

Red Hat Deep Dive Sessions

Linux on System z

Shawn Wells (swells@redhat.com)

W/W Lead Architect, Linux on System z

Team Lead, System z SMEs

Introduction

- **Shawn Wells** (swells@redhat.com)
Lead Architect, Linux on System z
Team Lead, System z SMEs

Phone: +1 443 534 0130

Agenda

- **Scheduled questions to be answered in this session:**
 - What's the Linux on System z development process?
 - What's in RHEL now? What's on the roadmap?
 - Provisioning & Patch Management in RHN
 - Security Update
 - SELinux, Audit, etc





Linux on System z Development Process

Linux on System z Development

Community

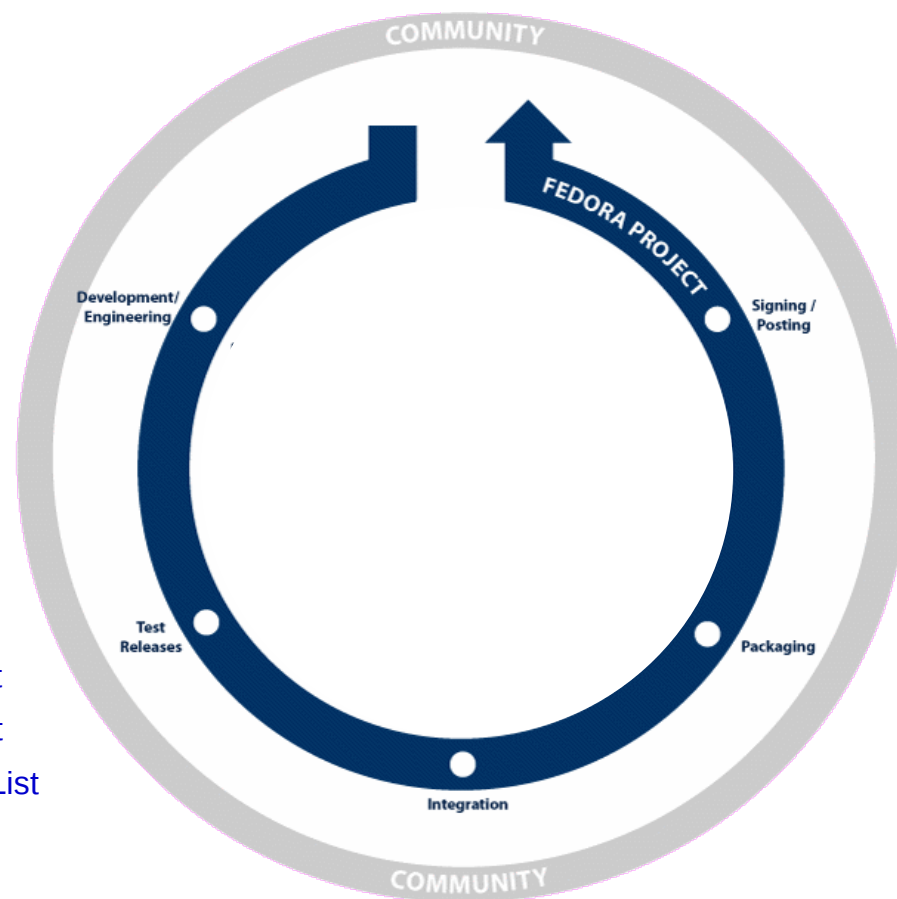
- Development with “upstream” communities
- Kernel, glibc, etc
- Collaboration with partners, IBM, open source contributors



Linux on System z Development

Fedora

- Bleeding Edge
- Sets direction for RHEL technologies
- Community Supported
- Released ~6mo cycles
- Fedora 8,9,10 = RHEL6



Fedora 8; <http://fedoraproject.org/wiki/Releases/8/FeatureList>

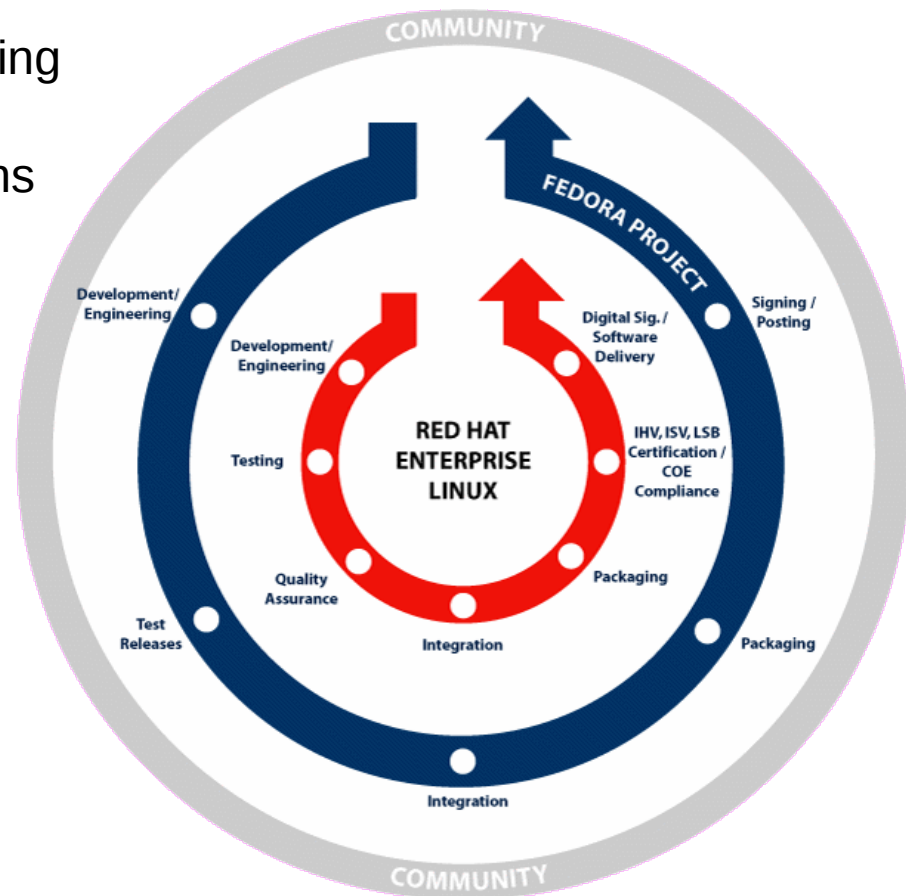
Fedora 9; <http://fedoraproject.org/wiki/Releases/9/FeatureList>

Fedora 10; <http://fedoraproject.org/wiki/Releases/10/FeatureList>

Linux on System z Development


Red Hat Enterprise Linux

- Stable, mature, commercial product
- Extensive Q&A, performance testing
- Hardware & Software Certifications
- 7yr maintenance
- Core ABI compatibility guarantee
- Major releases 2-3yr cycle



Support Cycle

Extended Product Lifecycle



	Production 1	Production 2	Production 3
Security Patches	X	X	X
Bug Fixes	X	X	X
Hardware Enablement	Full	Partial	None
Software Enhancements	X		

Linux on System z Subscriptions

- No Upgrade Costs
- No Client Access Fee
- Unlimited Support Incidents

For System z:

- Priced Per IFL
- Unlimited VMs per IFL

Customers can consolidate subscriptions *to or from* other platforms

Red Hat Enterprise Linux Subscription

S U P P O R T	PREMIUM	24x7 Phone/Web 1 Hour SLAs
	STANDARD	Phone/Web 1-4 Business Hour SLA
	BASIC	Web Support. 2 Day SLA

Security, Bug Fixes
Regular H/W & S/W Updates

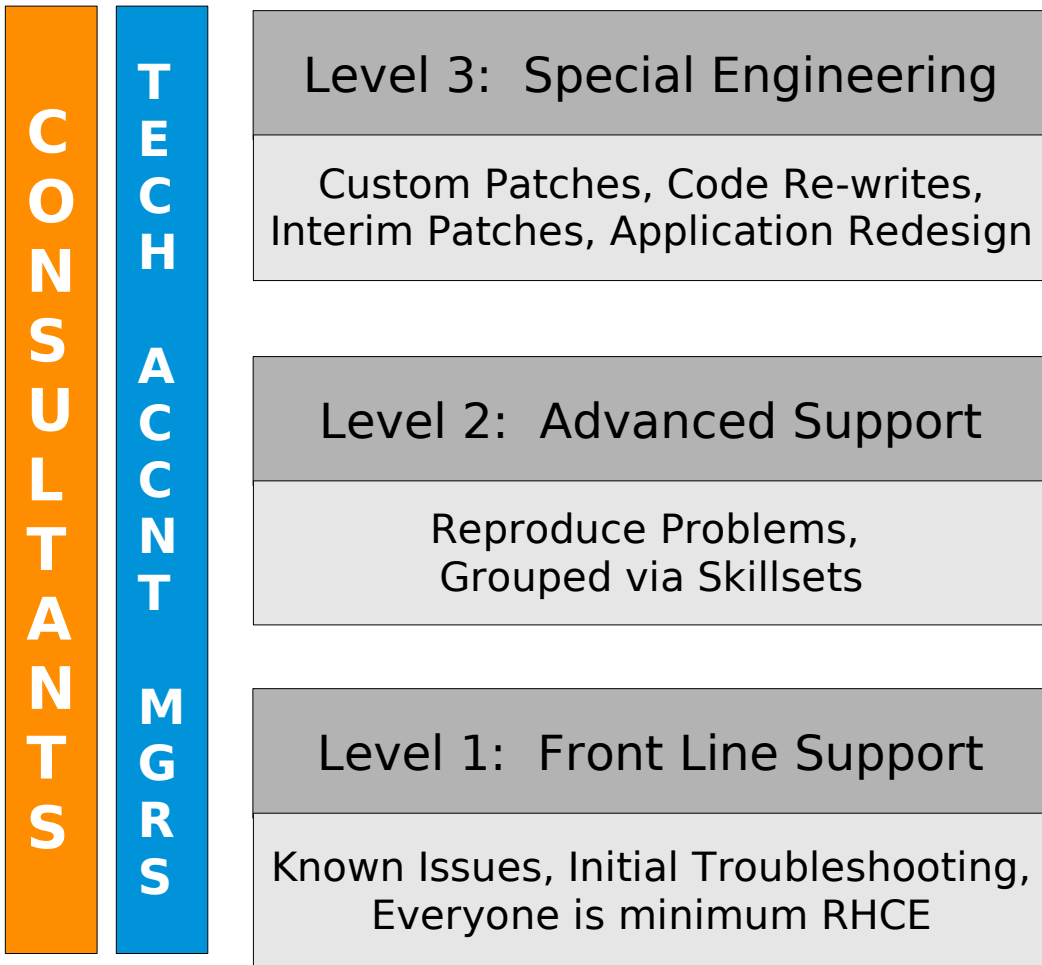
Hardware & Application Certifications

Stable Application Interfaces

Upgrades to New Versions

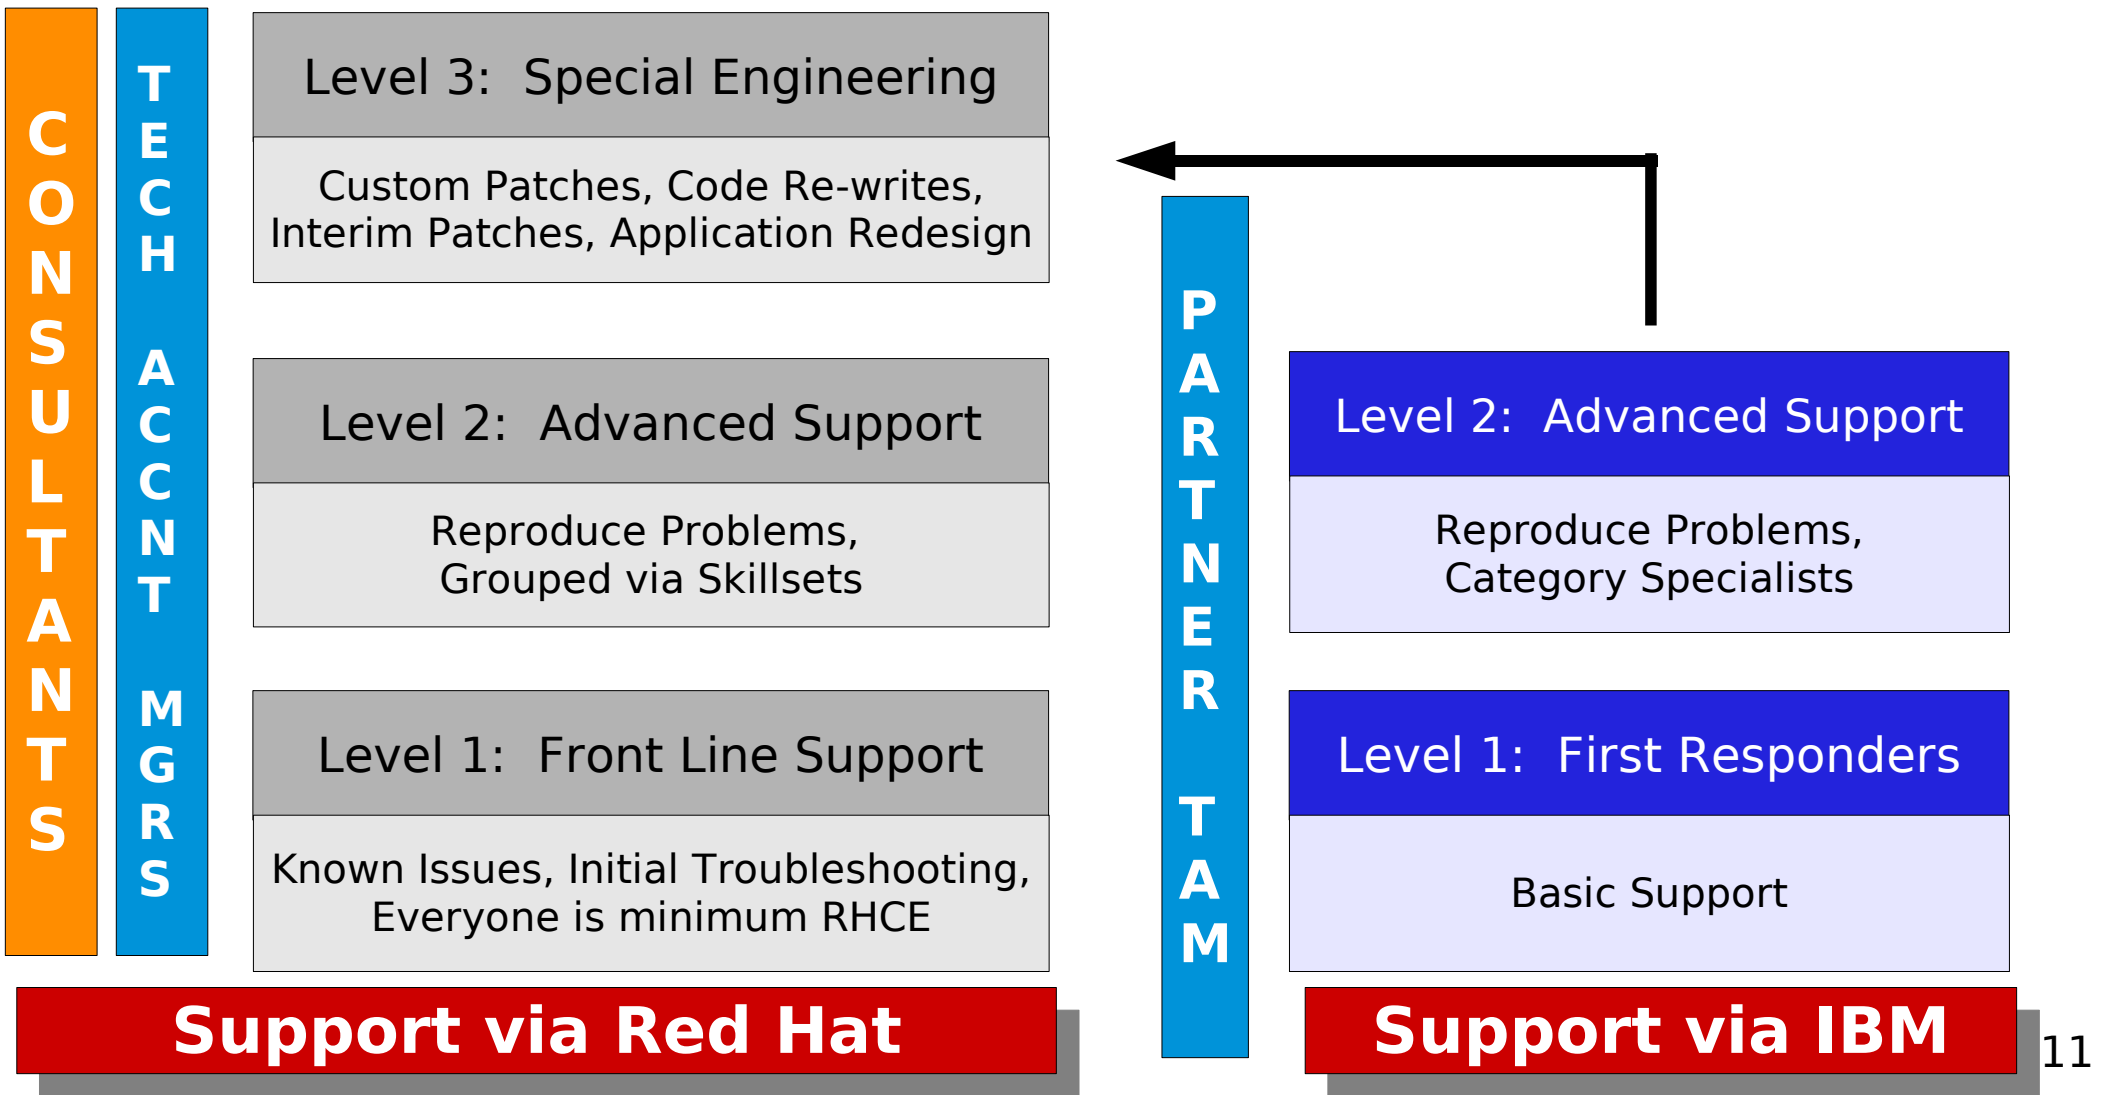
Product Source & Binaries

Linux on System z Support



Support via Red Hat

Linux on System z Support






What's in RHEL now?
What's on the road map?

the proof
is in
the
pudding...



IBM Changes to 2.6.x Kernel



RHEL Now: RHEL 5.2

- Support for z10
- Dynamic CHPID reconfiguration
- Improved “ssh -X” with VPN during installation process
- Better network performance with skb scatter-gather support
- Implementation of SCSI dump infrastructure

RHEL Now: RHEL 5.2

- Accelerated in-kernel Crypto
 - Support for crypto algorithms of z10
 - SHA-512, SHA-384, AES-192, AES-256

- Two OSA ports per CHPID; Four port exploitation
 - Exploit next OSA adapter generation which offers two ports within one CHPID. The additional port number 1 can be specified with the qeth sysfs-attribute “portno”
 - Support is available only for OSA-Express3 GbE SX and LX on z10, running in LPAR or z/VM guest (PFT for z/VM APAR VM64277 required!)

RHEL Now: RHEL 5.2

■ Large Page Support

- This adds hugetblfs support on System z, using both hardware large page support if available, and software large page emulation (with shared hugetblfs pagetables) on older hardware

■ skb scatter-gather support for large incoming messages

- This avoids allocating big chunks of consecutive memory and should increase networking throughput in some situations for large incoming packets

Full Release Notes At: [redhat.com](http://www.redhat.com)

http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Release_Notes/s390x/index.html

RHEL Now: RHEL 5.2

- Lightweight userspace priority inheritance (PI) support for futexes, useful for realtime applications (2.6.18)
 - Assists priority inversion handling. Ref: <http://lwn.net/Articles/178253/>
- High resolution timers (2.6.16)
 - Provide fine resolution and accuracy depending on system configuration and capabilities - used for precise in-kernel timing
- New Pipe implementation (2.6.11)
 - 30-90% perf improvement in pipe bandwidth
 - Circular buffer allow more buffering rather than blocking writers
- "Big Kernel Semaphore": Turns the Big Kernel Lock into a semaphore
 - Latency reduction, by breaking up long lock hold times and adds voluntary preemption

RHEL Now: RHEL 5.2

- Process Events Connector (2.6.15)
 - Reports fork, exec, id change, and exit events for all processes to userspace
 - Useful for accounting/auditing (e.g. ELSA), system activity monitoring, security, and resource management
- kexec & kdump (2.6.13)
 - Provide new crash-dumping capability with reserved, memory-resident kernel
- Extended device mapper multipath support
- Address space randomization:
 - Address randomization of multiple entities – including stack & mmap() region (used by shared libraries) (2.6.12; more complete implementation than in RHEL4)
 - Greatly complicates and slows down hacker attacks
- Audit subsystem
 - Support for process-context based filtering (2.6.17)
 - More filter rule comparators (2.6.17)

RHEL Now: RHEL 5.2

- Add nf_conntrack subsystem: (2.6.15)
 - Common IPv4/IPv6 generic connection tracking subsystem
 - Allows IPv6 to have a stateful firewall capability (not previously possible)
 - Increased security
 - Enables analysis of whole streams of packets, rather than only checking the headers of individual packets

- SELinux per-packet access controls
 - Replaces old packet controls
 - Add Secmark support to core networking
 - Allows security subsystems to place security markings on network packets (2.6.18)

RHEL Tomorrow: RHEL 5.3

- Currently in beta
 - Interested in being a beta tester?

- NSS

- CPU Affinity

- ETR Support

- Device-multipath support for xDR
 - RHT BugZilla: [184770](#)
 - IBM LTC 18425-62140

RHEL Tomorrow: Fedora

Fedora is Red Hat's bleeding edge, an incubator for new technologies and features

Fedora sets our direction for Red Hat Enterprise Linux, and gives you a good idea of what will be in our next RHEL release (... and in other Linux distros, too)

Fedora 8; <http://fedoraproject.org/wiki/Releases/8/FeatureList>

Fedora 9; <http://fedoraproject.org/wiki/Releases/9/FeatureList>

Fedora 10; <http://fedoraproject.org/wiki/Releases/10/FeatureList>

Fedora 8,9,10 = RHEL6

RHEL Tomorrow: “In Place” Upgrade


Currently a beta feature in RHEL 5.3

“In Place” Upgrades: preupgrade

- Will download files needed to upgrade,
 - Store them locally on disk
 - Reboot you into the installer
 - Not a true in-place upgrade (yet)!
-
- Benefit
 - The longest part of an install is when packages are downloaded to the local machine
 - Pre-Upgrade downloads and stores packages locally, while the machine is running/in production
 - Reboot directly into the installer

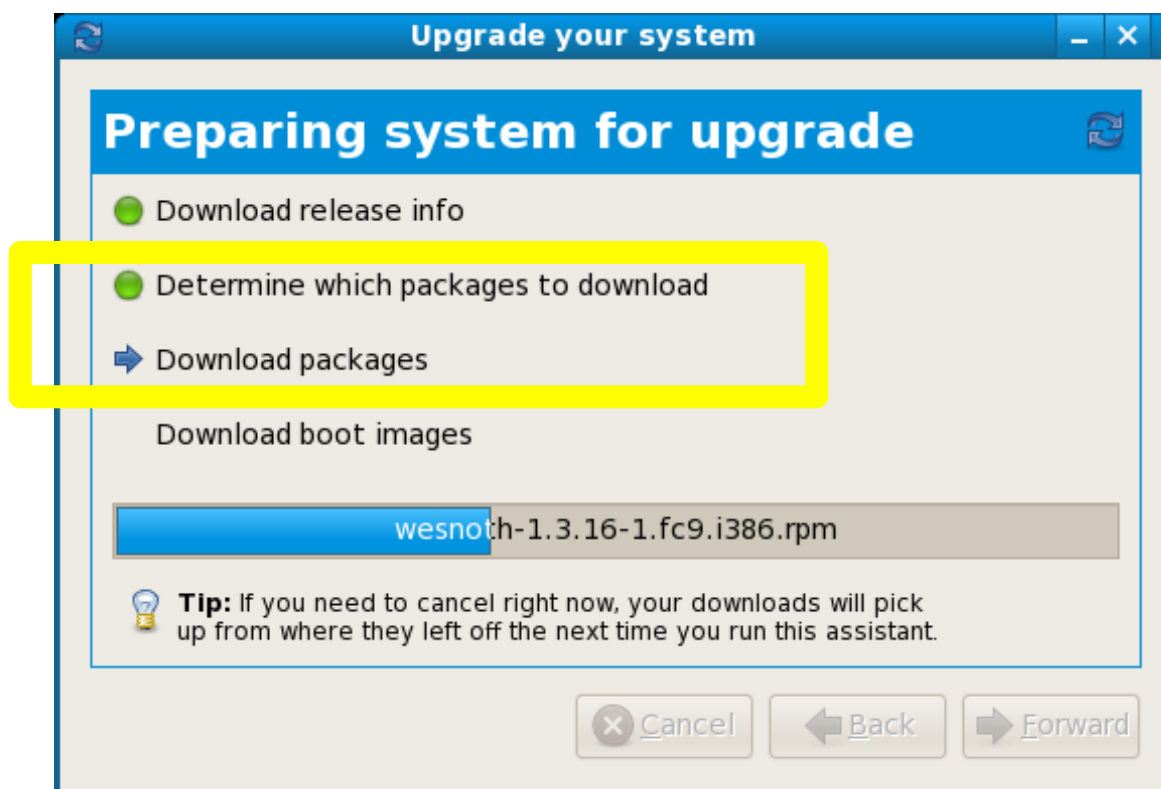
RHEL Tomorrow: “In Place” Upgrade

- Select Target Version



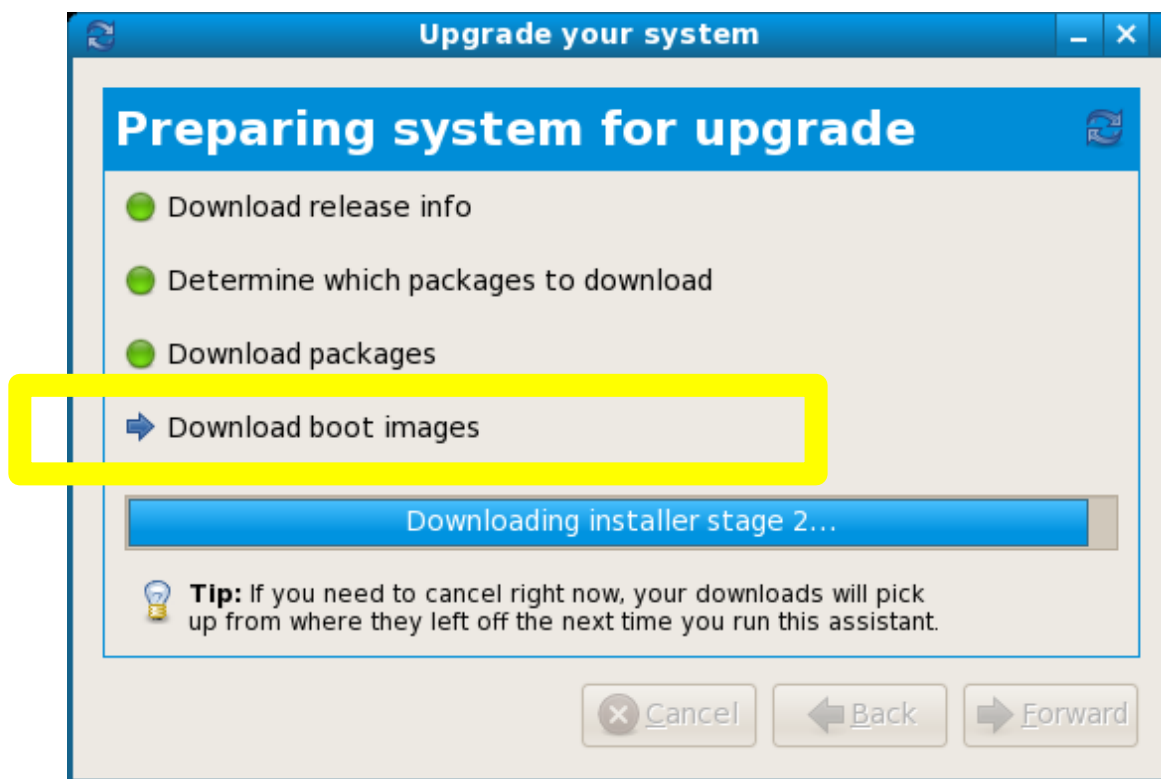
RHEL Tomorrow: “In Place” Upgrade

- Determines which packages need upgrading, and downloads them



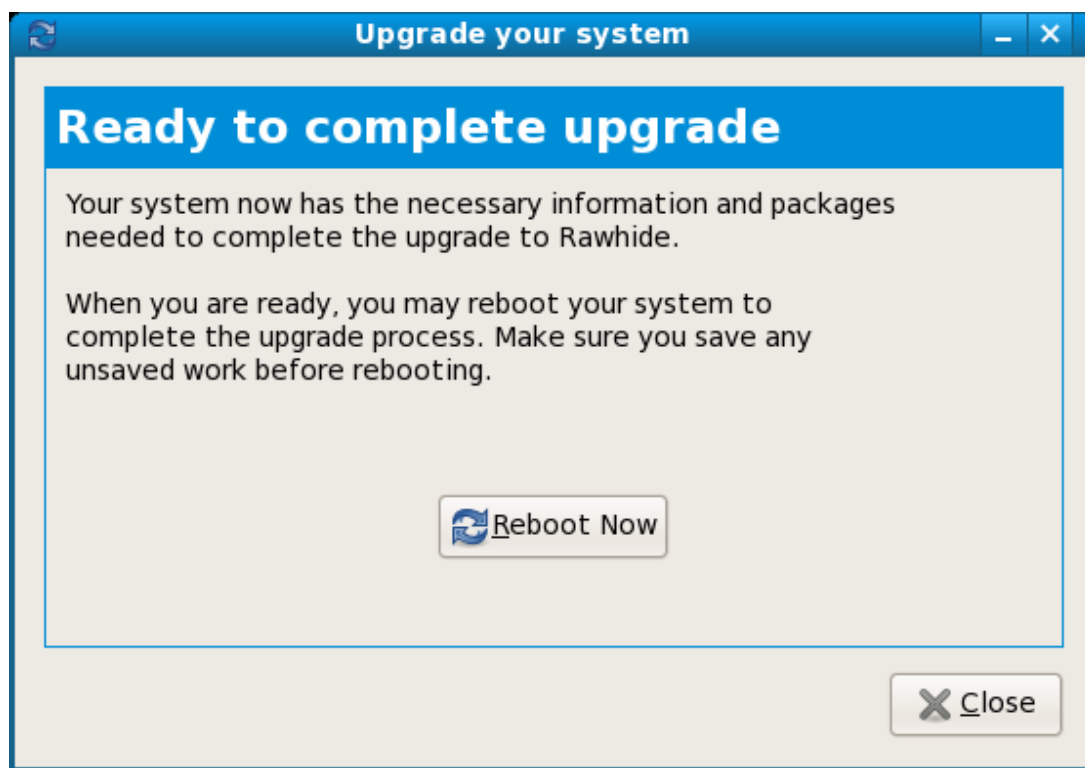
RHEL Tomorrow: “In Place” Upgrade

- Downloads new initrd & kernel images



RHEL Tomorrow: “In Place” Upgrade

- User reboots, brought into installer



RHEL Tomorrow: `gnome-control-center`

- `gnome-control-center`
 - It is not YaST (yet)
 - It is a unified GUI for package management and system configuration
- Benefit
 - Progress towards a YaST-like tool in RHEL (currently we have the `system-config-*` GUIs/TUIs)

RHEL Tomorrow: PackageKit

- PackageKit
 - Abstraction layer for YUM, apt, conary, etc
 - Provides a common set of abstractions that can be used by GUI/TUI package managers

rpm

dpkg

ipkg


RHEL Tomorrow: PackageKit

- PackageKit
 - Abstraction layer for YUM, apt, conary, etc
 - Provides a common set of abstractions that can be used by GUI/TUI package managers



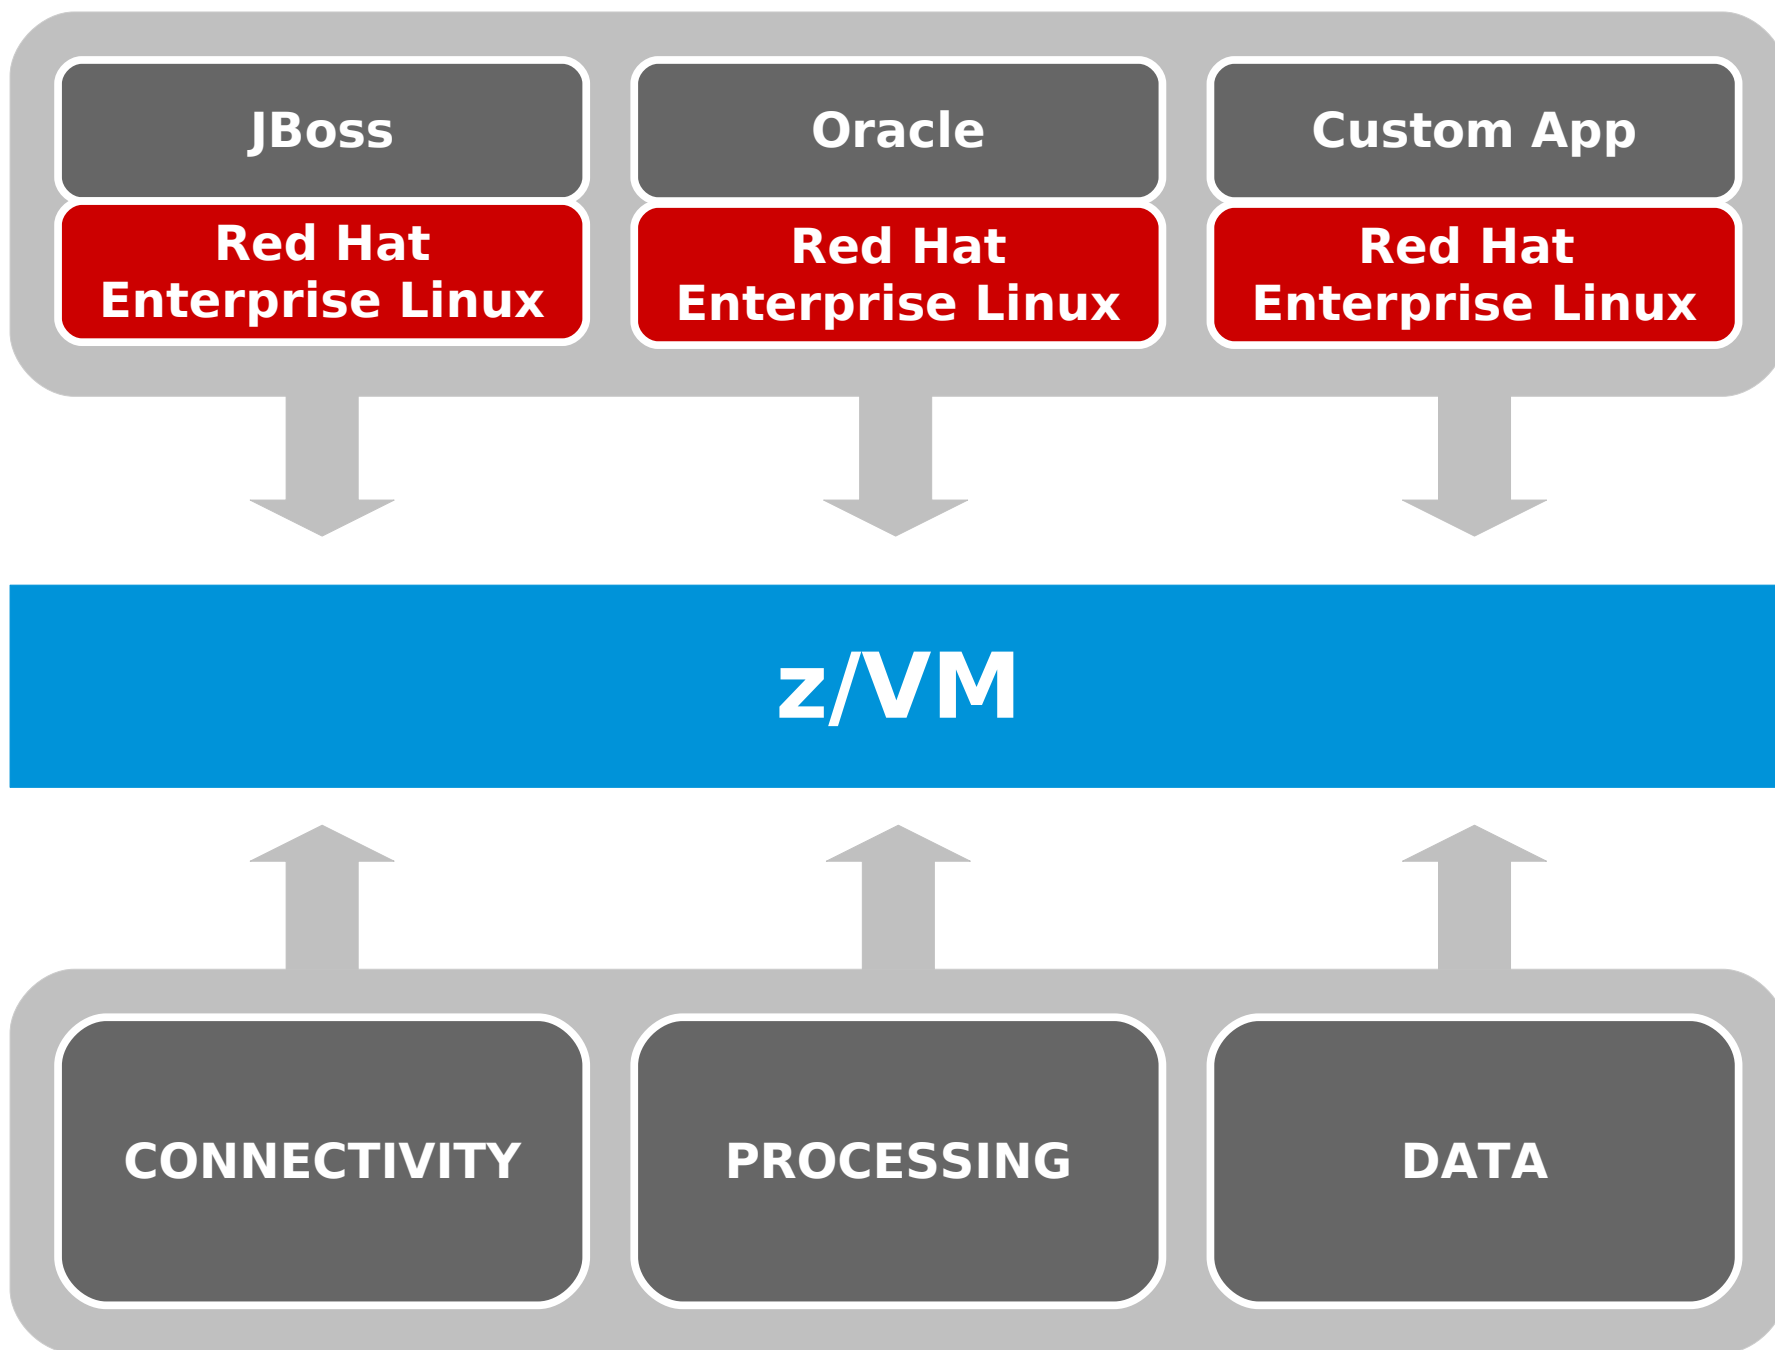
RHEL Tomorrow: PackageKit

- PackageKit
 - Abstraction layer for YUM, apt, conary, etc
 - Provides a common set of abstractions that can be used by GUI/TUI package managers





Linux Virtualization on System z





Using RHN Satellite to Manage s390/s390x & distributed

Red Hat Network Satellite



Red Hat Network Satellite

Update



Manage



Red Hat Network Satellite

Update



Manage



Provision



Red Hat Network Satellite

Update



Manage



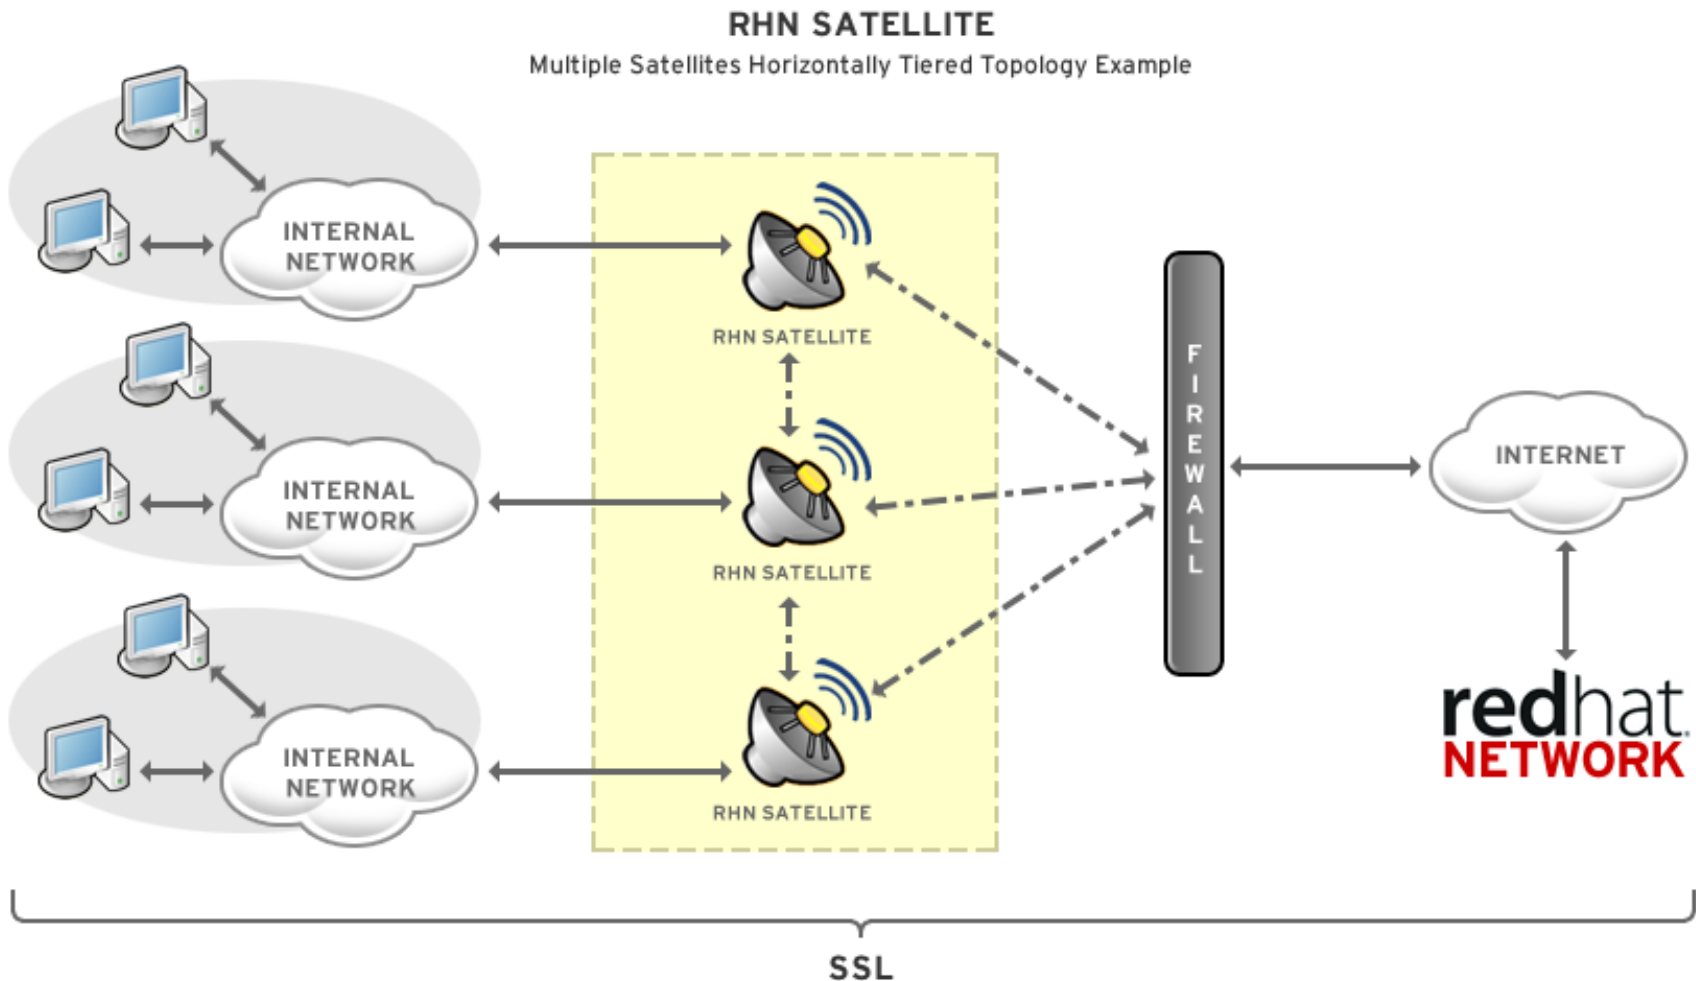
Provision



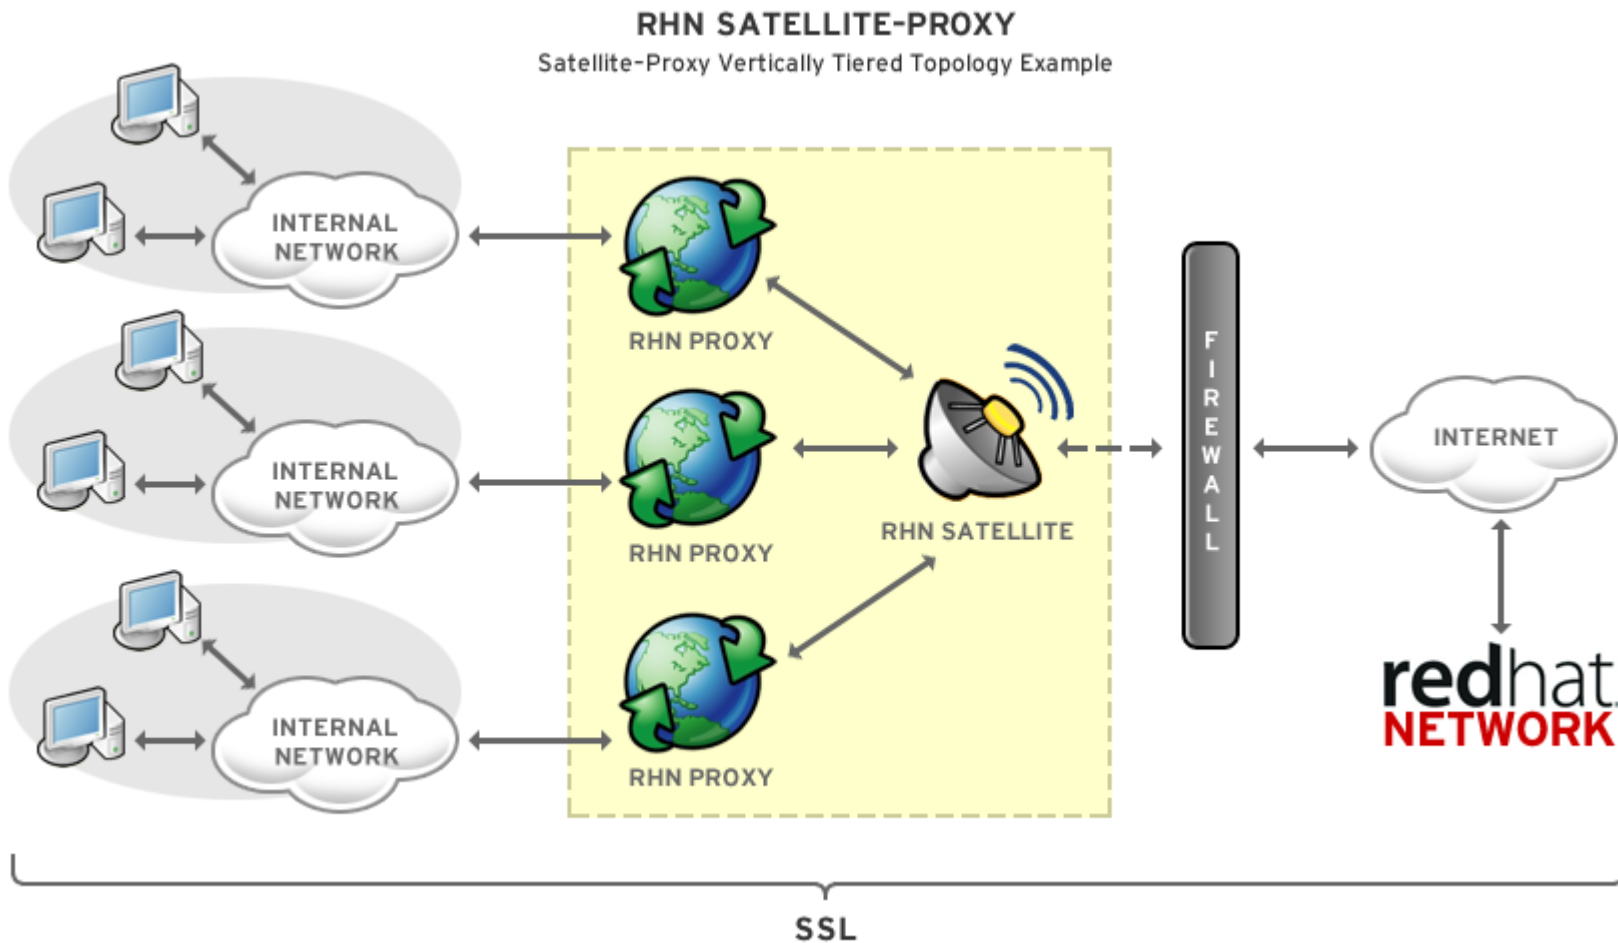
Monitor



Red Hat Network Satellite

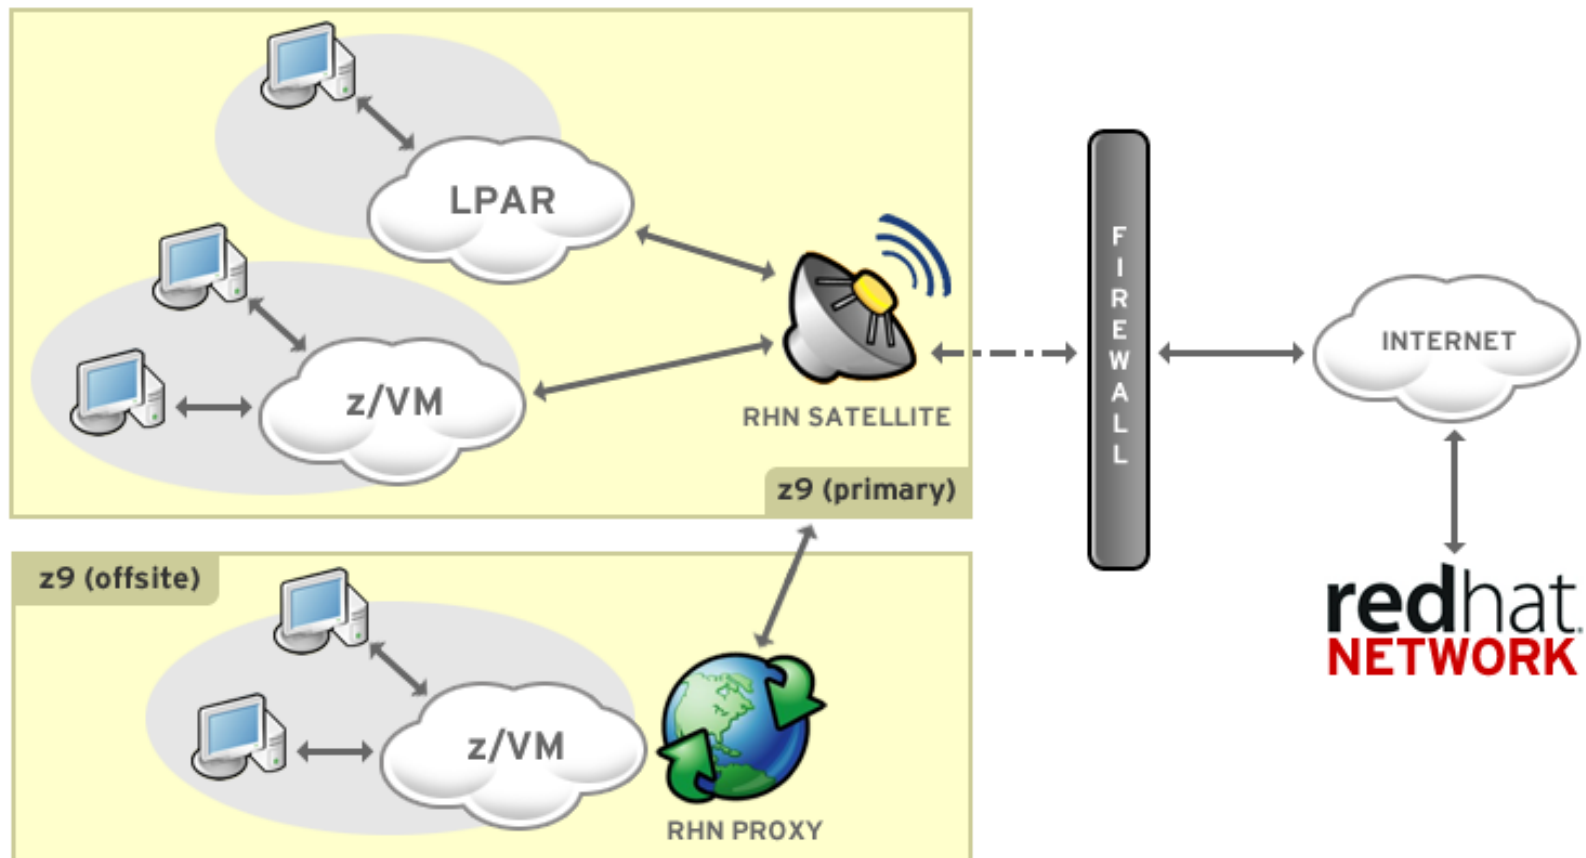


Red Hat Network Satellite



Red Hat Network Satellite

RHN SATELLITE-PROXY
Satellite-Proxy System z Topology Example



RHN Installation Requirements

- **Software**
 - RHEL 4 (31-bit or 64-bit)
 - @Base install
- **Hardware**
 - 1 to 2 (virtual) IFLs
 - 2 to 4 GB storage (memory)
 - 1 GB swap (combination VDISK, disk)
 - 1 x mod3 for OS install
 - Estimated 12 GB disk space for embedded database
 - 6 GB per channel (disk)

- Overview
- Systems**
- All
- Virtual Systems
- Out of Date
- Unentitled
- Ungrouped
- Inactive
- Recently Registered
- Proxy
- System Groups
- System Set Manager
- Advanced Search
- Activation Keys
- Stored Profiles
- Custom System Info
- Kickstart

 Virtual Systems ?Filter by System:

1 - 20 of 66 (2 selected) << < > >>

<input type="checkbox"/>	System	Updates	Status	Base Software Channel
<input type="checkbox"/>	● Host: inf01.coe.muc.redhat.com 3 Active Virtual Systems, 4 Total. (View All)			
<input type="checkbox"/>	└ vinf02.coe.muc.redhat.com		Running	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vinf03.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux AS (v. 4 for 32-bit x86)
<input type="checkbox"/>	● Host: storage03.coe.muc.redhat.com 0 Active Virtual Systems, 0 Total. (View All)			
<input type="checkbox"/>	● Host: inf02.coe.muc.redhat.com 2 Active Virtual Systems, 2 Total. (View All)			
<input type="checkbox"/>	└ vinf05.coe.muc.redhat.com	<input checked="" type="checkbox"/>	Running	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vinf06.coe.muc.redhat.com	<input checked="" type="checkbox"/>	Running	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	● Host: hv001.coe.muc.redhat.com 1 Active Virtual Systems, 28 Total. (View All)			
<input type="checkbox"/>	└ vm019.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux AS (v. 4 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vm040.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux AS (v. 4 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)
<input type="checkbox"/>	└ vm003.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux AS (v. 4 for 32-bit x86)
<input type="checkbox"/>	└ vm045.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux AS (v. 4 for 32-bit x86)
<input type="checkbox"/>	└ vm013.coe.muc.redhat.com		Stopped	Red Hat Enterprise Linux (v. 5 for 32-bit x86)

System Legend

- OK
- Critical
- Warning

- Overview
- Systems
- System Groups
- System Set Manager
- Advanced Search
- Activation Keys
- Stored Profiles
- Custom System Info
- Kickstart**
- Profiles
- Bare Metal
- GPG and SSL Keys
- Distributions
- File Preservation



Kickstart: rhel-5-i386-server_default_part_novirt

- [Kickstart Details](#)
[System Details](#)
[Software](#)
[Activation Keys](#)
[Scripts](#)
[Kickstart File](#)

Kickstart File

The kickstart file generated by this kickstart profile is viewable below:

[Download Kickstart File](#)

```
# Kickstart config file generated by RHN Config Management
#
# Profile Name : rhel-5-i386-server_default_part_novirt
# Profile Label : rhel-5-i386-server_default_part_novirt
# Date Created : 2008-06-03 20:40:03.0
#

install
text
network --bootproto dhcp
url --url http://devel13.z900.redhat.com/ty/MwPJrTGI
lang en_US
langsupport --default en_US en_US
keyboard us
mouse none
zerombr yes
clearpart --all
part /boot --fstype=ext3 --size=200
part pv.01 --size=1000 --grow
part swap --size=1000 --maxsize=2000
volgroup myvg pv.01
logvol / --vgname=myvg --name=rootvol --size=1000 --grow
bootloader --location mbr
timezone America/New_York
auth --enablemd5 --enablesshadow
rootpw --iscrypted $1$0KAZmj1I$V05gL5mVVj9T09GidA/Y6/
selinux --permissive
reboot
firewall --disabled
skipx
repo --name=Cluster --baseurl=http://devel13.z900.redhat.com/kickstart/dist/ks-rhel-i386-server-5-u1/Cluster
repo --name=ClusterStorage --baseurl=http://devel13.z900.redhat.com/kickstart/dist/ks-rhel-i386-server-5-u1/ClusterStorage
repo --name=VT --baseurl=http://devel13.z900.redhat.com/kickstart/dist/ks-rhel-i386-server-5-u1/VT
repo --name=Workstation --baseurl=http://devel13.z900.redhat.com/kickstart/dist/ks-rhel-i386-server-5-u1/Workstation
```

RHN Satellite Is Now Open Source

<http://spacewalk.redhat.com>

- Announced at Red Hat Summit 2008
 - remember the Fedora -> RHEL model?



Security

Agenda

- **Why do we need SELinux? What are the principal concepts?**
- **SELinux Details**
 - **Type Enforcement**
 - **What are the available policies?**
 - **What's a policy actually made of?**
 - **How do I {add, change} a policy?**
 - **What's the associated overhead?**
- **Usage**
 - **User Perspective**
 - **Admin Perspective**
- **Scenarios**
 - **Fixing the RHT Corporate VPN “update”**



Why do we need SELinux?

Linux Access Control Problems

- Access is based off users' access

Example: Firefox can read SSH keys

```
# ps -x | grep firefox  
shawn 21375 1 35 11:38 ? 00:00:01 firefox-bin
```

```
# ls -l id_rsa  
-rw----- 1 shawn shawn 1743 2008-08-10 id_rsa
```

Fundamental Problem: Security properties not specific enough. Kernel can't distinguish applications from users.

Linux Access Control Problems

2) Processes can change security properties

Example: Mail files are readable only by me..... but
Thunderbird could make them world readable

Fundamental Problems:

- Standard access control is discretionary
- Includes concept of “resource ownership”
- Processes can escape security policy

Linux Access Control Problems

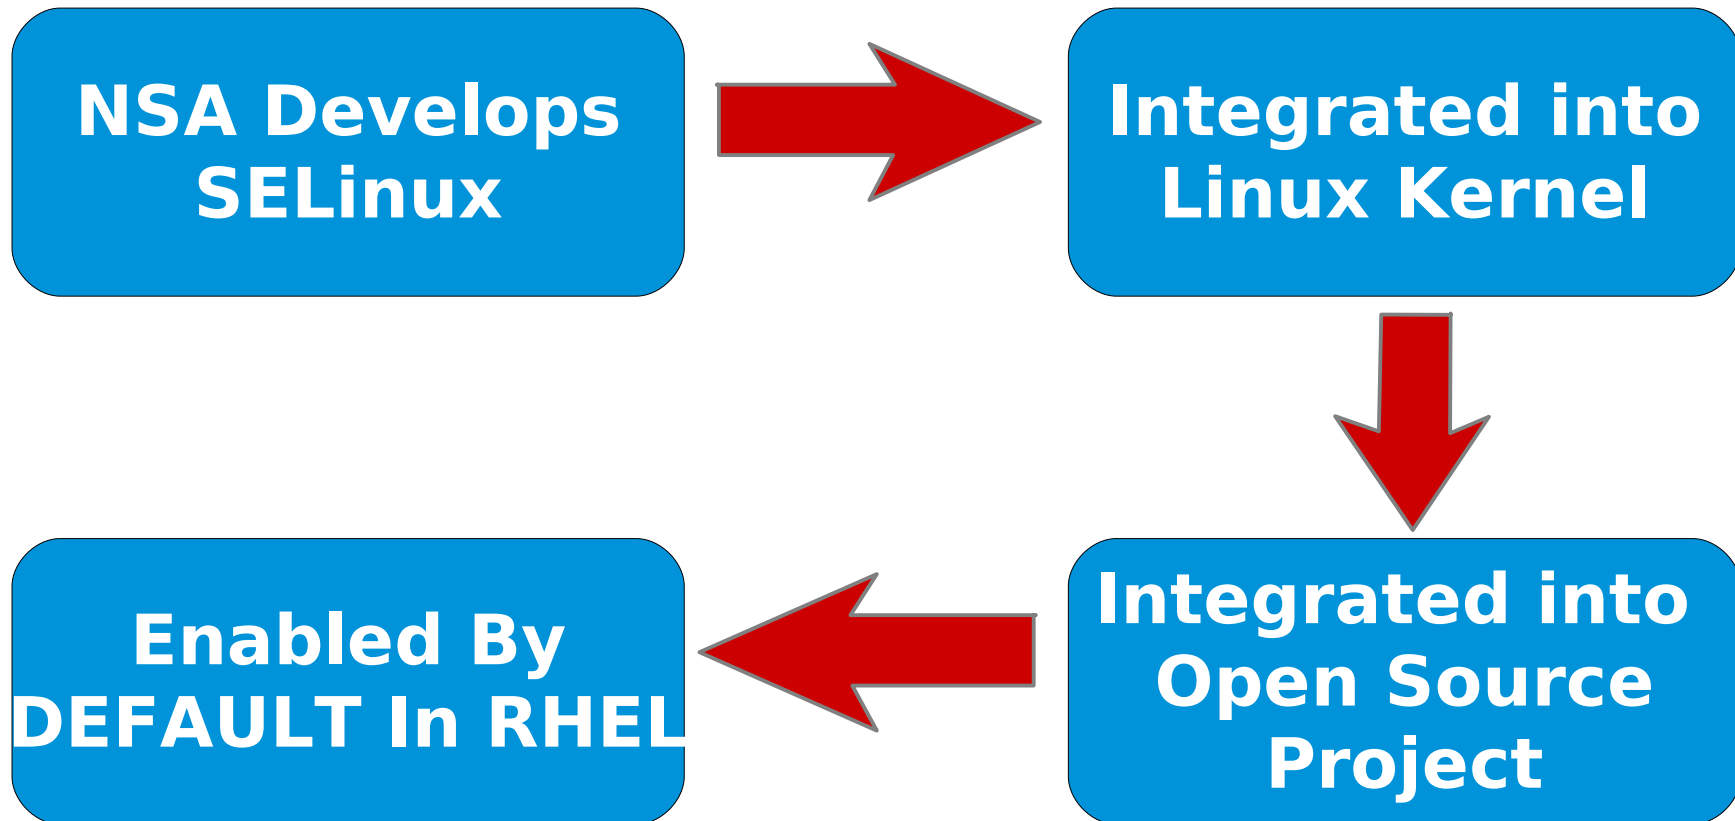
3) Only two privilege levels: User & root

Example: Apache gets hacked, allowing remote access to root. Entire system is compromised.

Fundamental Problems:

- Simplistic security policy
- No way to enforce least-privilege

SELinux: Building Security Openly



Customers, NSA, Community, and Red Hat continue evolution

Red Hat Security Certifications

NIAP/Common Criteria: The most evaluated operating system platform

- Red Hat Enterprise Linux 2.1 – EAL 2 (Completed: February 2004)
- Red Hat Enterprise Linux 3 EAL 3+/CAPP (Completed: August 2004)
- Red Hat Enterprise Linux 4 EAL 4+/CAPP (Completed: February 2006)
- Red Hat Enterprise Linux 5 EAL4+/CAPP/LSPP/RBAC (Completed: June 2007)

DII-COE

- Red Hat Enterprise Linux 3 (Self-Certification Completed: October 2004)
- Red Hat Enterprise Linux: First Linux platform certified by DISA

DCID 6/3

- Currently PL3/PL4: ask about kickstarts.
- Often a component in PL5 systems

DISA SRRs / STIGs

- Ask about kickstarts.

FIPS 140-2

- Red Hat / NSS Cryptography Libraries certified Level 2

Security Standards Work

Extensible Configuration Checklist Description Format (XCCDF)

- Enumeration for configuration requirements
- DISA FSO committed to deploying STIG as XCCDF
- Others working with NIST
- Security policy becomes one file

Open Vulnerability & Assessment Language (OVAL)

- Machine-readable versions of security advisories

Common Vulnerability and Exposures (CVE) Compatibility

- Trace a vulnerability through multiple vendors



How's it work?

Linux Access Control Introduction

Linux access control involves the kernel controlling

- **Processes** (running programs), which try to access...
 - **Resources** (files, directories, sockets, etc)

For example:

- Apache (process) can read web files
- But **not** the `/etc/shadow` file (resource)

Traditional methods do not clearly separate the privileges of users and applications acting on the users behalf, increasing the damage that can be caused by application exploits.

So, how should these decisions be made?

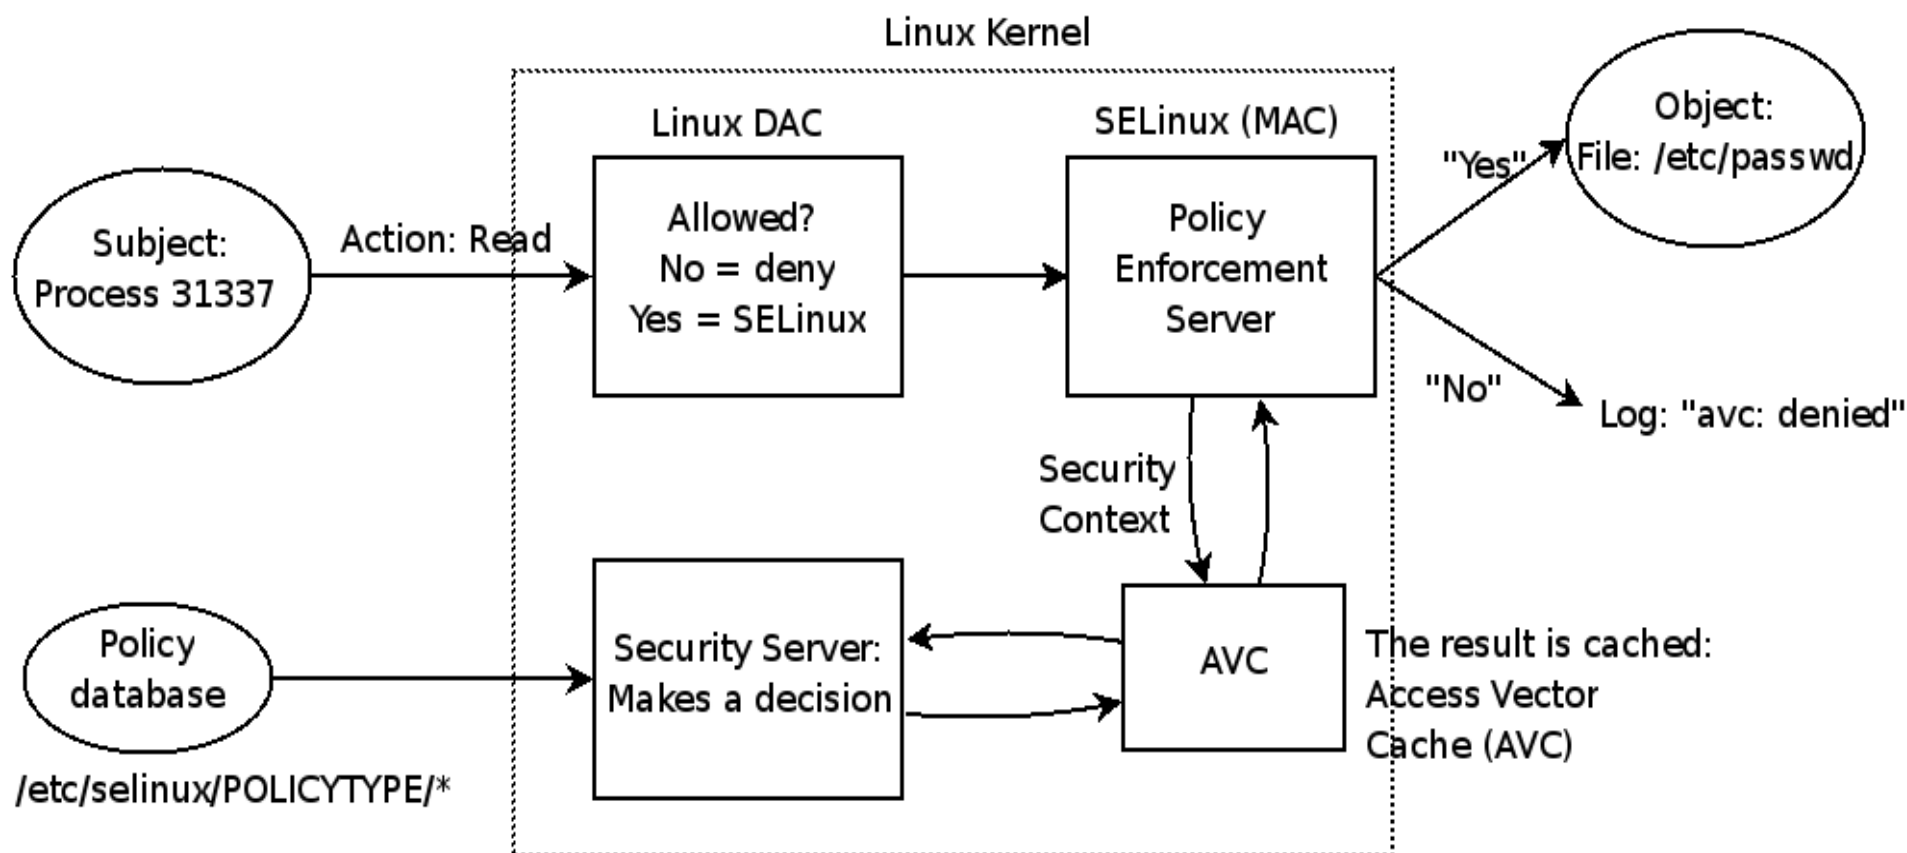
Security Architecture

Every subject (i.e process) and object (i.e. data files) are assigned collections of security attributes, called a **security context**

- 1) Security context of subject & object passed to SELinux
- 2) Kernel/SELinux check, verify access
 - 2a) Grant access. Record allowance in AVC (Access Vector Cache)
 - 2b) Deny access, log error

Security Architecture

Or in picture view...





SELinux Details

SELinux Contexts

root:object r:sysadm home t:s0:c0

- The above is an SELinux context
- user_t
- role_t
- file_t
- Sensitivity
- category

Role Based Access Control (RBAC)

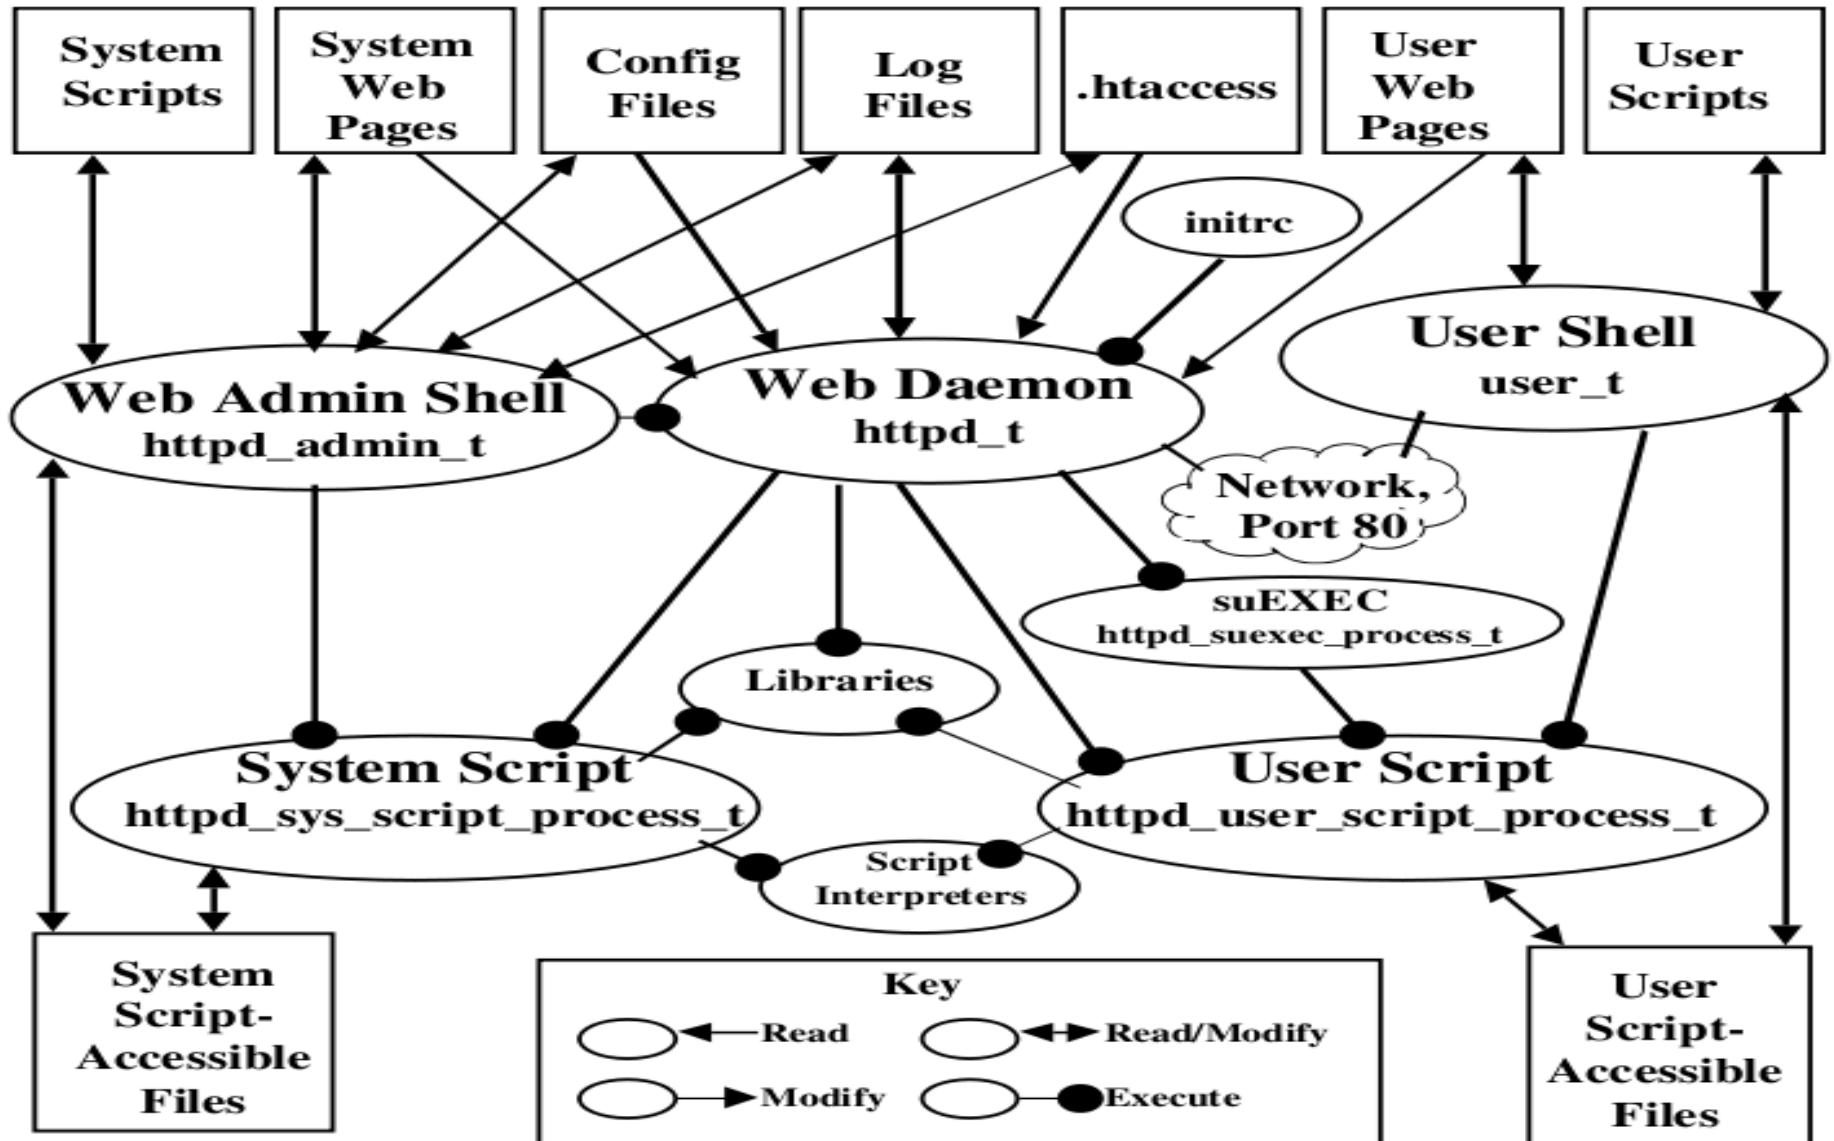
“root” really isn't “root”

i.e:

root_u:**WebServerAdmin_r**:SysAdmin_t

root_u:**OracleDBAdmin_r**:SysAdmin_t

SELinux Contexts



SELinux Policy

- Policies are matrices of statements which tell SELinux if certain actions are allowed based on the context of the objects attempting those actions.
- There are three SELinux Policy Types

The Three SELinux Policy Types

■ Targeted Policy

- *Default policy in RHEL5. Supported by HelpDesk.*
- Targets specific applications to lock down.
- Allows all other applications to run in the unconfined domain (`unconfined_t`)
- Applications running in the unconfined domain run as if SELinux were disabled

The Three SELinux Policy Types

2) Strict Policy

- Denies access to everything by default
- Complete protection for all processes on the system
- Requires that policies be written for **all** applications, often requires customization
- Strict is type enforcement with added types for users (e.g. `user_t` and `user_firefox_t`).
- Not enabled by Red Hat as default

The Three SELinux Policy Types

3) Multi-Level Security (MLS)


- Focuses on confidentiality (i.e. separation of multiple classifications of data)
- Ability to manage {processes, users} with varying levels of access. (i.e. “*the need to know*”)
- Uses category & sensitivity levels

The Three SELinux Policy Types

3) Multi-Level Security (MLS)

(a) Sensitivity Labels

- Mostly used by the government – Top Secret, Secret, Unclassified, etc



The Three SELinux Policy Types

3) Multi-Level Security (MLS)

(b) Category Labels

- Separation of data types, compartments, projects, etc

s0

Unclassified

s1

Secret

c0

Project A

c1

Project B

s1

Top Secret

c0

Alpha

c1

Bravo

c2

Charlie

c3

Delta

The Three SELinux Policy Types

3) Multi-Level Security (MLS)

(b) Polyinstantiation & pam_namespace

- The pam_namespace PAM module sets up a private namespace for a session with polyinstantiated directories
- A polyinstantiated directory provides a different instance of itself based on user name, or when using SELinux, user name, security context or both

The Three SELinux Policy Types

3) Multi-Level Security (MLS)

(b) Polyinstantiation & pam_namespace

```
# id -Z
```

```
staff_u:WebServer_Admin_r:WebServer_Admin_t:s0:c0
```

```
# ls -l /data
```

```
secret-file-1
```

```
secret-file 2
```

```
# id -Z
```

```
staff_u:WebServer_Admin_r:WebServer_Admin_t:s1:c0
```

```
# ls -l /data
```

```
secret-file-1
```

```
secret-file 2
```


```
top-secret-file-1
```

The Three SELinux Policy Types

Multi-Level Security (MLS) & Common Criteria

- The Common Criteria (CC) is an international security standard against which systems are evaluated. Many government customers require CC evaluated systems.
- Red Hat Enterprise Linux 5 meets EAL4+ with RBAC/LSPP/CAPP endorcements

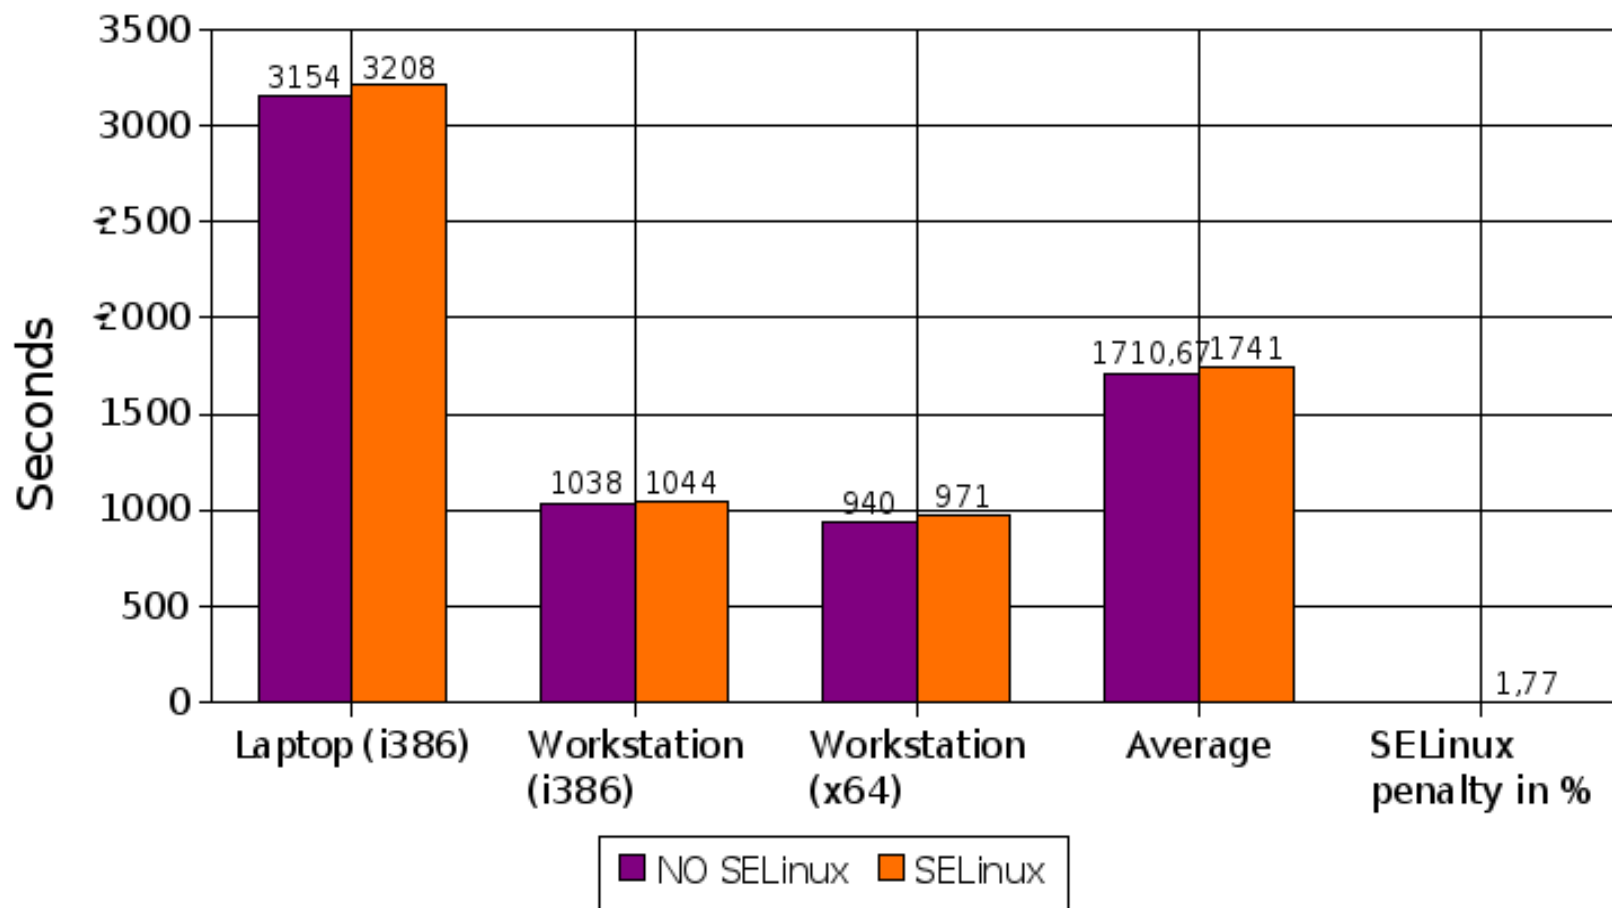
What's the Performance Overhead?



What's the Performance Overhead?

RHEL5 SELinux: MySQL 5.0.22

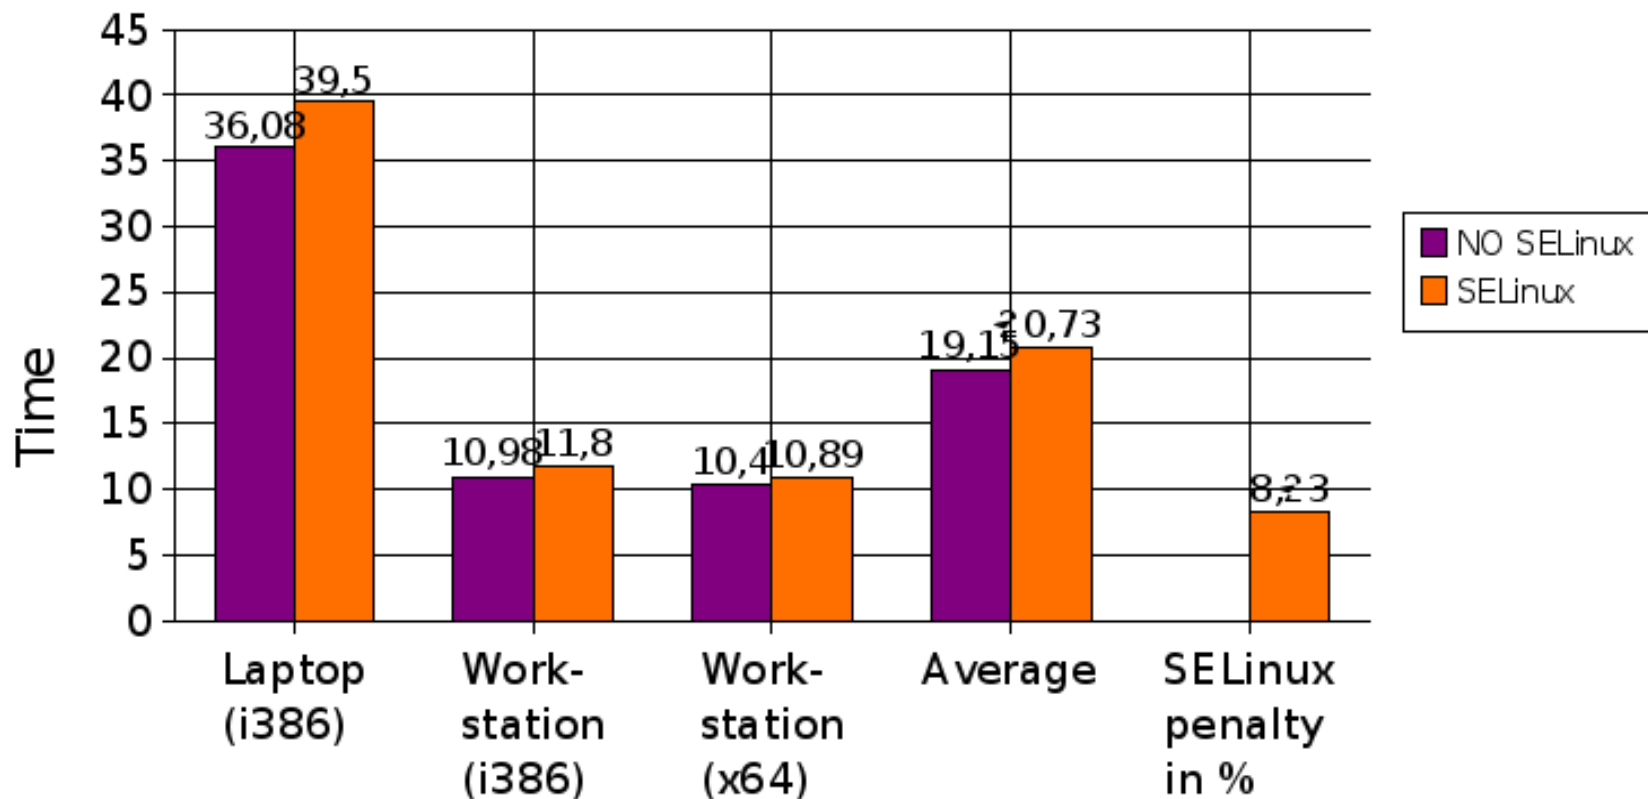
MySQL Benchmark suite: run-all-tests. Lower is better.



What's the Performance Overhead?

RHEL5 SELinux: Apache 2.2.3 (worker)

11 tests: 100000 requests with 1-255 concurrent connections. Lower is better.



What's the Performance Overhead?

- **Not official statistics**
- **Laptop = 2GHz, 2x 1GB RAM**
- **Workstation = 2.13GHz, 4x 1GB RAM**
- **Apache = Lots of threads**
- **MySQL = Lots of disk I/O**

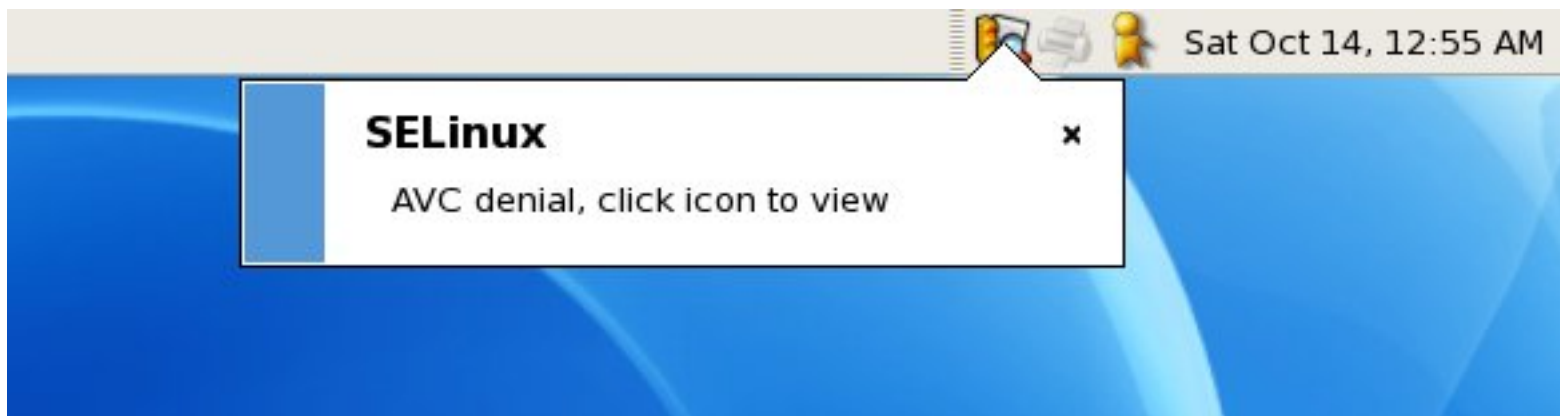


SELinux Usage

(GUI & console)

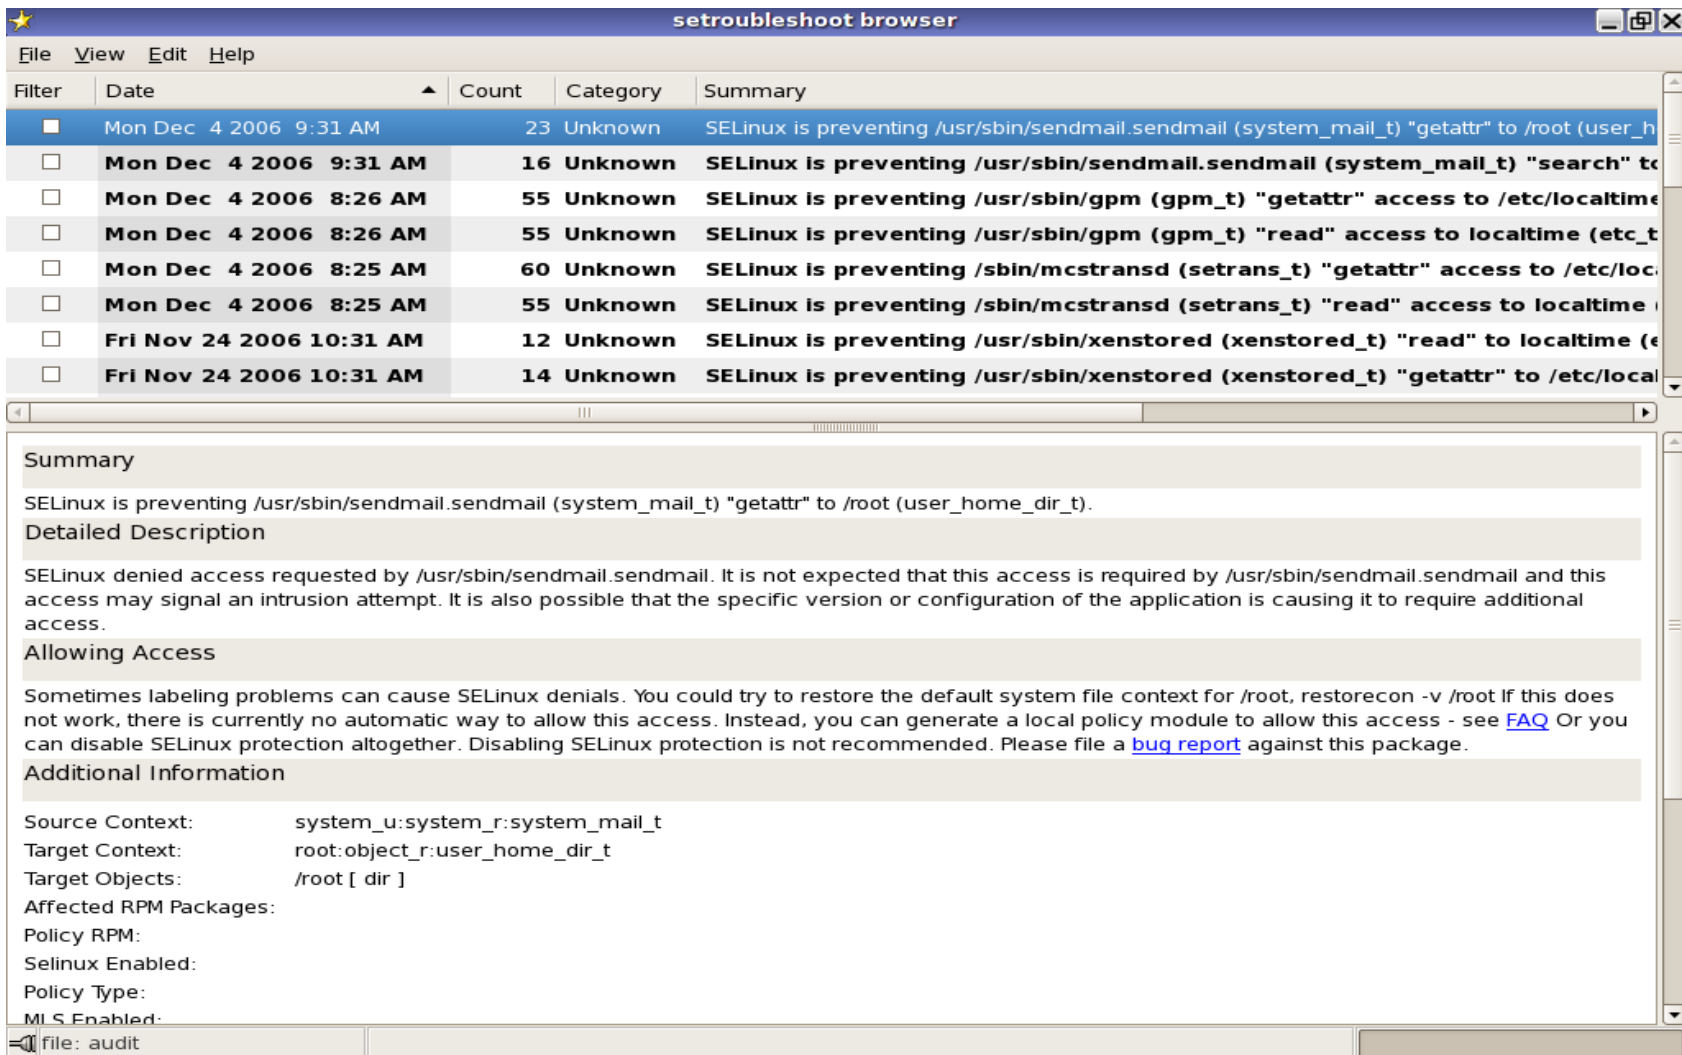
End-User Perspective

- **sealert Notifications**



End-User Perspective

- **sealert Browser**



The screenshot shows the 'setroubleshoot browser' window. At the top, there is a menu bar with 'File', 'View', 'Edit', and 'Help'. Below the menu is a table with columns: Filter, Date, Count, Category, and Summary. The table lists several SELinux denials. The first entry is selected, and its details are shown in a scrollable area below the table.

Filter	Date	Count	Category	Summary
<input checked="" type="checkbox"/>	Mon Dec 4 2006 9:31 AM	23	Unknown	SELinux is preventing /usr/sbin/sendmail.sendmail (system_mail_t) "getattr" to /root (user_h
<input type="checkbox"/>	Mon Dec 4 2006 9:31 AM	16	Unknown	SELinux is preventing /usr/sbin/sendmail.sendmail (system_mail_t) "search" to
<input type="checkbox"/>	Mon Dec 4 2006 8:26 AM	55	Unknown	SELinux is preventing /usr/sbin/gpm (gpm_t) "getattr" access to /etc/localtime
<input type="checkbox"/>	Mon Dec 4 2006 8:26 AM	55	Unknown	SELinux is preventing /usr/sbin/gpm (gpm_t) "read" access to localtime (etc_t
<input type="checkbox"/>	Mon Dec 4 2006 8:25 AM	60	Unknown	SELinux is preventing /sbin/mcstransd (setrans_t) "getattr" access to /etc/loc
<input type="checkbox"/>	Mon Dec 4 2006 8:25 AM	55	Unknown	SELinux is preventing /sbin/mcstransd (setrans_t) "read" access to localtime
<input type="checkbox"/>	Fri Nov 24 2006 10:31 AM	12	Unknown	SELinux is preventing /usr/sbin/xenstored (xenstored_t) "read" to localtime (e
<input type="checkbox"/>	Fri Nov 24 2006 10:31 AM	14	Unknown	SELinux is preventing /usr/sbin/xenstored (xenstored_t) "getattr" to /etc/local

Summary

SELinux is preventing /usr/sbin/sendmail.sendmail (system_mail_t) "getattr" to /root (user_home_dir_t).

Detailed Description

SELinux denied access requested by /usr/sbin/sendmail.sendmail. It is not expected that this access is required by /usr/sbin/sendmail.sendmail and this access may signal an intrusion attempt. It is also possible that the specific version or configuration of the application is causing it to require additional access.

Allowing Access

Sometimes labeling problems can cause SELinux denials. You could try to restore the default system file context for /root, restorecon -v /root If this does not work, there is currently no automatic way to allow this access. Instead, you can generate a local policy module to allow this access - see [FAQ](#) Or you can disable SELinux protection altogether. Disabling SELinux protection is not recommended. Please file a [bug report](#) against this package.

Additional Information

Source Context: system_u:system_r:system_mail_t
 Target Context: root:object_r:user_home_dir_t
 Target Objects: /root [dir]
 Affected RPM Packages:
 Policy RPM:
 Selinux Enabled:
 Policy Type:
 MLS Enabled:

file: audit

System Administrator Perspective

- **sealert + EMail Notifications**

The screenshot shows the Evolution email client interface. The main window displays an email from SELinux_Troubleshoot@redhat.com with the subject "[SELinux AVC Alert] SELinux is preventing the ftp daemon from reading users home directories (/)".

From: SELinux_Troubleshoot@redhat.com
To: jdennis@redhat.com
Subject: [SELinux AVC Alert] SELinux is preventing the ftp daemon from reading users home directories (/).
Date: Sat, 14 Oct 2006 05:02:35 -0000 (01:02 EDT)

Summary
 SELinux is preventing the ftp daemon from reading users home directories (/).

Detailed Description
 SELinux has denied the ftp daemon access to users home directories (/). Someone is attempting to login via your ftp daemon to a user account. If you only setup ftp to allow anonymous ftp, this could signal a intrusion attempt.

Allowing Access
 If you want ftp to allow users access to their home directories you need to tum on the ftp_home_dir boolean: "setsebool -P ftp_home_dir=1"

The following command will allow this access:
 setsebool -P ftp_home_dir=1

Additional Information

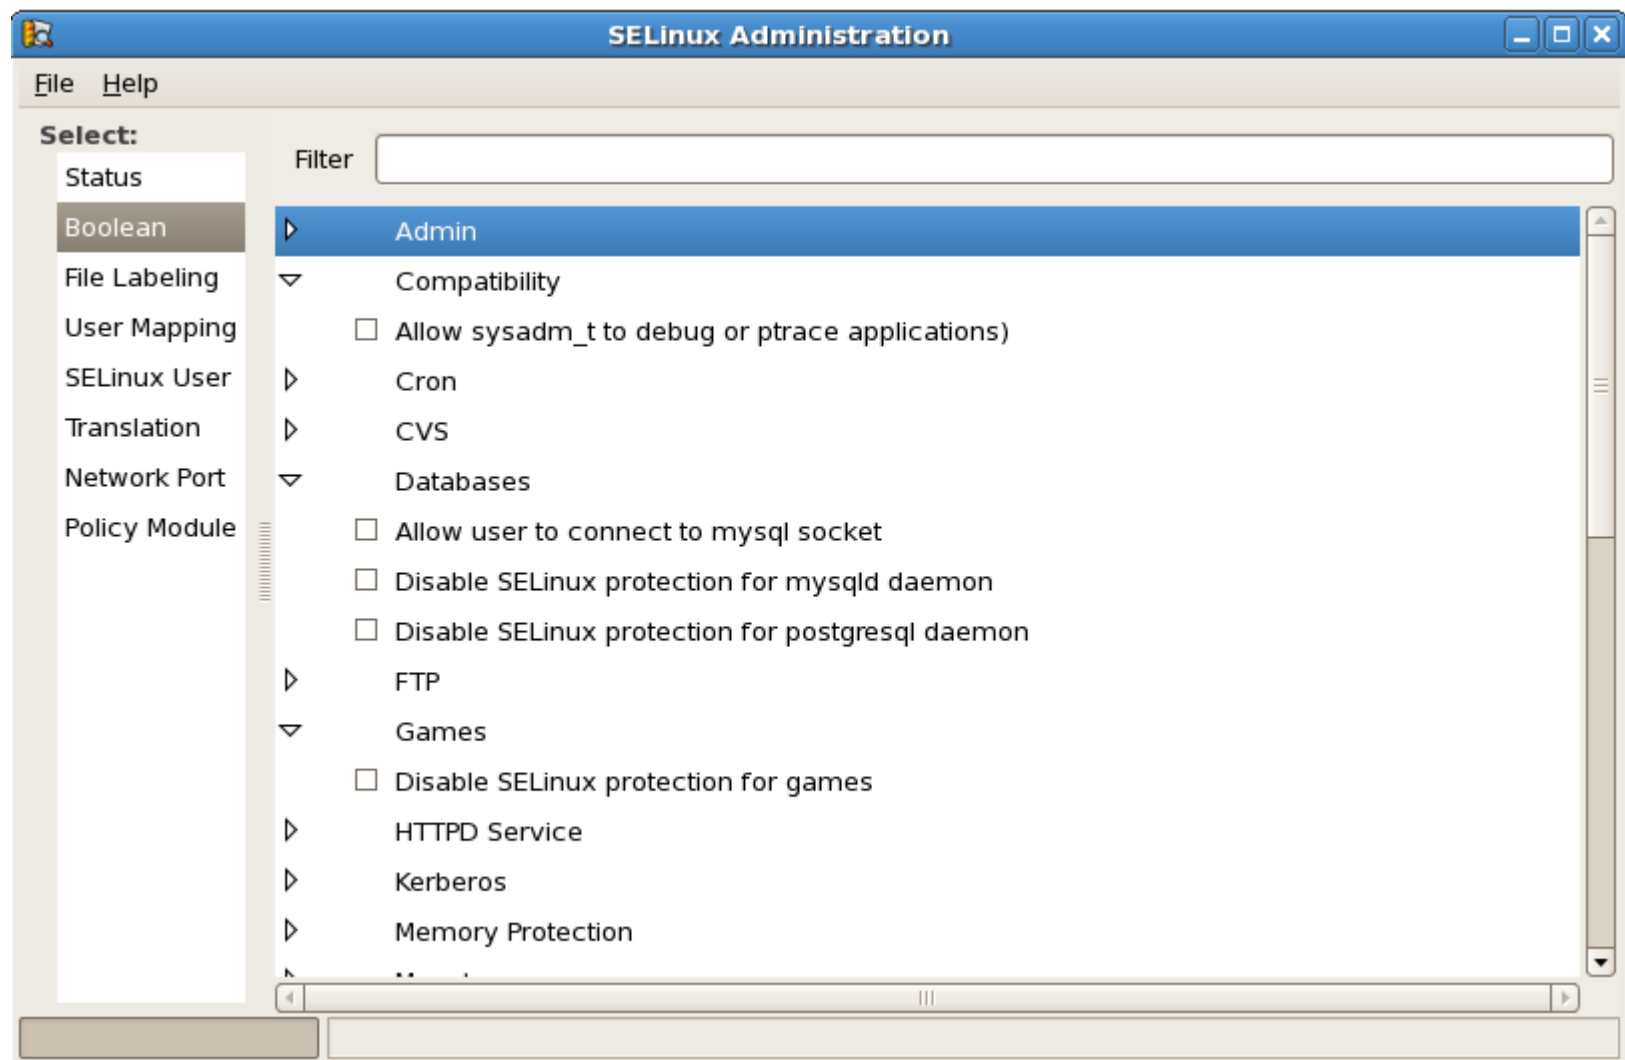
```

Source Context:      system_u:system_r:ftpd_t
Target Context:     system_u:object_r:home_root_t
Target Objects:     / [ dir ]
Affected RPM Packages: vsftpd-2.0.5-8 [application]filesystem-2.3.7-2.1 [target]
Policy RPM:         selinux-policy-2.3.18-8
Selinux Enabled:    1
Policy Type:        targeted
MLS Enabled:        1
Enforcing Mode:     Enforcing
Plugin Name:        plugins.ftp_home_dir
Host Name:          localhost.localdomain
Platform:           Linux localhost.localdomain 2.6.18-1.2726.fc6 #1 SMP Mon Oct 2 19:27:36 EDT 2006 i686 i686
  
```

Raw Audit Messages:

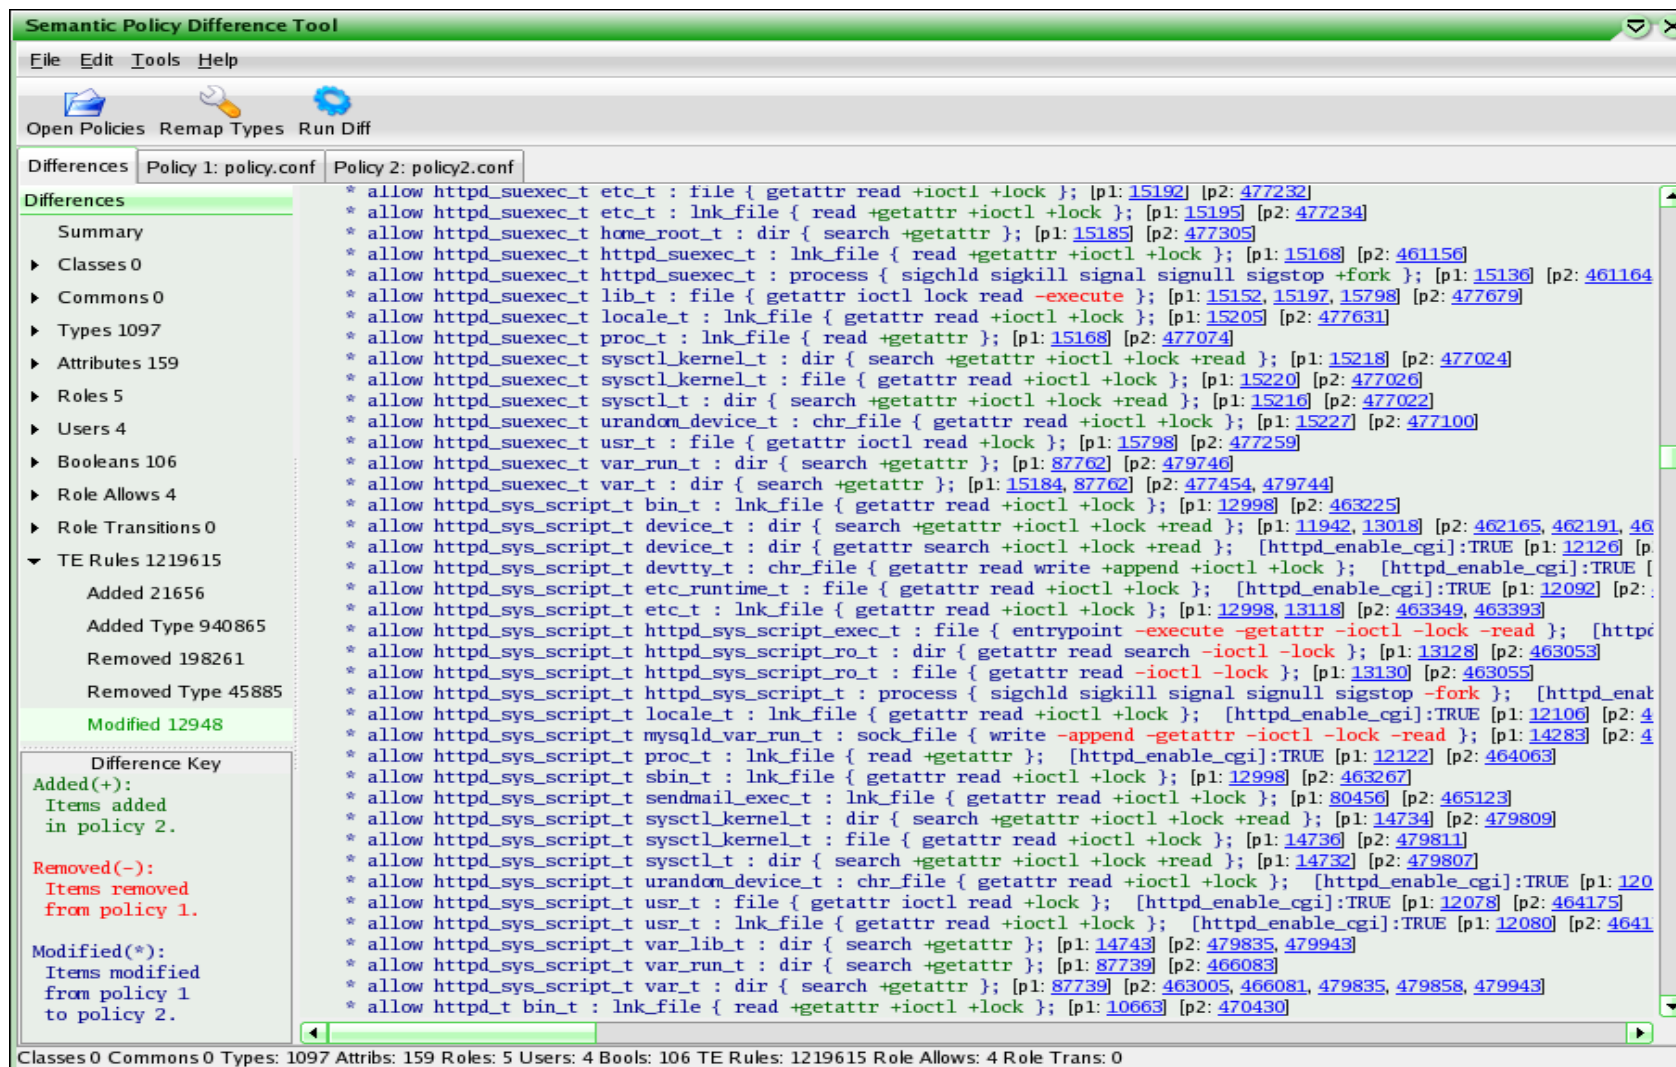
System Administrator Perspective

- **system-config-selinux**



System Administrator Perspective

- **sediffx**



Semantic Policy Difference Tool

File Edit Tools Help

Open Policies Remap Types Run Diff

Differences Policy 1: policy.conf Policy 2: policy2.conf

Differences

Summary

- Classes 0
- Commons 0
- Types 1097
- Attributes 159
- Roles 5
- Users 4
- Booleans 106
- Role Allows 4
- Role Transitions 0
- TE Rules 1219615
 - Added 21656
 - Added Type 940865
 - Removed 198261
 - Removed Type 45885
 - Modified 12948

Difference Key

Added(+):
Items added
in policy 2.

Removed(-):
Items removed
from policy 1.

Modified(*):
Items modified
from policy 1
to policy 2.

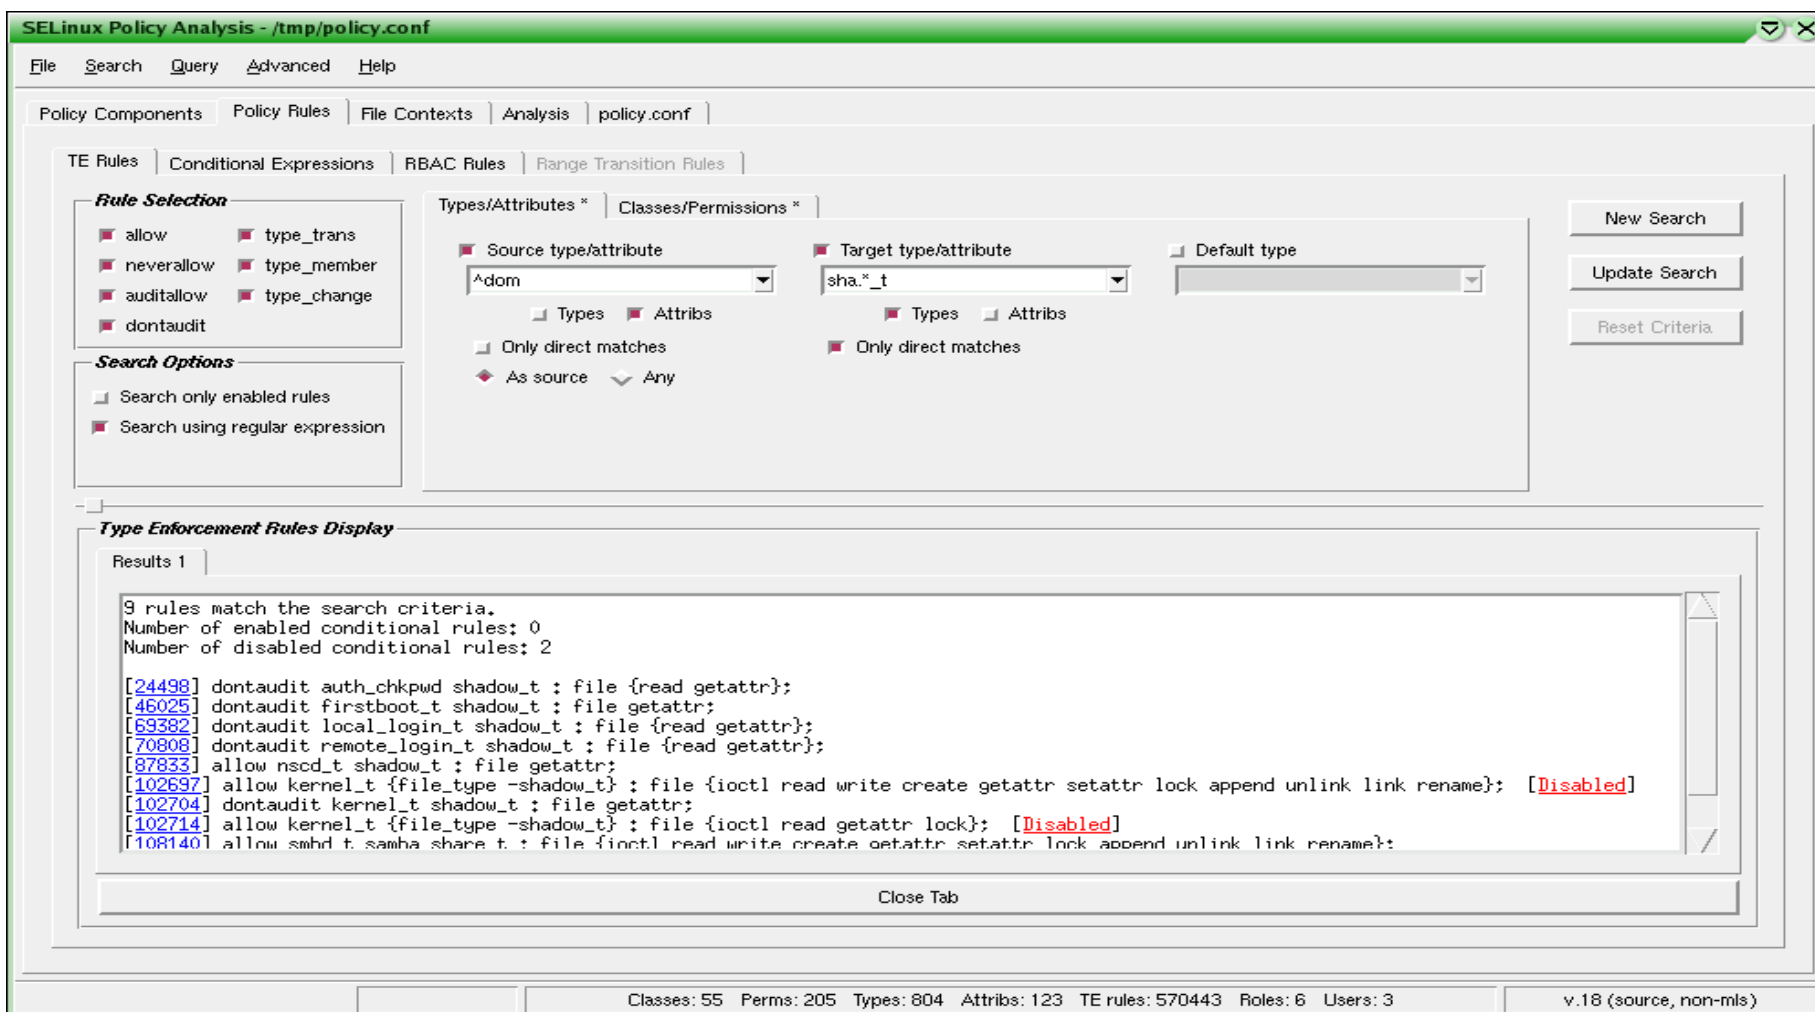
```

* allow httpd_suexec_t etc_t : file { getattr read +ioctl +lock }; [p1: 15192] [p2: 477232]
* allow httpd_suexec_t etc_t : lnk_file { read +getattr +ioctl +lock }; [p1: 15195] [p2: 477234]
* allow httpd_suexec_t home_root_t : dir { search +getattr }; [p1: 15185] [p2: 477305]
* allow httpd_suexec_t httpd_suexec_t : lnk_file { read +getattr +ioctl +lock }; [p1: 15168] [p2: 461156]
* allow httpd_suexec_t httpd_suexec_t : process { sigchld sigkill signal signull sigstop +fork }; [p1: 15136] [p2: 461164]
* allow httpd_suexec_t lib_t : file { getattr ioctl lock read -execute }; [p1: 15152, 15197, 15798] [p2: 477679]
* allow httpd_suexec_t locale_t : lnk_file { getattr read +ioctl +lock }; [p1: 15205] [p2: 477631]
* allow httpd_suexec_t proc_t : lnk_file { read +getattr }; [p1: 15168] [p2: 477074]
* allow httpd_suexec_t sysctl_kernel_t : dir { search +getattr +ioctl +lock +read }; [p1: 15218] [p2: 477024]
* allow httpd_suexec_t sysctl_kernel_t : file { getattr read +ioctl +lock }; [p1: 15220] [p2: 477026]
* allow httpd_suexec_t sysctl_t : dir { search +getattr +ioctl +lock +read }; [p1: 15216] [p2: 477022]
* allow httpd_suexec_t urandom_device_t : chr_file { getattr read +ioctl +lock }; [p1: 15227] [p2: 477100]
* allow httpd_suexec_t usr_t : file { getattr ioctl read +lock }; [p1: 15798] [p2: 477259]
* allow httpd_suexec_t var_run_t : dir { search +getattr }; [p1: 87762] [p2: 479746]
* allow httpd_suexec_t var_t : dir { search +getattr }; [p1: 15184, 87762] [p2: 477454, 479744]
* allow httpd_sys_script_t bin_t : lnk_file { getattr read +ioctl +lock }; [p1: 12998] [p2: 463225]
* allow httpd_sys_script_t device_t : dir { search +getattr +ioctl +lock +read }; [p1: 11942, 13018] [p2: 462165, 462191, 46]
* allow httpd_sys_script_t device_t : dir { getattr search +ioctl +lock +read }; [httpd_enable_cgi]:TRUE [p1: 12126] [p2: 463053]
* allow httpd_sys_script_t devtty_t : chr_file { getattr read write +append +ioctl +lock }; [httpd_enable_cgi]:TRUE [p1: 12092] [p2: 463053]
* allow httpd_sys_script_t etc_runtime_t : file { getattr read +ioctl +lock }; [httpd_enable_cgi]:TRUE [p1: 12092] [p2: 463053]
* allow httpd_sys_script_t etc_t : lnk_file { getattr read +ioctl +lock }; [p1: 12998, 13118] [p2: 463349, 463393]
* allow httpd_sys_script_t httpd_sys_script_exec_t : file { endpoint -execute -getattr -ioctl -lock -read }; [httpd_enable_cgi]:TRUE [p1: 12092] [p2: 463053]
* allow httpd_sys_script_t httpd_sys_script_ro_t : file { getattr read -ioctl -lock }; [p1: 13130] [p2: 463055]
* allow httpd_sys_script_t httpd_sys_script_t : process { sigchld sigkill signal signull sigstop -fork }; [httpd_enable_cgi]:TRUE [p1: 12092] [p2: 463053]
* allow httpd_sys_script_t locale_t : lnk_file { getattr read +ioctl +lock }; [httpd_enable_cgi]:TRUE [p1: 12106] [p2: 463053]
* allow httpd_sys_script_t mysql_var_run_t : sock_file { write -append -getattr -ioctl -lock -read }; [p1: 14283] [p2: 464063]
* allow httpd_sys_script_t proc_t : lnk_file { read +getattr }; [httpd_enable_cgi]:TRUE [p1: 12122] [p2: 464063]
* allow httpd_sys_script_t sbin_t : lnk_file { getattr read +ioctl +lock }; [p1: 12998] [p2: 463267]
* allow httpd_sys_script_t sendmail_exec_t : lnk_file { getattr read +ioctl +lock }; [p1: 80456] [p2: 465123]
* allow httpd_sys_script_t sysctl_kernel_t : dir { search +getattr +ioctl +lock +read }; [p1: 14734] [p2: 479809]
* allow httpd_sys_script_t sysctl_kernel_t : file { getattr read +ioctl +lock }; [p1: 14736] [p2: 479811]
* allow httpd_sys_script_t sysctl_t : dir { search +getattr +ioctl +lock +read }; [p1: 14732] [p2: 479807]
* allow httpd_sys_script_t urandom_device_t : chr_file { getattr read +ioctl +lock }; [httpd_enable_cgi]:TRUE [p1: 120] [p2: 464175]
* allow httpd_sys_script_t usr_t : file { getattr ioctl read +lock }; [httpd_enable_cgi]:TRUE [p1: 12078] [p2: 464175]
* allow httpd_sys_script_t usr_t : lnk_file { getattr read +ioctl +lock }; [httpd_enable_cgi]:TRUE [p1: 12080] [p2: 464175]
* allow httpd_sys_script_t var_lib_t : dir { search +getattr }; [p1: 14743] [p2: 479835, 479943]
* allow httpd_sys_script_t var_run_t : dir { search +getattr }; [p1: 87739] [p2: 466083]
* allow httpd_sys_script_t var_t : dir { search +getattr }; [p1: 87739] [p2: 463005, 466081, 479835, 479858, 479943]
* allow httpd_t bin_t : lnk_file { read +getattr +ioctl +lock }; [p1: 10663] [p2: 470430]
    
```

Classes 0 Commons 0 Types: 1097 Attrs: 159 Roles: 5 Users: 4 Bools: 106 TE Rules: 1219615 Role Allows: 4 Role Trans: 0

System Administrator Perspective

- apol



The screenshot shows the SELinux Policy Analysis tool interface. The title bar reads "SELinux Policy Analysis - /tmp/policy.conf". The menu bar includes "File", "Search", "Query", "Advanced", and "Help". The breadcrumb trail is "Policy Components | Policy Rules | File Contexts | Analysis | policy.conf".

The "TE Rules" section is active, showing search criteria:

- Source type/attribute: ^dom
- Target type/attribute: sha.*_t
- Default type: (empty)
- Options: Only direct matches, As source

The "Type Enforcement Rules Display" section shows the results:


```

    Results 1
    9 rules match the search criteria.
    Number of enabled conditional rules: 0
    Number of disabled conditional rules: 2

    [24498] dontaudit auth_chkpwd shadow_t : file {read getattr};
    [46025] dontaudit firstboot_t shadow_t : file getattr;
    [69382] dontaudit local_login_t shadow_t : file {read getattr};
    [70808] dontaudit remote_login_t shadow_t : file {read getattr};
    [87833] allow nscd_t shadow_t : file getattr;
    [102697] allow kernel_t {file_type -shadow_t} : file {ioctl read write create getattr setattr lock append unlink link rename}; [Disabled]
    [102704] dontaudit kernel_t shadow_t : file getattr;
    [102714] allow kernel_t {file_type -shadow_t} : file {ioctl read getattr lock}; [Disabled]
    [108140] allow smbd_t.samba.share.t : file {ioctl read write create getattr setattr lock append unlink link rename};
    
```

The status bar at the bottom shows: "Classes: 55 Perms: 205 Types: 804 Attribs: 123 TE rules: 570443 Roles: 6 Users: 3" and "v.18 (source, non-mls)".

SELinux Usage

(Hints & Tips)

System Administrator Perspective

- **semanage**
Configure elements of SELinux policy without modification/recompilation of policy sources
. . . . aka on the fly

Example: Dynamically Allowing Apache to listen on port 1234

```
# semanage port -a -t httpd_port_t -p tcp 1234
```

System Administrator Perspective

- **semanage** (more examples)

Example: Allow shawn to join “webadmin_u” group

```
# semanage login -a -s webadmin_u shawn
```

Example: Relabel files for access by Apache

```
# semanage fcontext -a -t \  
httpd_sys_content_t "/data/webpages(/.*)?"
```

System Administrator Perspective

- **semanage** (most important example)

You don't need to disable SELinux to fix a single error!

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfce1bc a3=bfce1bc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

The Fix:

```
# semanage permissive -a ifconfig_t
```

System Administrator Perspective

- **audit2allow**

Allows generation of SELinux policy rules from logs of denied operations

Example: Fix all the errors on the system (completely not a good idea on a real system)

```
# cat /var/log/audit/audit.log | audit2allow -M FixAll
```

```
Generating type enforcement file: FixAll.te
```

```
Compiling policy: checkmodule -M -m -o FixAll.mod FixAll.te
```

```
Building package: semodule_package -o FixAll.pp -m FixAll.mod
```

```
# semodule -i FixAll.pp
```

Scenarios

Scenario: Fixing the RHT corporate VPN “update”

- Red Hat has a Corporate Standard Build (CSB) for desktop environments
- Red Hat pushes updates to said CSB
- I “tweak” my configuration files
- When RHT pushed a CSB update, it broke my VPN settings

Scenario: Fixing the RHT corporate VPN “update”

/var/log/messages:

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfce1bc a3=bfce1bc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

Now what?

Scenario: Fixing the RHT corporate VPN “update”

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfceclbc a3=bfceclbc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

What I Know:

- 1) AVC Event ID 738
- 2) syscall=54 (I'd have to google this)
- 3) root (or an application on its behalf) was running /sbin/ip
- 4) context = user_u:system_r:ifconfig_t:s0

Scenario: Fixing the RHT corporate VPN “update”

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfceclbc a3=bfceclbc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

My Options:

1) Create a SELinux Policy Module

```
# ausearch -x "/sbin/ip" | audit2allow -M MyVPNFix
```

Scenario: Fixing the RHT corporate VPN “update”

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfceclbc a3=bfceclbc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

My Options:

1) Create a SELinux Policy Module

```
# ausearch -x "/sbin/ip" | audit2allow -M MyVPNFix
# semodule -i MyVPNFix.pp
```

Scenario: Fixing the RHT corporate VPN “update”

```
type=SYSCALL msg=audit(1204719775.306:738): arch=40000003 syscall=54
success=no exit=-19 a0=4 a1=8933 a2=bfceclbc a3=bfceclbc items=0
ppid=3900 pid=5003 auid=501 uid=0 gid=0 euid=0 suid=0 fsuid=0 egid=0
sgid=0 fsgid=0 tty=(none) comm="ip" exe="/sbin/ip"
subj=user_u:system_r:ifconfig_t:s0 key=(null)
```

My Options:

2) Disable enforcement of `ifconfig_t` (**there is no need to turn SELinux completely off!**)

```
# semanage permissive -a ifconfig_t
```



What'd I forget? Open Discussion