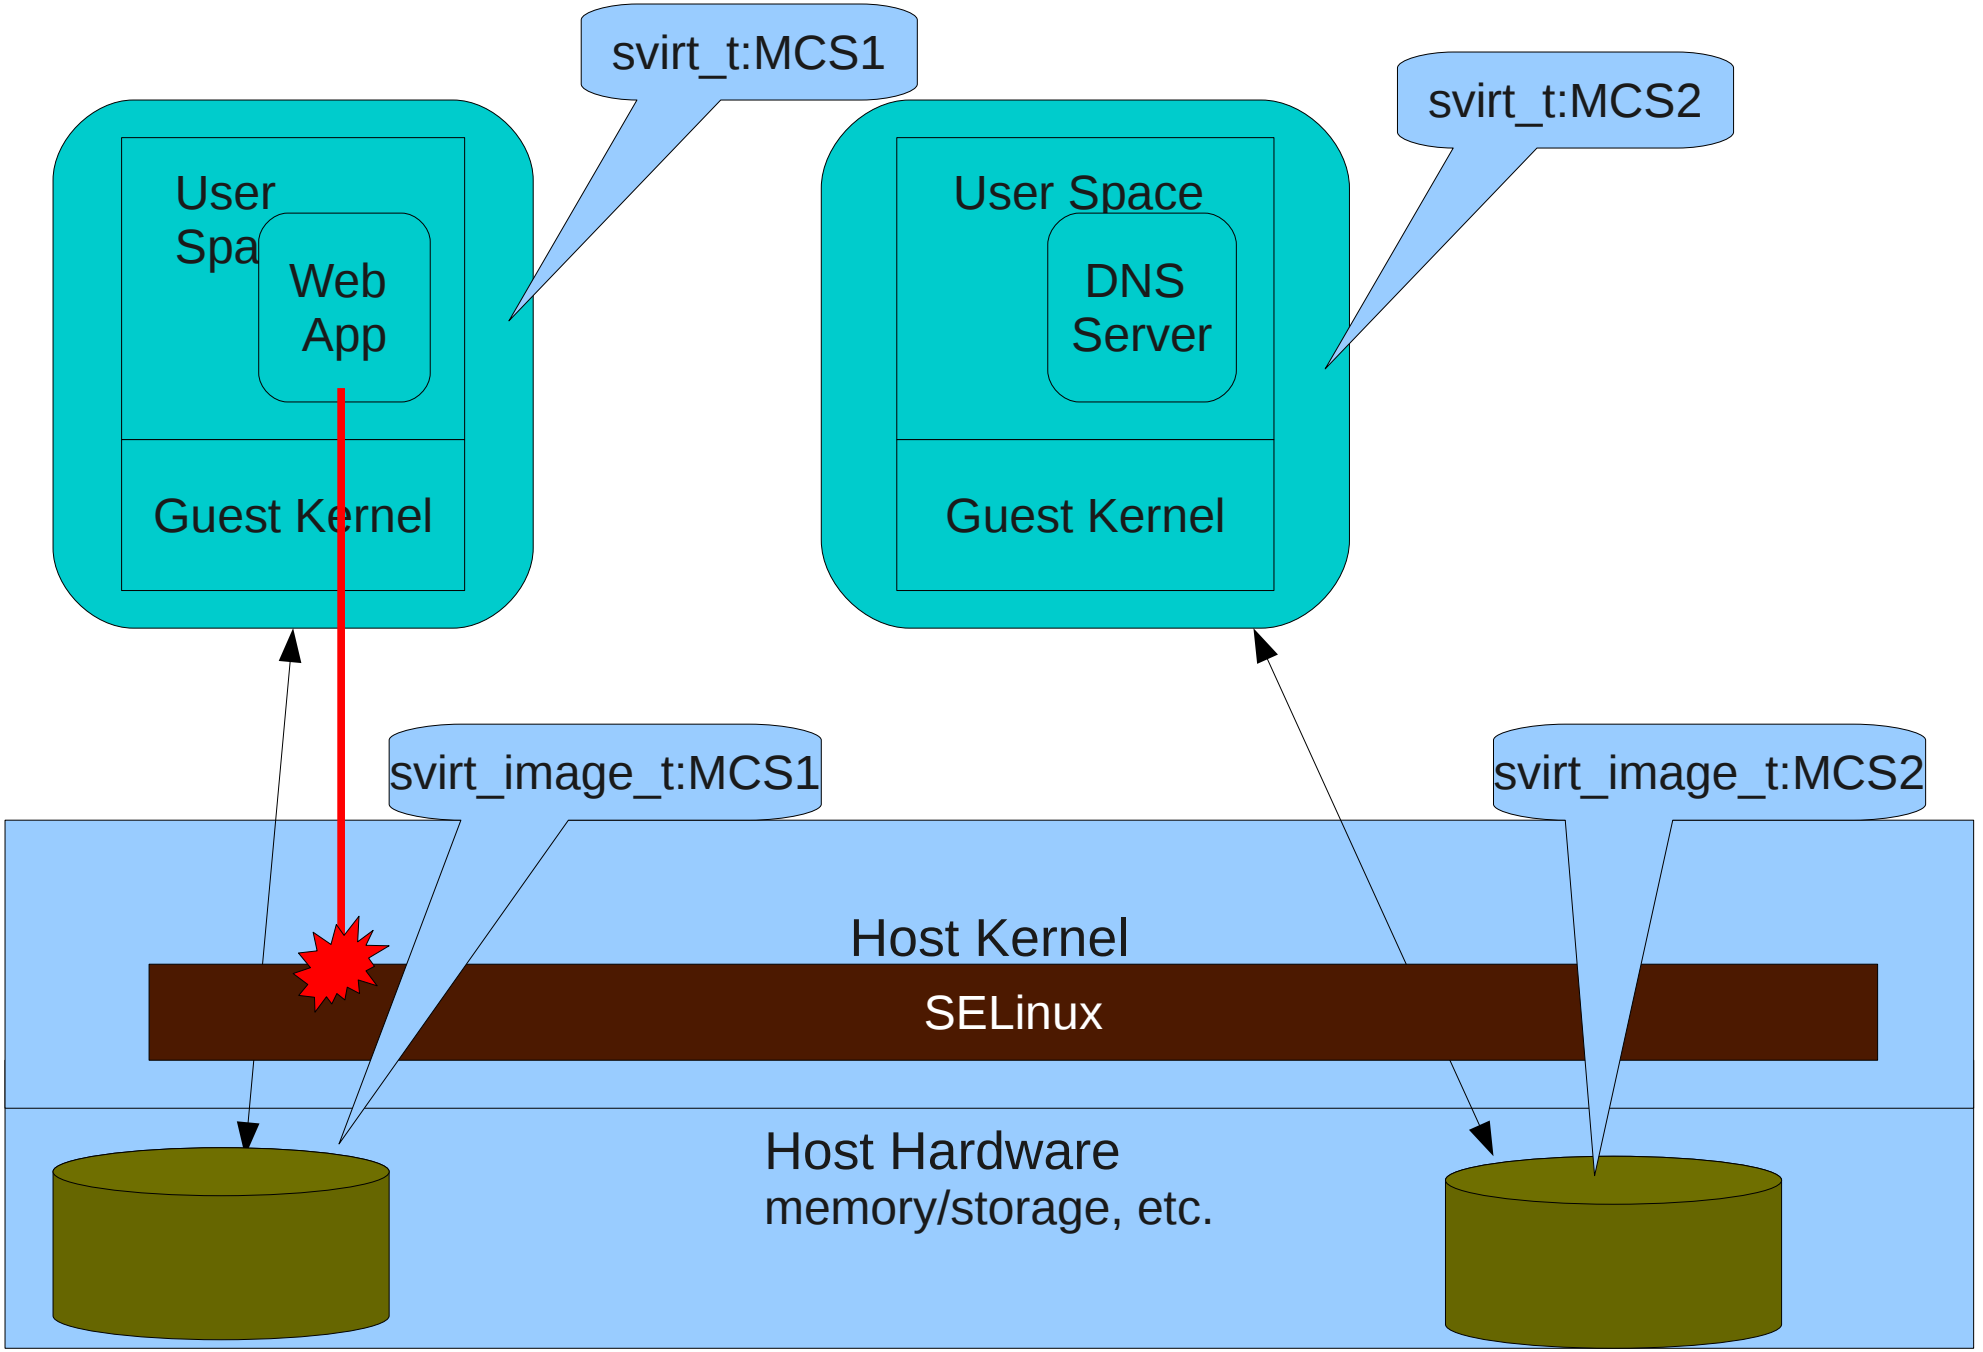

Libvirt – Static Label - MLS

Multi-Level Security

- Administrator must specify image label svirt_t:TopSecret
- Launches the image - svirt_t:TopSecret
- Libvirt will **NOT** label any content. Administrator responsible for labeling content.

Virt Manager

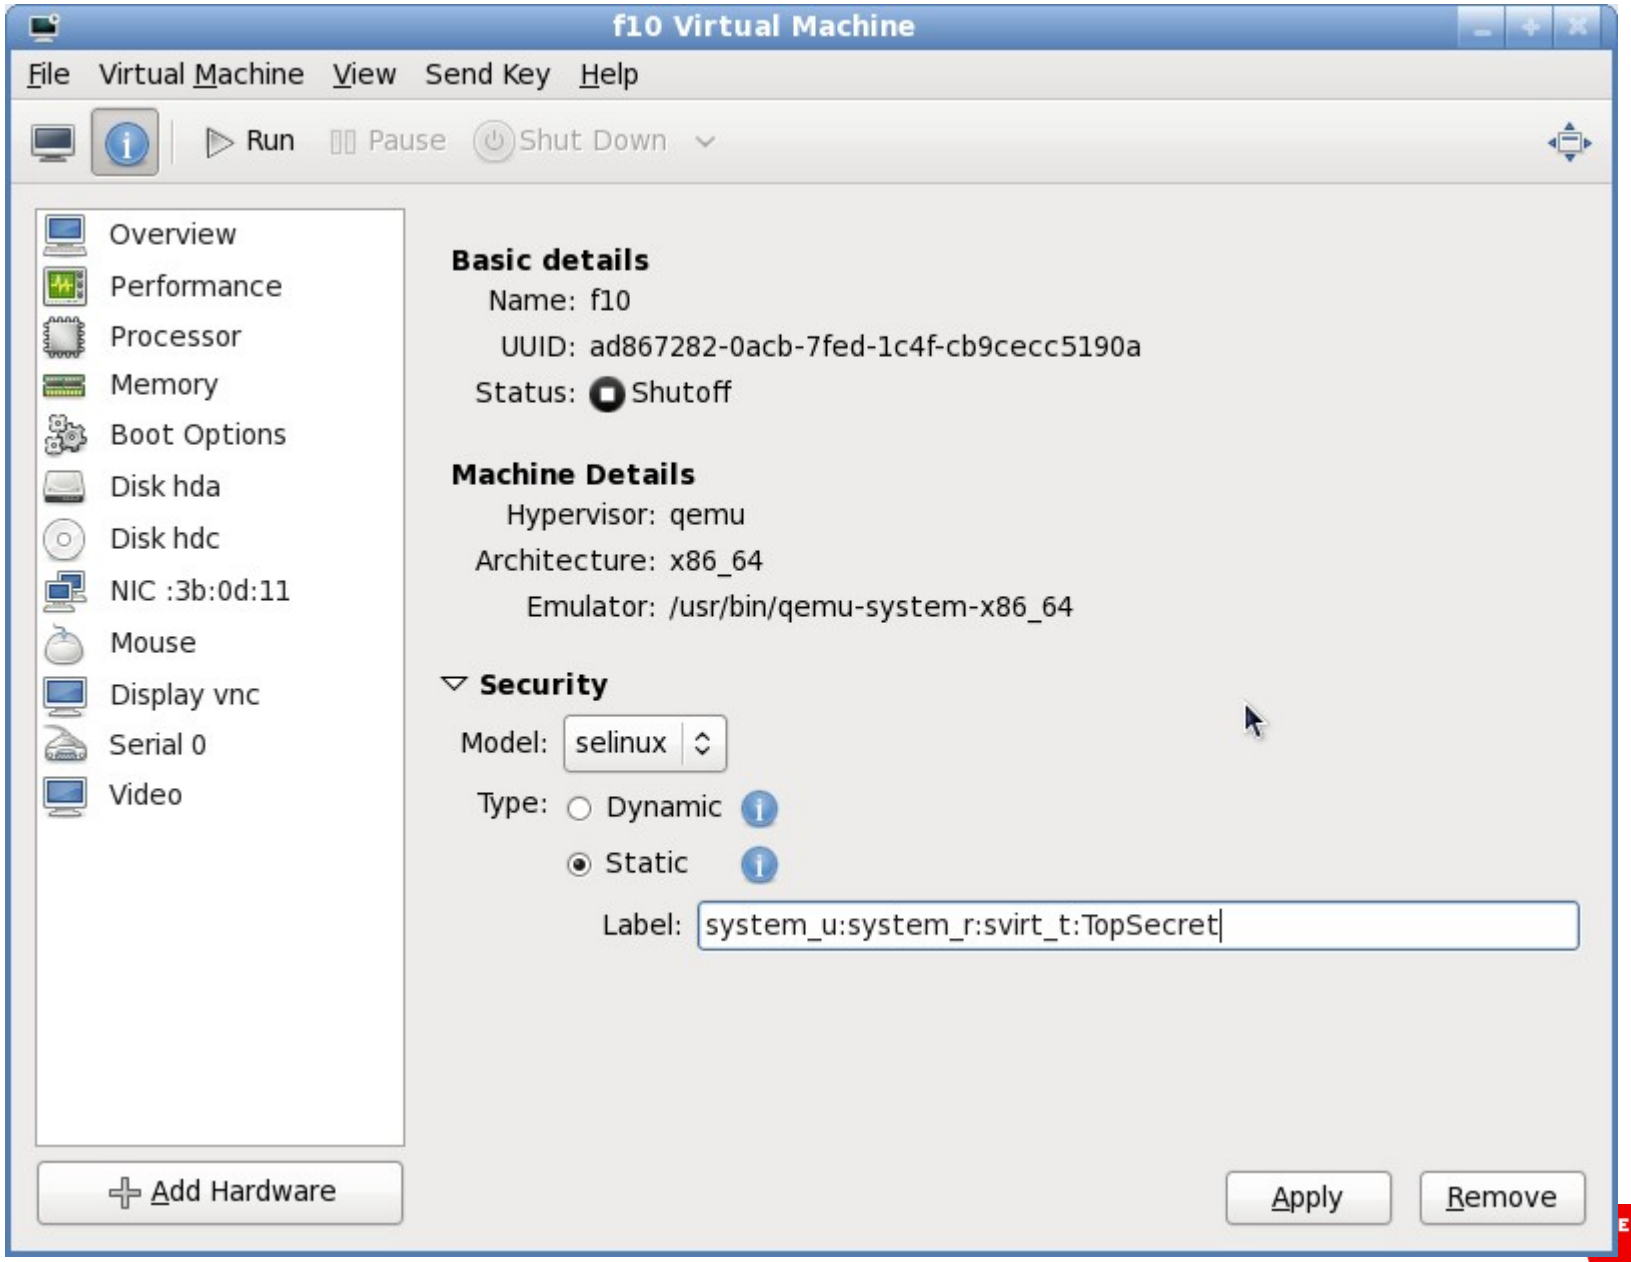

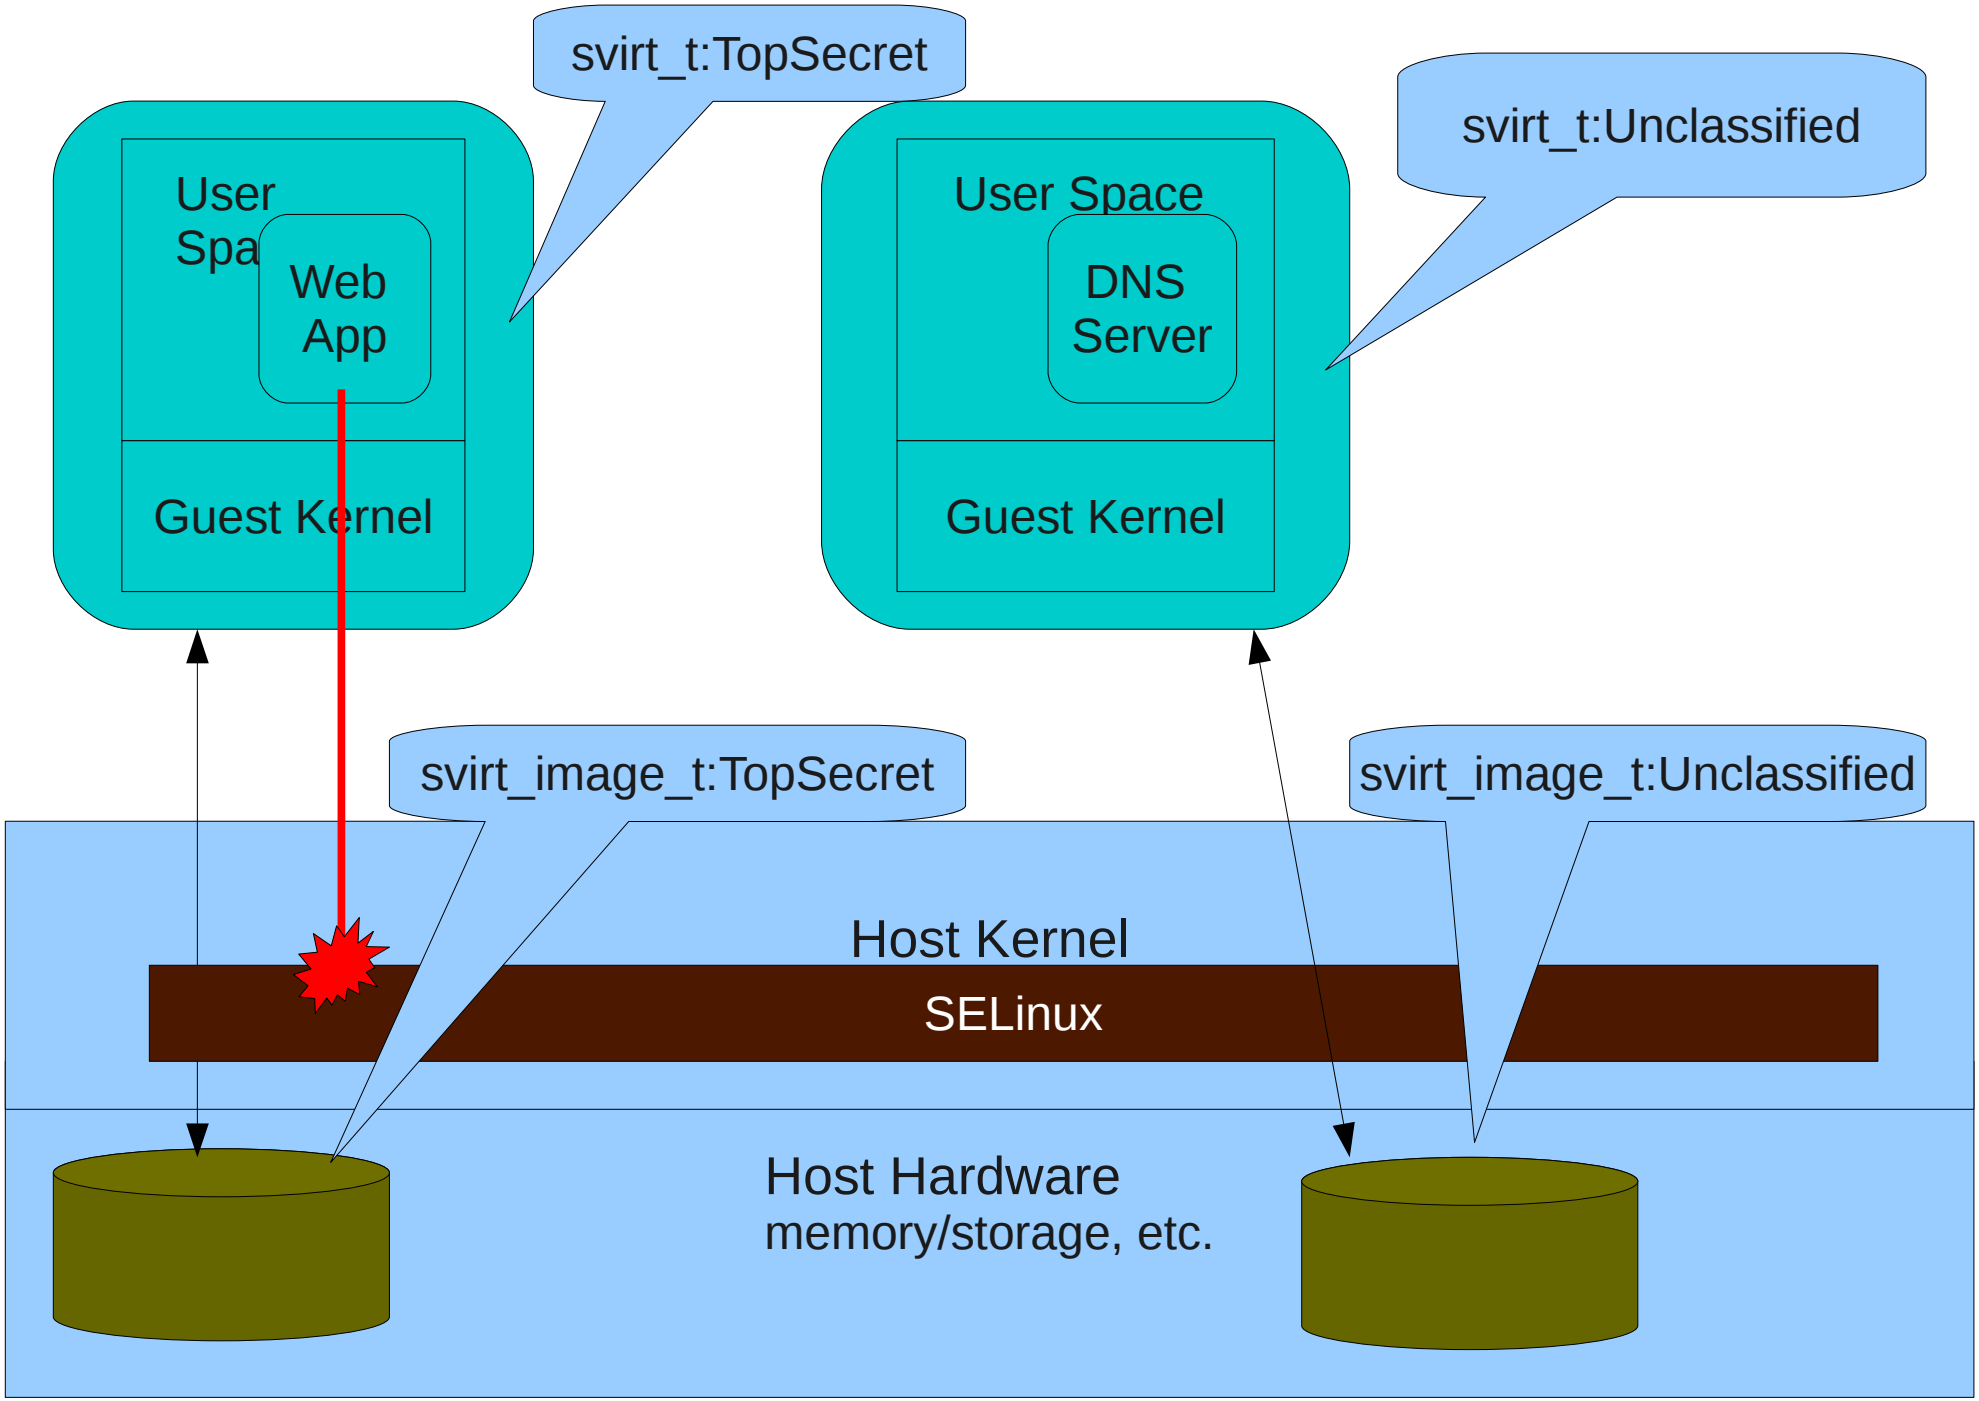
svirt_t:TopSecret
svirt_t:Unclassified
User Spa
Web App
Guest Kernel
User Space
DNS Server
Guest Kernel
svirt_image_t:TopSecret
svirt_image_t:Unclassified
Host Kernel
SELinux
Host Hardware
memory/storage, etc.

DEMO

Future Enhancements

- Different Types for confined guest
  - svirt_web_t – type
    - only allow a guest virtual machine to listen on web ports
  - Confine a Windows 2003 box to only run as a ISS server
    - If corrupted it could not become a Spam Bot.

sVirt Project Page

- http://selinuxproject.org/page/SVirt