PRESENTED BY RED HAT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.

www.theredhatsummit.com

Not your Grandfather's SELinux!

Daniel J Walsh
SELinux Lead Engineer
Red Hat


Wait, I'm using SELinux?

Top new risks being mitigated by SELinux.

Red Hat Enterprise Linux 6 introduces many new SELinux capabilities.

- Virtualization
- Users
- Enabling the Administrator
- Desktop Applications
- Other...

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 their potential usefulness in other cases.

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

2 The nature of the vulnerability

The Challenges posed by SELinux are taken into consideration.

Before Virtualization

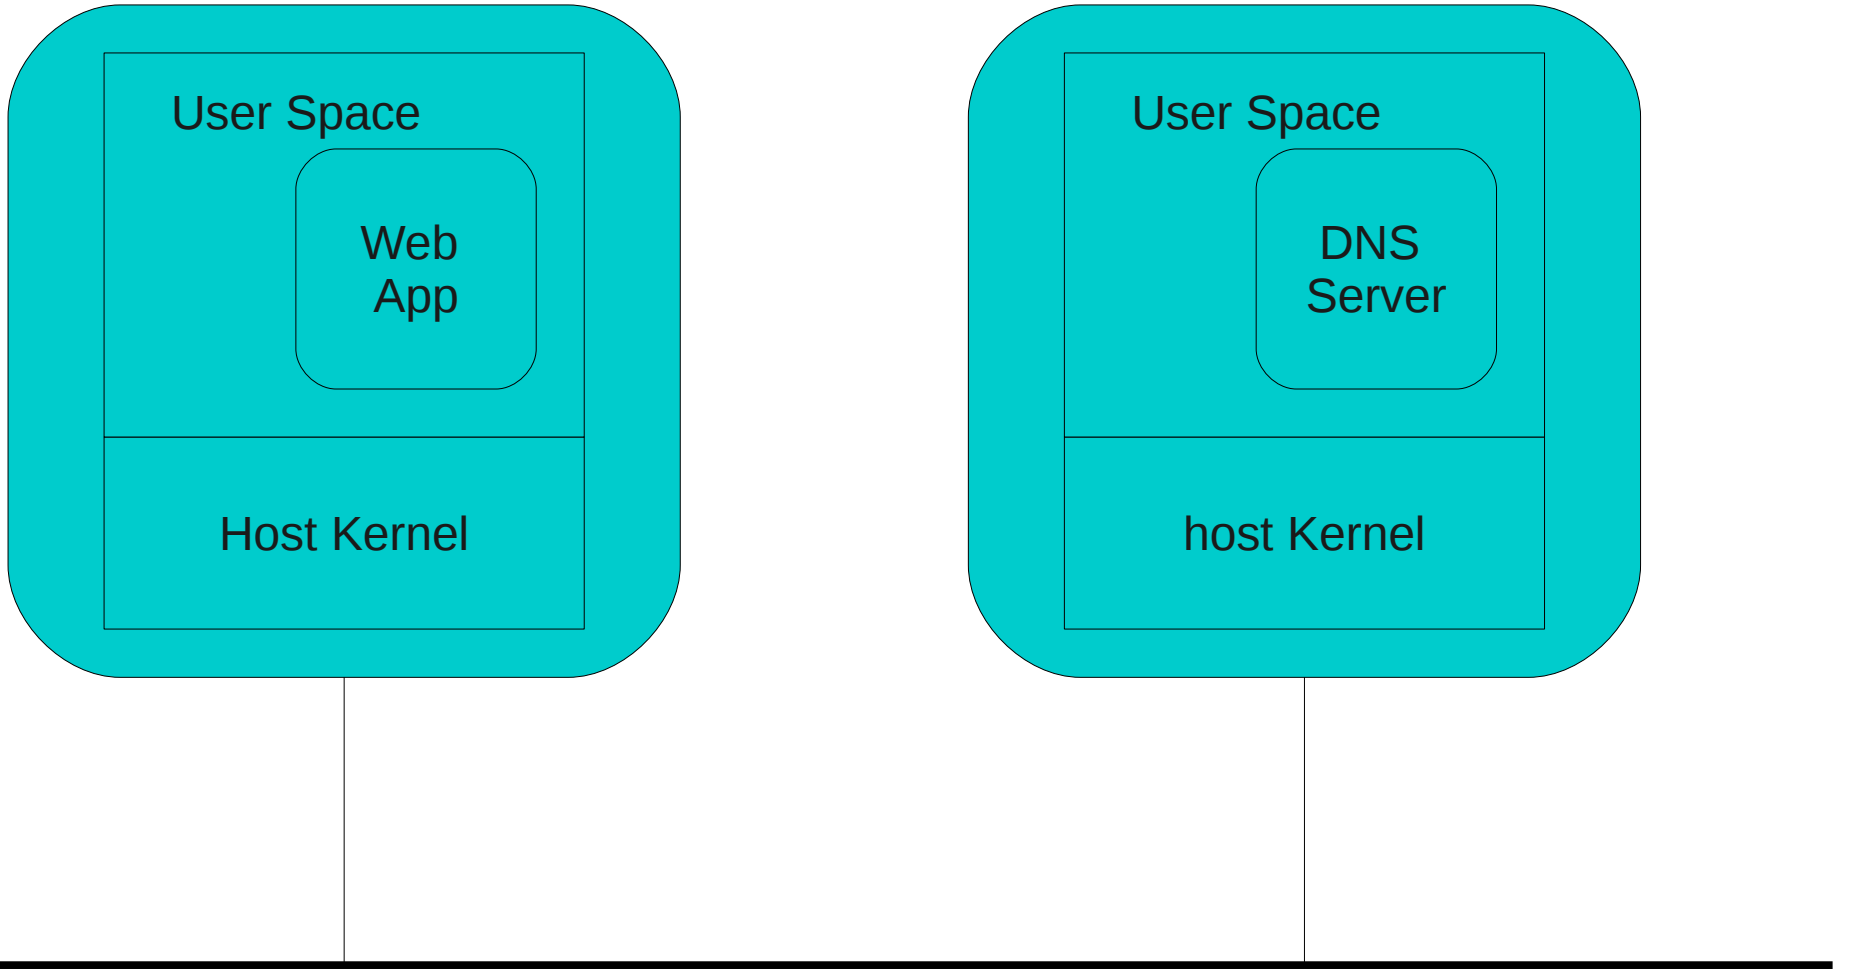

After Virtualization

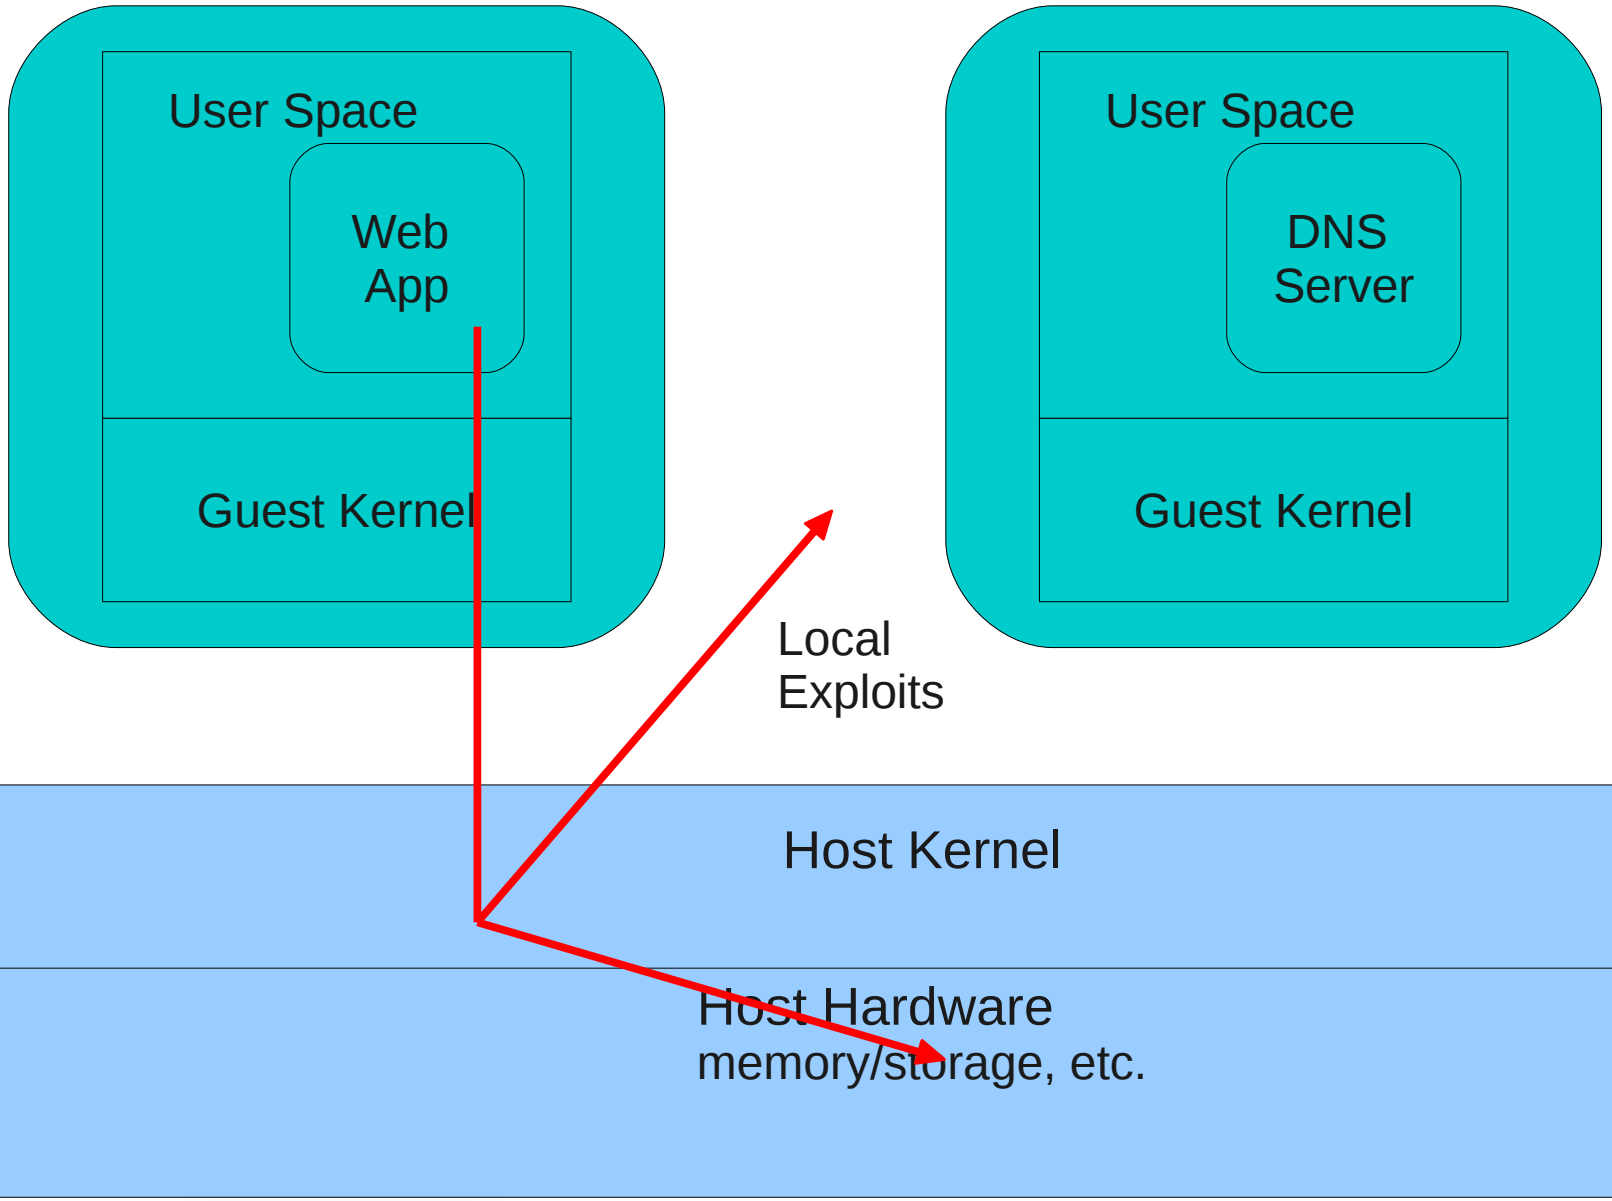

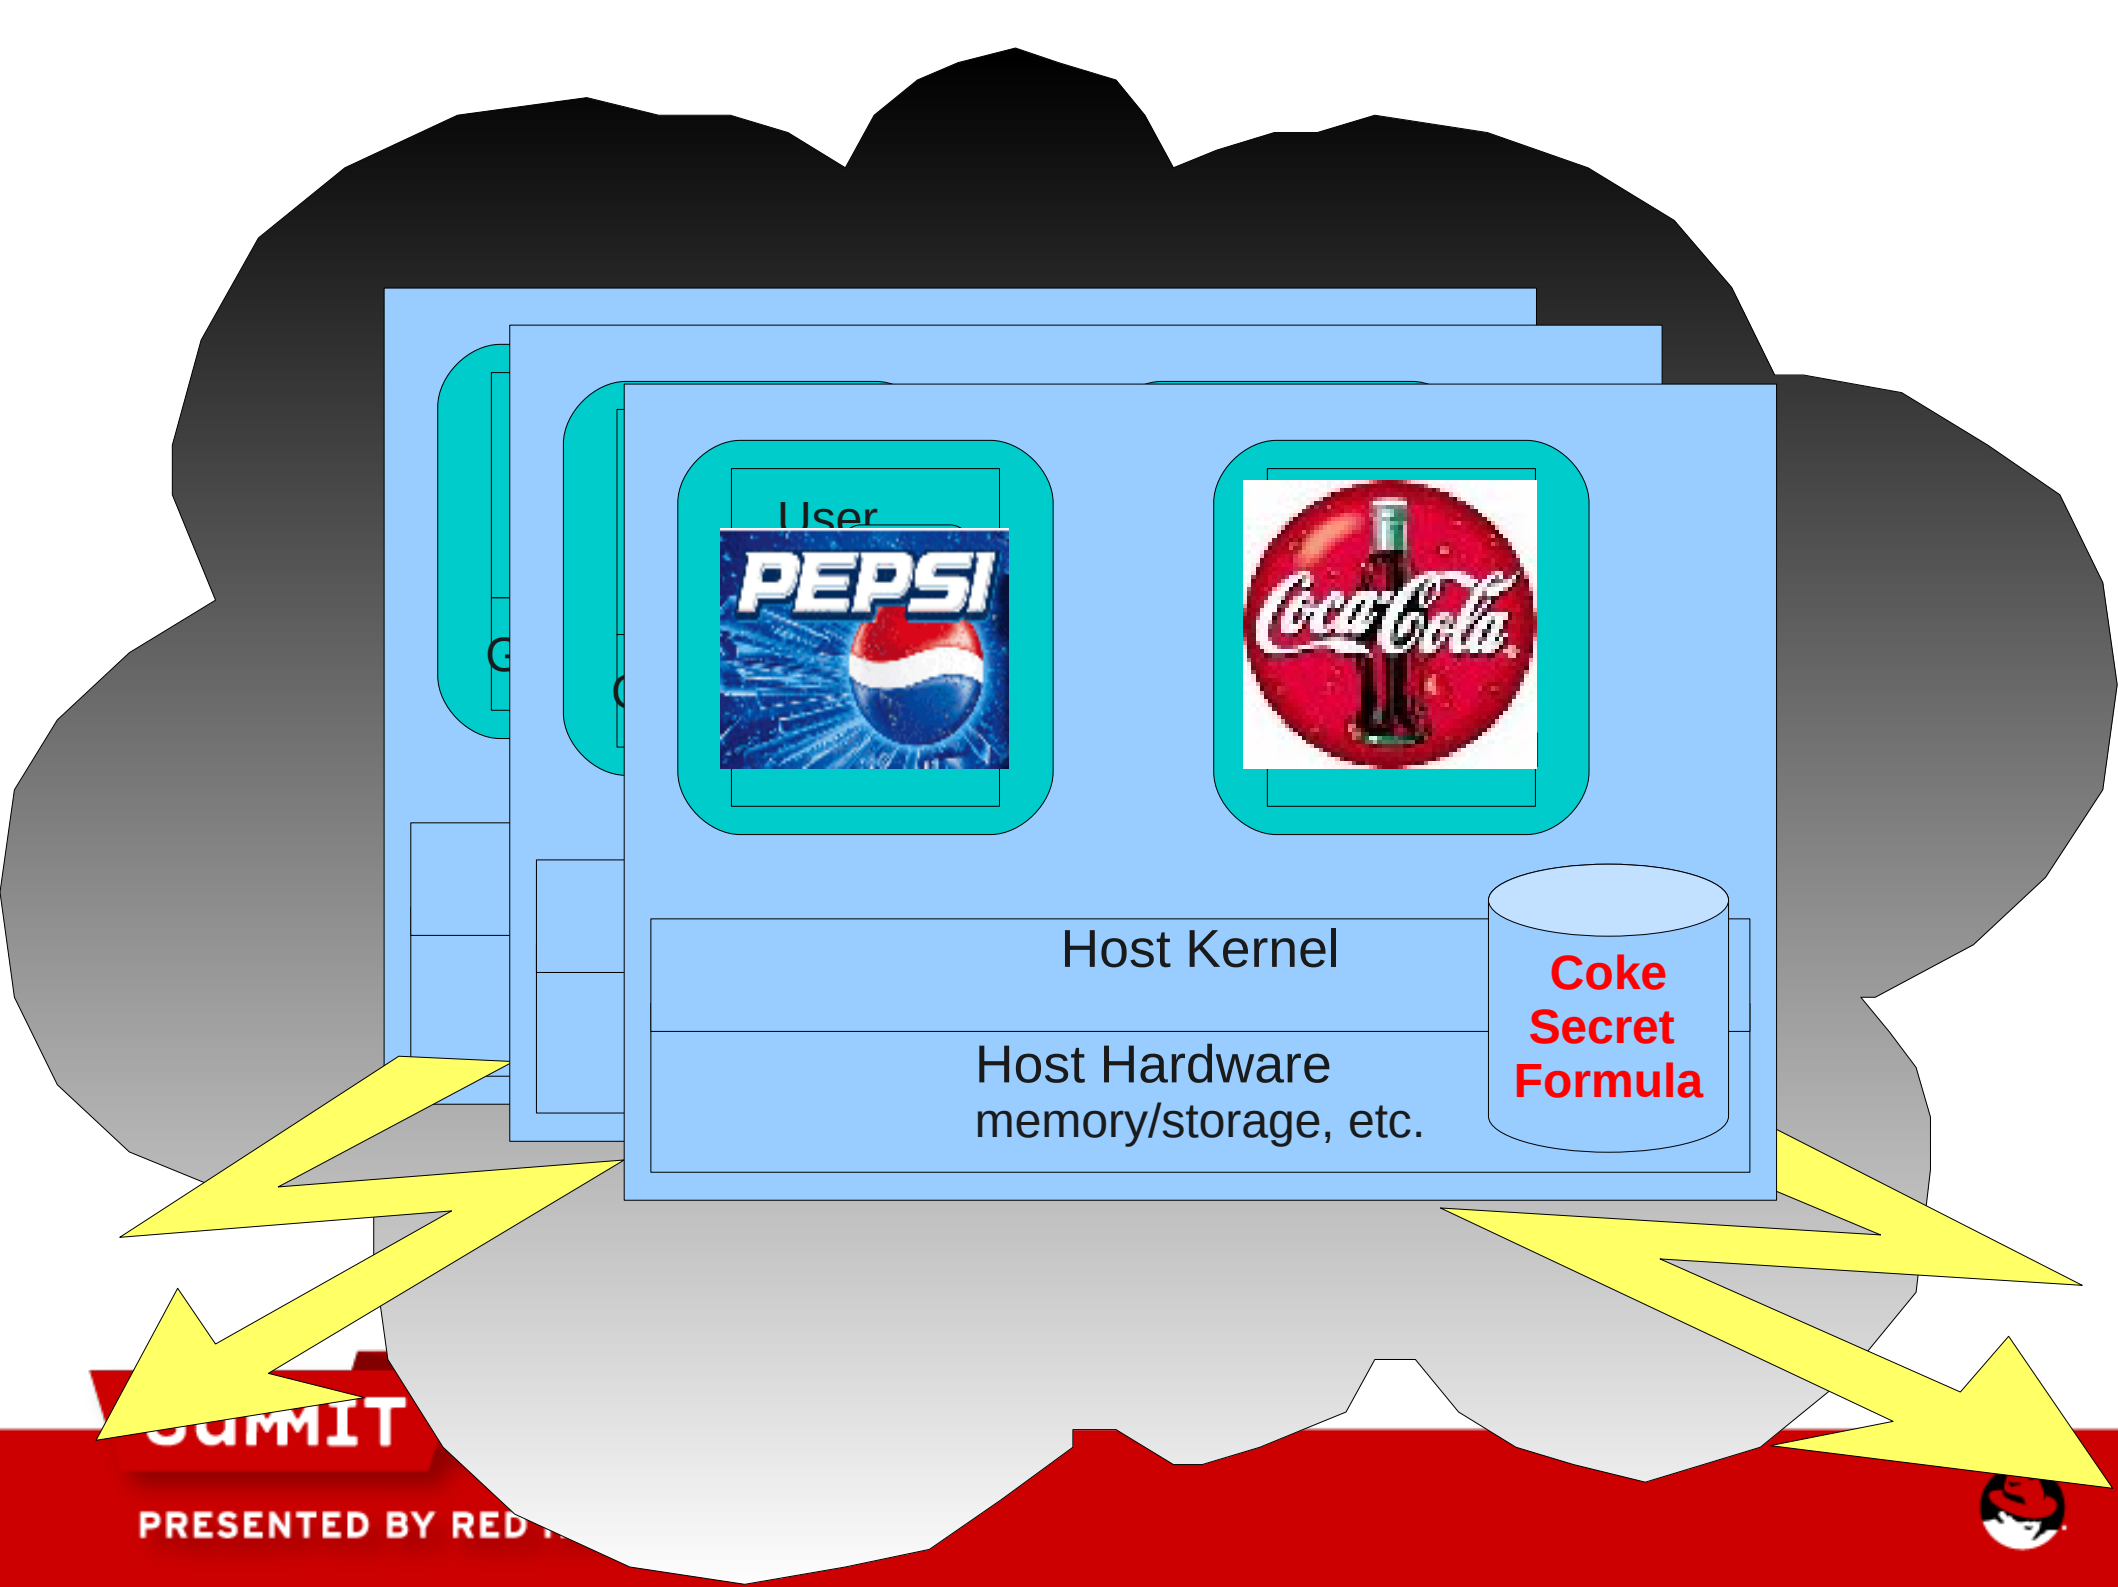
User
PEPSI
Host Kernel
Host Hardware
memory/storage, etc.
Coke
Secret
Formula

Enter SELinux..

SELinux is all about labeling

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

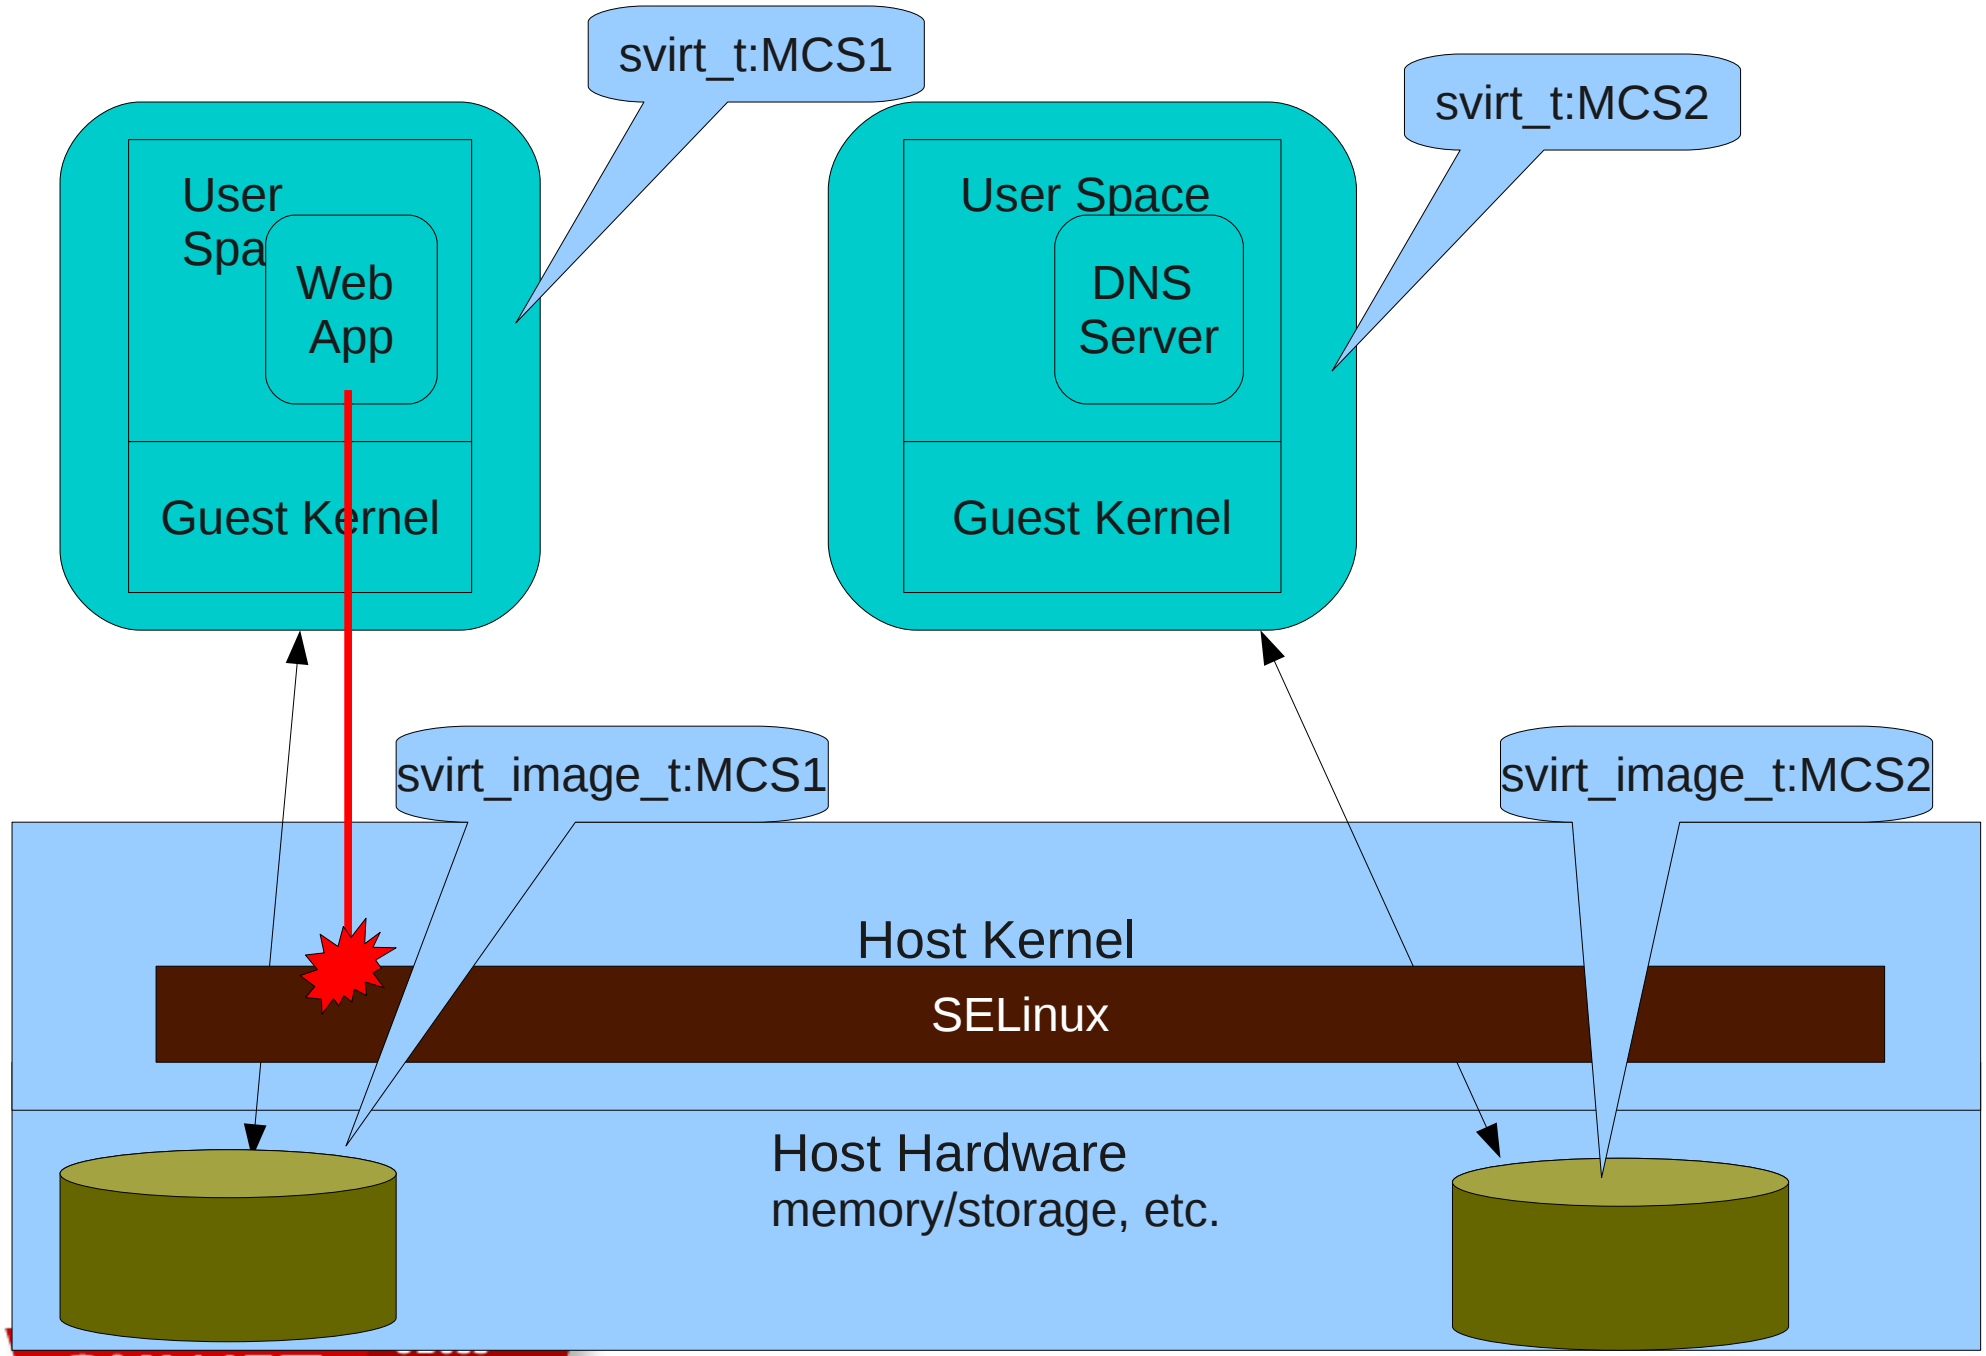
svirt_t:MCS1
svirt_t:MCS2
User Spa
Web App
Guest Kernel
User Space
DNS Server
Guest Kernel
svirt_image_t:MCS1
svirt_image_t:MCS2
Host Kernel
SELinux
Host Hardware
memory/storage, etc.

Vulnerability - Users

- Computers would be a lot more secure if we could eliminate the users...

Dave Bowman: Hello, HAL. Do you read me, HAL?
HAL: Affirmative, Dave. I read you.
Dave Bowman: Open the pod bay doors, HAL.
HAL: I'm sorry, Dave. I'm afraid I can't do that.
Dave Bowman: What's the problem?
HAL: I think you know what the problem is just as well as I do.
Dave Bowman: What are you talking about, HAL?
HAL: This mission is too important for me to allow you to jeopardize it.
- 2001: A Space Odyssey

Confining Users - RHEL5

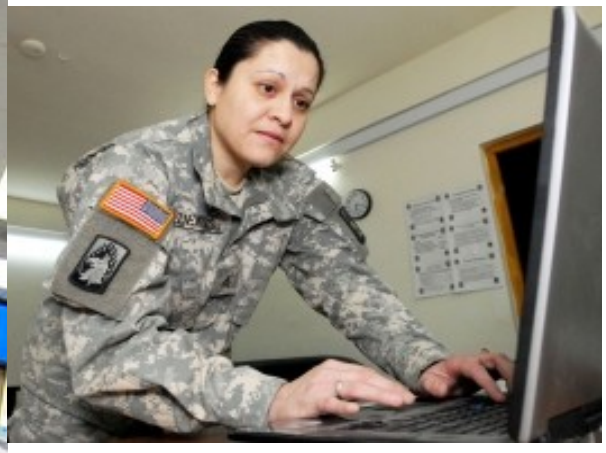

Default SELinux User Types - RHEL6

- Terminal user/ssh - guest_t
  - No Network, No setuid, no exec in homedir
- Browser user/kiosk - xguest_t
  - Web access ports only. No setuid, no exec in homedir
- Full Desktop user - User_t
  - Full Network, No SETUID.
- Confined Admin/Desktop User - Staff_t
  - Full Network, sudo to admin only, no root password. Usually a confined admin
- Unconfined user - unconfined_t (Default)
  - SELinux does not block access.

Confining Users - RHEL6

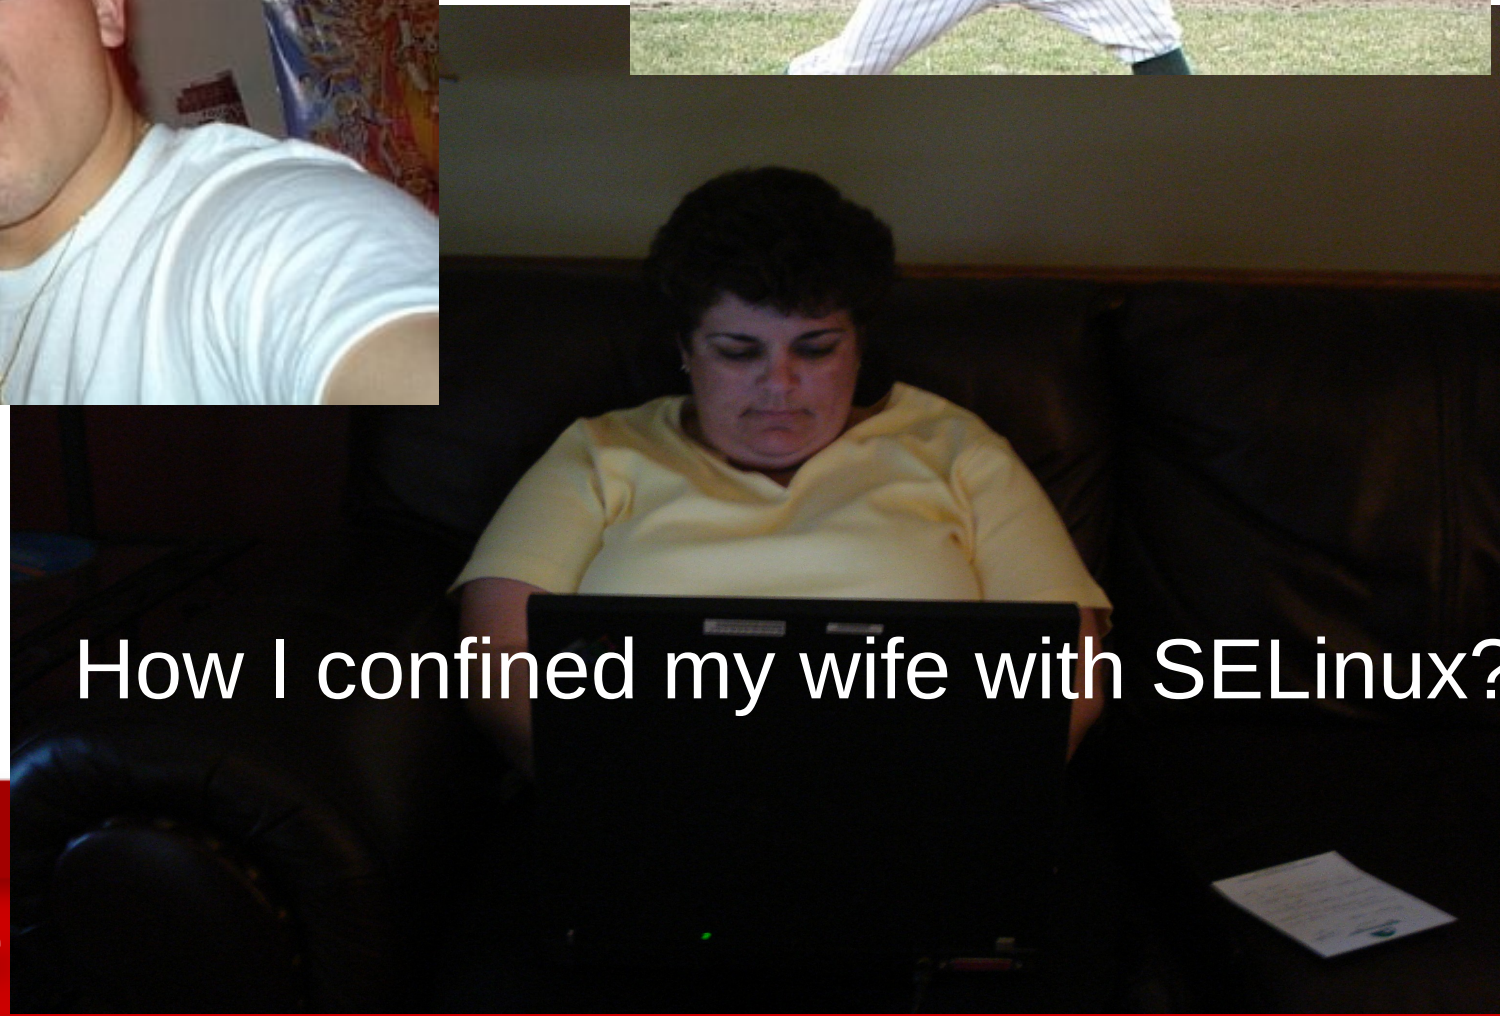

Red Hat Non-Technical Staff

Bosses

CEOs?

Not Likely...


HOLY CROSS
I even confine this guy!!!

Unknown/Untrusted Users?
Kiosk Mode

Vulnerabilities

- Application vulnerability triggered by content can cause application to do bad things.
  - tcpdump vulnerability CVE-2007-3798
  - DOD – Crossed Domain Solutions
  - Web Browser Vulnerabilities
  - Adobe Reader Vulnerabilities
  - OpenOffice
  - How can I trust my machine to be used for grid jobs?

SELinux

- SELinux difficult to use on random applications.
  - Transitions process to locked down environment
  - Policy needs to be written
  - Somewhat hard coded
  - Does not lend it self easily to scripting
  - Processes with the same type can attack each other

Introducing sandbox tool

- Run any application in a locked down environment.
  - Block Network? Access to Processes? Access to files? Homedir? X? dbus?
- Run untrusted applications?
- Run filters on untrusted data?

Standard SELinux Sandbox

- Execution any app within SELinux Confinement
  - Blocks “open” call
  - Allows read/write on inherited file descriptors
  - Temporary storage allowed
- cat untrusted.txt | sandbox filter > trusted.txt

Standard SELinux Sandbox

- Uses MCS labels for separation
  - Based on same technology as svirt/libvirt
  - Apps have same types/access but can not interact.
- Excellent for scripting
  - Pipe apps read stdin/write stdout
- Confinement of grid jobs
  - Wrap grid jobs in sandbox wrapper

What about the desktop?

- How do I confine acroread?
- Large communications paths
  - X server
  - File system
    - Home Directory
    - /tmp
  - gconf
  - Dbus

Sandbox -X

- Creates and populates temporary $HOME and /tmp
  - Each sandbox has their own /tmp and $HOME
  - Applications create content within the new homedir
- Different X server per sandbox
- SELinux labels confinement like standard sandbox
- Temporary $HOME & /tmp deleted on exit

sandbox -X DEMO

SELinux Policy Changes

Category	RHEL5	RHEL6
Executable Types	327	630
Types	1787	3053
Roles	6	13
Allow Rules	124658	255093
Booleans	259	161
File Contexts	2366	3654

Why did Booleans go down?

- DOMAIN_disable_trans booleans removed
  - Bad Idea. Caused other problems with type transitions.
- RHEL6 introduces permissive domains
  - Domain is not blocked by SELinux, AVC still generated
  - semanage permissive -a httpd_t
    - Turns httpd_t (apache) to a permissive domain

SELinux Policy Changes

- Policy updates for
  - RHEL5, Fedora 7,8,9,10,11,12,13
  - Ubuntu, Debian, Gentoo
- open permission
- mac_admin permission
  - Livecd:
    - Support for building images with "foreign" policy.

Multi Level Security

- Desktop
  - X/ACE
- Multi-level virtualization
- mcstransd updated - complex label encodings
- Networking Controls
  - Removal of legacy network controls
  - New unified peer controls for xfrm and netlabel
  - New ingress/egress/forward controls.
  - Labeled networking/access control enhancements

Setroubleshoot Improvements

- New gui, bug reporting, fix it button
- Limited intruder analysis
- New Plugins

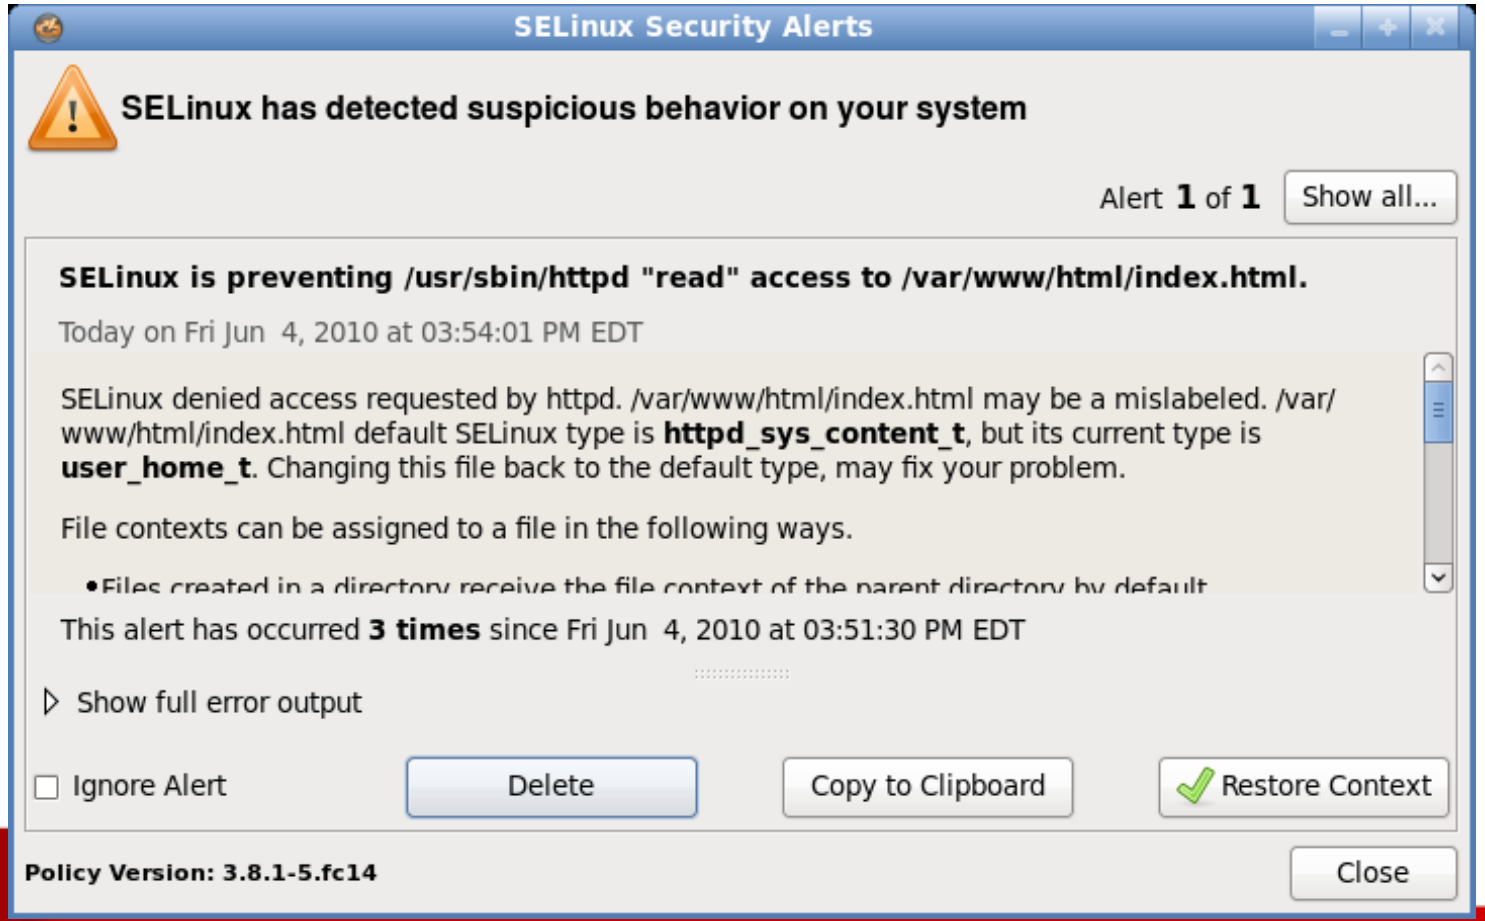

Tool chain improvements

- File context equivalence.
  - semanage fcontext -a -e /var/www /src/web
- Policy module compression.
- Audit2allow analysis
  - Boolean support
  - Constraints
- Writing Policy
  - sepolgen /usr/sbin/rwhod
  - permissive domains

Resources

- Documentation
  - New user guide
  - Managing confined services
- External resources
  - http://www.selinuxproject.org
  - SELinux Notebook
    - http://www.freetechbooks.com/the-selinux-notebook-the-foundations-t785.html