Practical SELinux: Writing Custom Application Policy

Miroslav Grepl
Lukas Vrabec
Simon Sekidde
Thursday, May 4, 10:15 AM - 12:15 PM

redhat.

Agenda

- Proactive Security
- SELinux Security Policy
- Updated Userspace with Easier Policy Customization
- SELinux Awareness
- Writing SELinux Policy
- Troubleshooting Existing Policy

redhat.

Proactive Security

redhat.

WHEN DO PEOPLE CARE ABOUT SECURITY?

redhat.

redhat.

HOW DO SECURITY ISSUES AFFECT ME?

redhat.

LOST/UNWITTINGLY SHARED/COMPROMISED PERSONAL DATA

redhat.

LOST/UNWITTINGLY SHARED/COMPROMISED PERSONAL DATA

ONLINE MONEY-RELATED THEFTS

redhat.

LOST/UNWITTINGLY SHARED/COMPROMISED

PERSONAL DATA

ONLINE MONEY-RELATED THEFTS

MOBILE DEVICE PROTECTION

redhat.

WHERE DO SECURITY ISSUES COME FROM?

redhat.

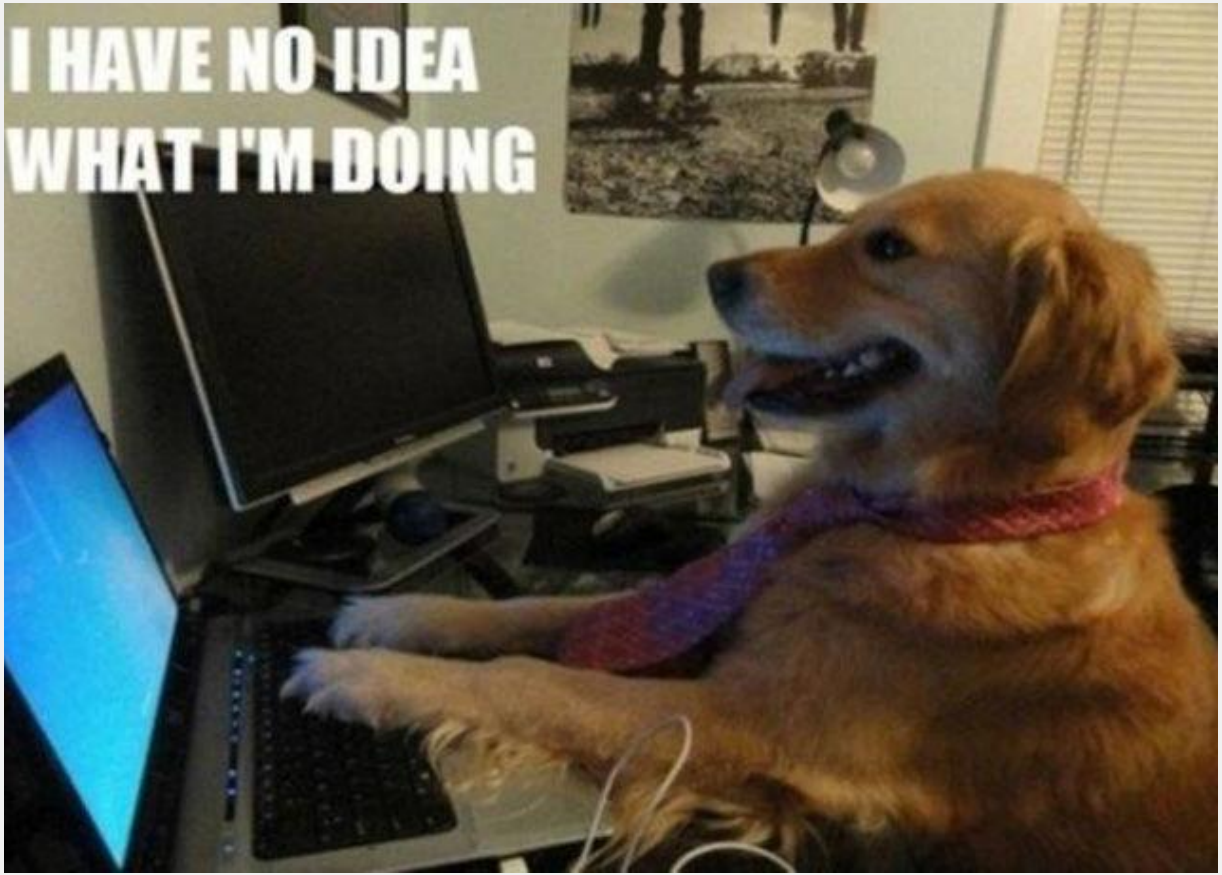
I HAVE NO IDEA
WHAT I'M DOING

HOW ARE THEY FIXED?

redhat.

REACTIVE SECURITY

redhat.

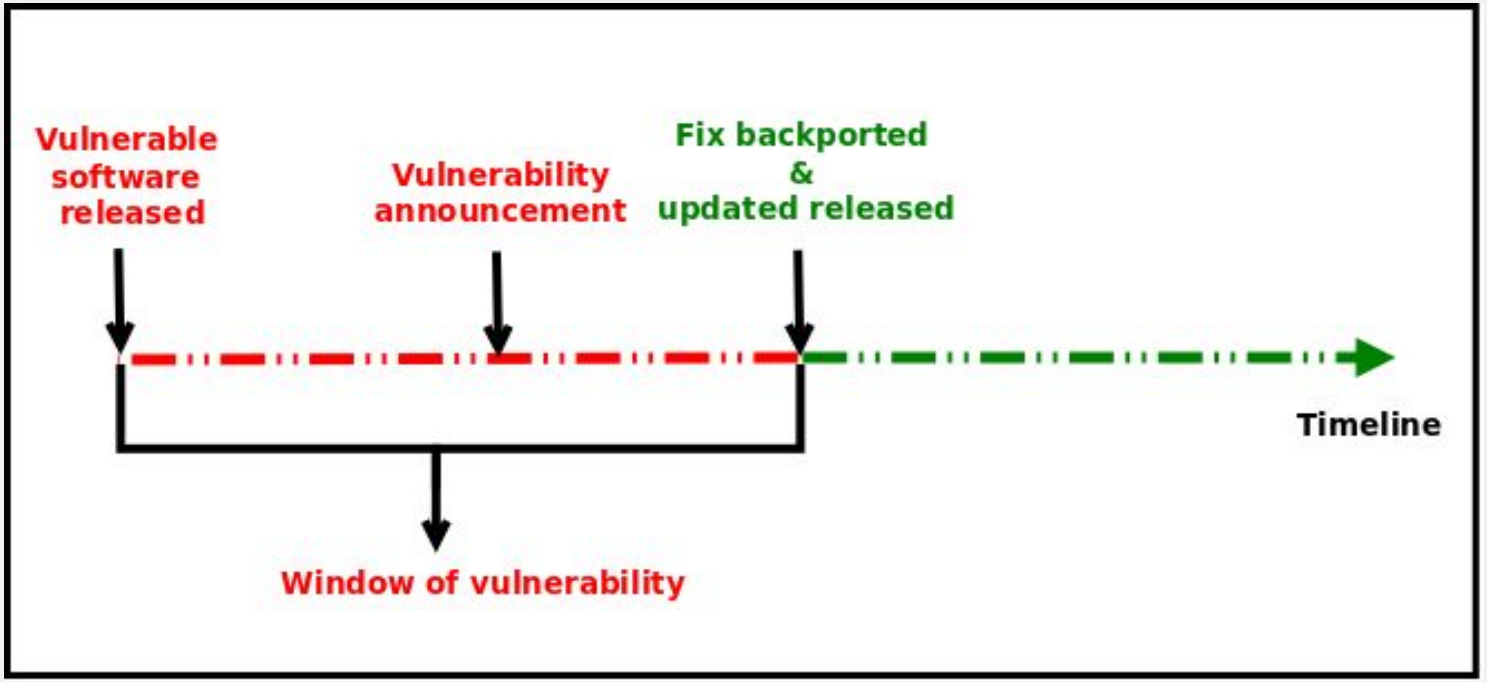
Vulnerable software released
Vulnerability announcement
Fix backported & updated released
Timeline
Window of vulnerability

YOUR SYSTEM **IS NOT PROTECTED** DURING THE WINDOW OF VULNERABILITY!

redhat.

PROACTIVE SECURITY

redhat.

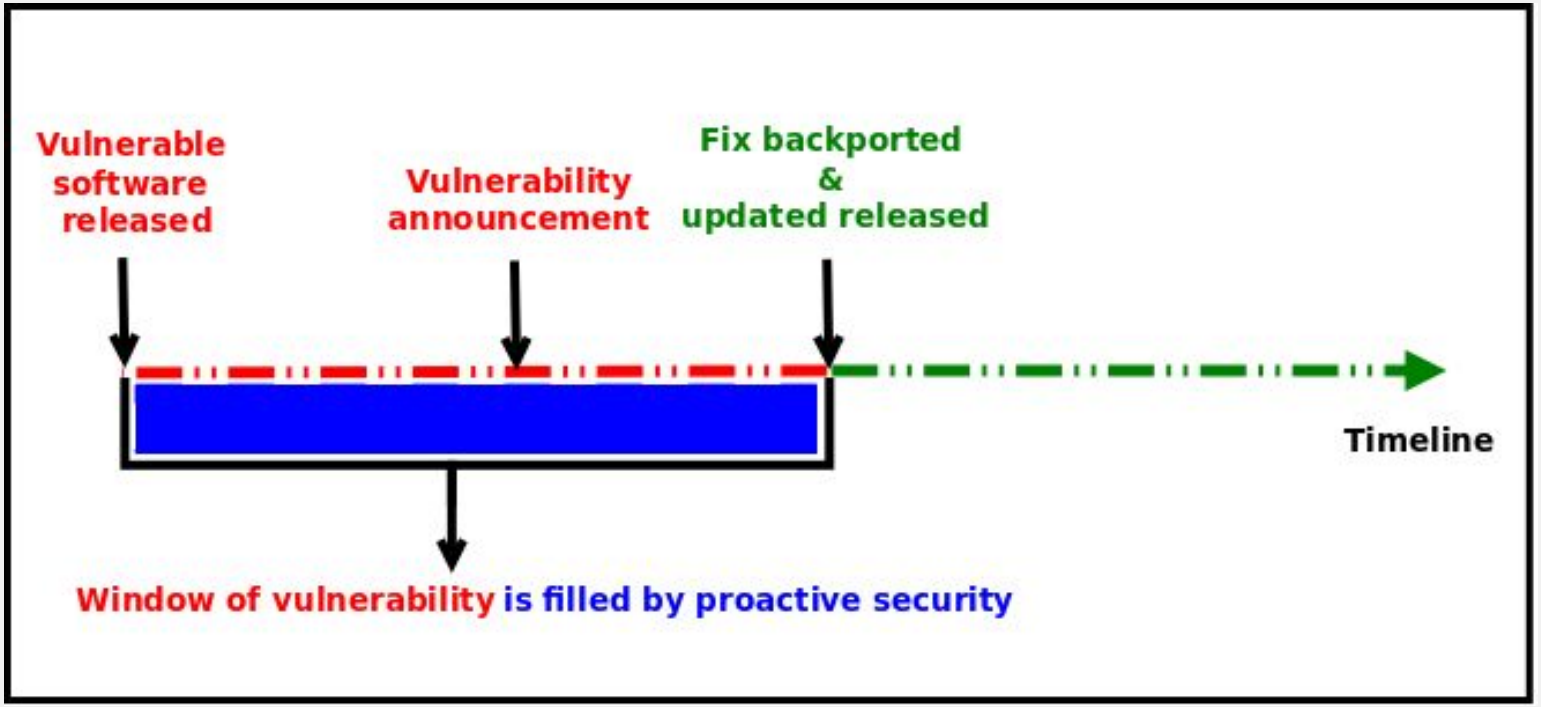
Vulnerable software released
Vulnerability announcement
Fix backported & updated released
Timeline
Window of vulnerability is filled by proactive security

PROACTIVE SECURITY HELPS TO **PROTECT** YOUR SYSTEM DURING THE WINDOW OF VULNERABILITY!

redhat.

SECURITY ENHANCED LINUX IS A SECURITY MECHANISM BRINGING PROACTIVE SECURITY FOR YOUR SYSTEM.

redhat.

EXPLOIT EXAMPLES WHERE SELINUX HELPED TO PROTECT YOUR SYSTEM

redhat.

VENOM

redhat.

VENOM

DOCKER CVE-2016-9962

redhat.

VENOM

DOCKER CVE-2016-9962

SHELLSHOCK

redhat.

HACKING TIME!

redhat.

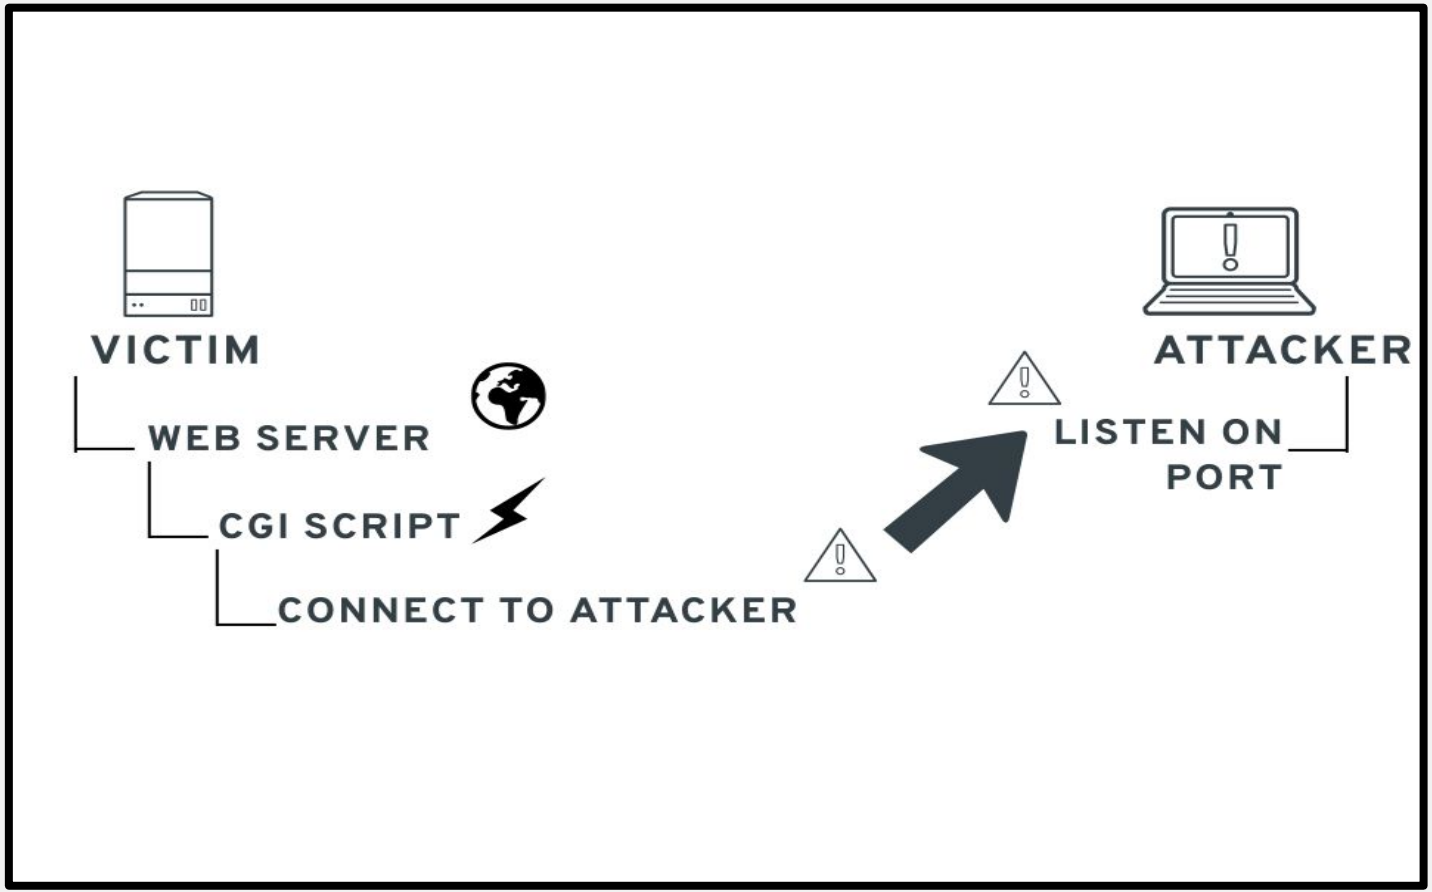
VICTIM
WEB SERVER
CGI SCRIPT
CONNECT TO ATTACKER
ATTACKER
LISTEN ON PORT

redhat.

DEMO TIME!

redhat.

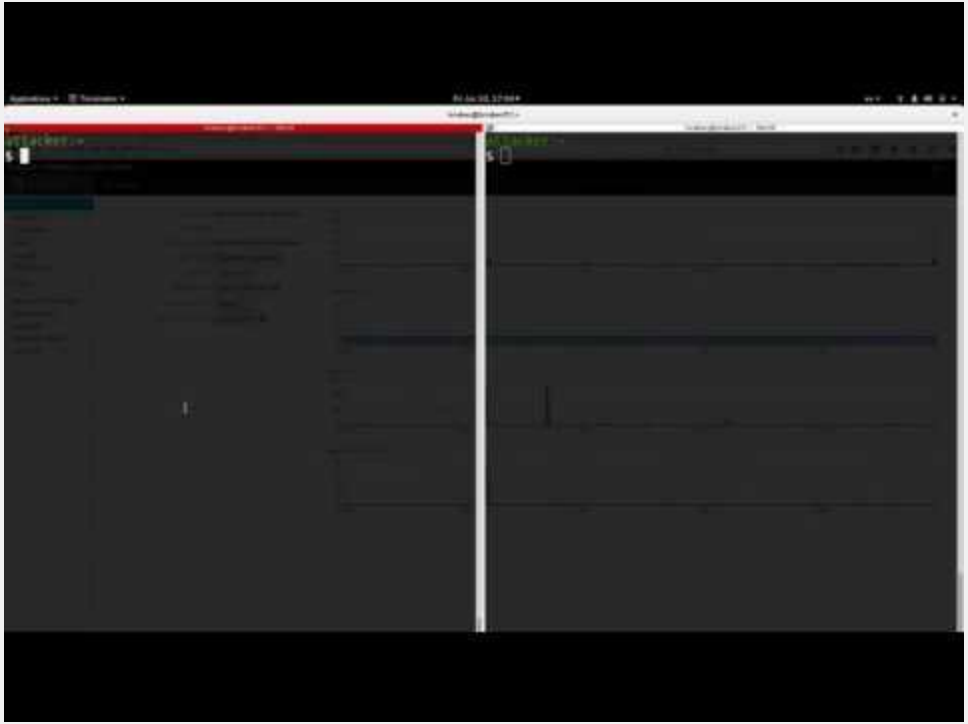

redhat.

CONCLUSION?

redhat.


IF YOU RUN LINUX WITH SELINUX DISABLED
INSTRUCTOR
YOUR GONNA HAVE A BAD TIME
memegenerator.net

SELinux Security Policy

redhat.

IS

redhat.

CORE COMPONENT OF SELINUX

redhat.

CORE COMPONENT OF SELINUX

COLLECTION OF SELINUX POLICY RULES

redhat.

CORE COMPONENT OF SELINUX

COLLECTION OF SELINUX POLICY RULES

LOADED INTO THE KERNEL BY SELINUX USERSPACE TOOLS

redhat.

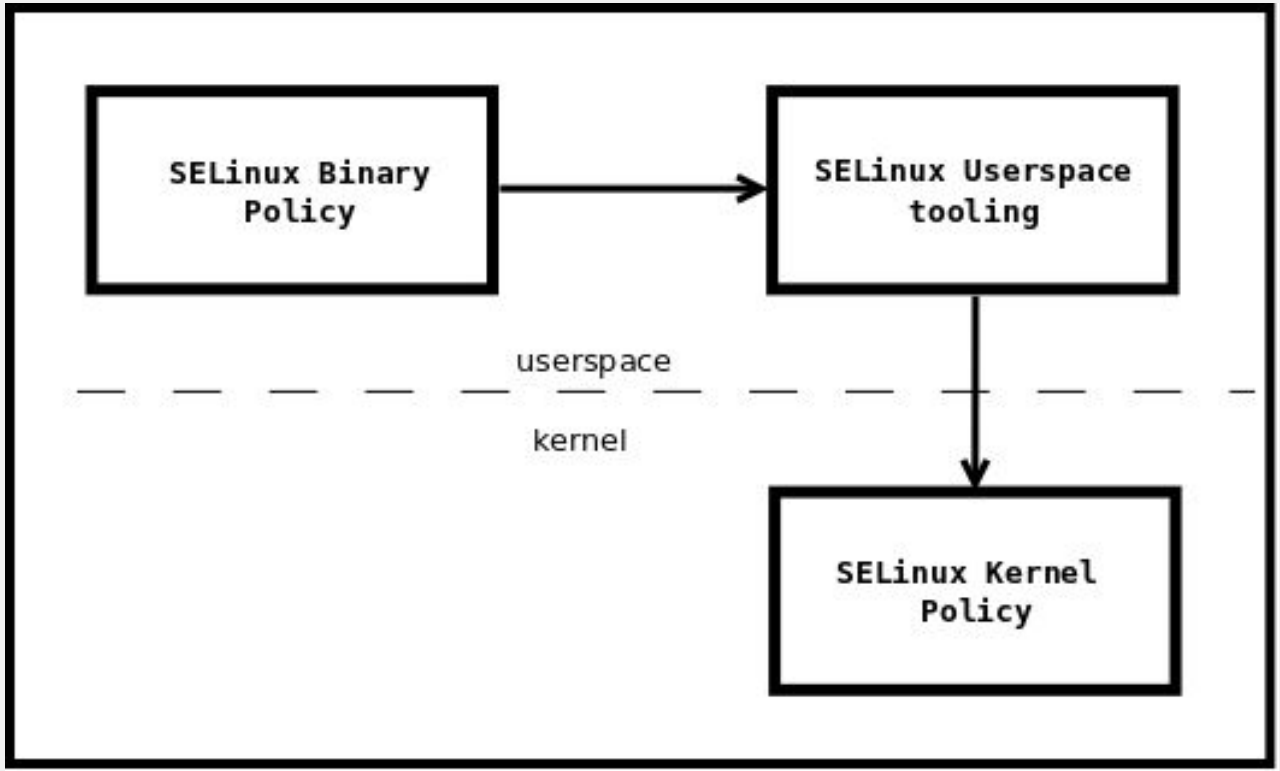
SELinux Binary Policy
SELinux Userspace tooling
userspace
kernel
SELinux Kernel Policy

redhat.

ENFORCED BY THE KERNEL

redhat.

ENFORCED BY THE KERNEL

USED TO AUTHORIZE ACCESS REQUESTS ON THE SYSTEM

redhat.

BY DEFAULT **EVERYTHING** IS DENIED AND YOU DEFINE POLICY RULES TO ALLOW CERTAIN REQUESTS.

redhat.

SELINUX POLICY RULES

redhat.

DESCRIBE AN **INTERACTION** BETWEEN PROCESSES AND SYSTEM RESOURCES

redhat.

SELINUX POLICY RULE IN HUMAN LANGUAGE

redhat.

"APACHE process can READ its LOGGING FILE"

redhat.

SELINUX VIEW OF THAT INTERACTION

redhat.

```
ALLOW apache_process apache_log:FILE
READ;
```


redhat.

`apache_process` `apache_log`

ARE **LABELS**

redhat.

LABELS

redhat.

ASSIGNED TO PROCESSES

redhat.

ASSIGNED TO PROCESSES

ASSIGNED TO SYSTEM RESOURCES

redhat.

ASSIGNED TO PROCESSES

ASSIGNED TO SYSTEM RESOURCES

BY SELINUX SECURITY POLICY

redhat.

ASSIGNED TO PROCESSES

ASSIGNED TO SYSTEM RESOURCES

BY SELINUX SECURITY POLICY

MAP REAL SYSTEM ENTITIES INTO THE SELINUX WORLD

redhat.

LABELS IN REALITY

redhat.

STORED IN EXTENDED ATTRIBUTES OF FILE SYSTEMS - EXT2,EXT3, EXT4 ...

redhat.

```
# getfattr -n security.selinux /etc/passwd
getfattr: Removing leading '/' from absolute path
names
# file: etc/passwd
security.selinux="system_u:object_r:passwd_file_t:s0"

# ls -Z /etc/passwd
system_u:object_r:passwd_file_t:s0 /etc/passwd
```


redhat.

SELINUX LABELS CONSIST OF **FOUR** PARTS

redhat.

<user>:<role>:<type>:<MLS/MCS>

redhat.

<user>:<role>:<type>:<MLS/MCS>

Not the same as Linux users

Several Linux users can be mapped to a single SELinux user

object_u is a placeholder for Linux system resources

system_u is a placeholder for Linux processes

Can be limited to a set of SELinux roles

redhat.

<user>:<role>:<type>:<MLS/MCS>

SELinux users can have multiple roles but only one can be active

object_r is a placeholder for Linux system resources

system_r is a placeholder for system processes

Can be limited to a set of SELinux types

redhat.

<user>:<role>:<type>:<MLS/MCS>

Security model known as **TYPE ENFORCEMENT**

In 99% you care only about TYPES

policy rules and interactions between types

redhat.

<user>:<role>:<type>:<MLS/MCS>

Multi Level Security

Only the MCS part is used in Targeted Policy with the default *s0* level

Allow users to mark resources with compartment tags *(MCS1*, *MCS2*)

Used for RHEL virtualization and for container security

s0:c1 can not access *s0:c2*

redhat.

IN RHEL7 WE SHIP THE **TARGETED** SELINUX POLICY BY DEFAULT

redhat.

WE MOSTLY CARE ONLY ABOUT **TYPES**

redhat.

SELINUX **ALLOW** RULE SYNTAX WITH **TYPES**

redhat.

```
ALLOW TYPE1 TYPE2:OBJECT_CLASS
PERMISSION;
```


redhat.

```
ALLOW APACHE_T APACHE_LOG_T:FILE READ;
```


redhat.

Updated Userspace with Easier Policy Customization

redhat.

FRIENDLY SELINUX?

redhat.

NEW SELINUX USERSPACE 2.5 INTRODUCED IN RHEL-7.3

redhat.

NEW COMMON INTERMEDIATE LANGUAGE - CIL

redhat.

”M4+COMPILATION” VS. CIL

redhat.

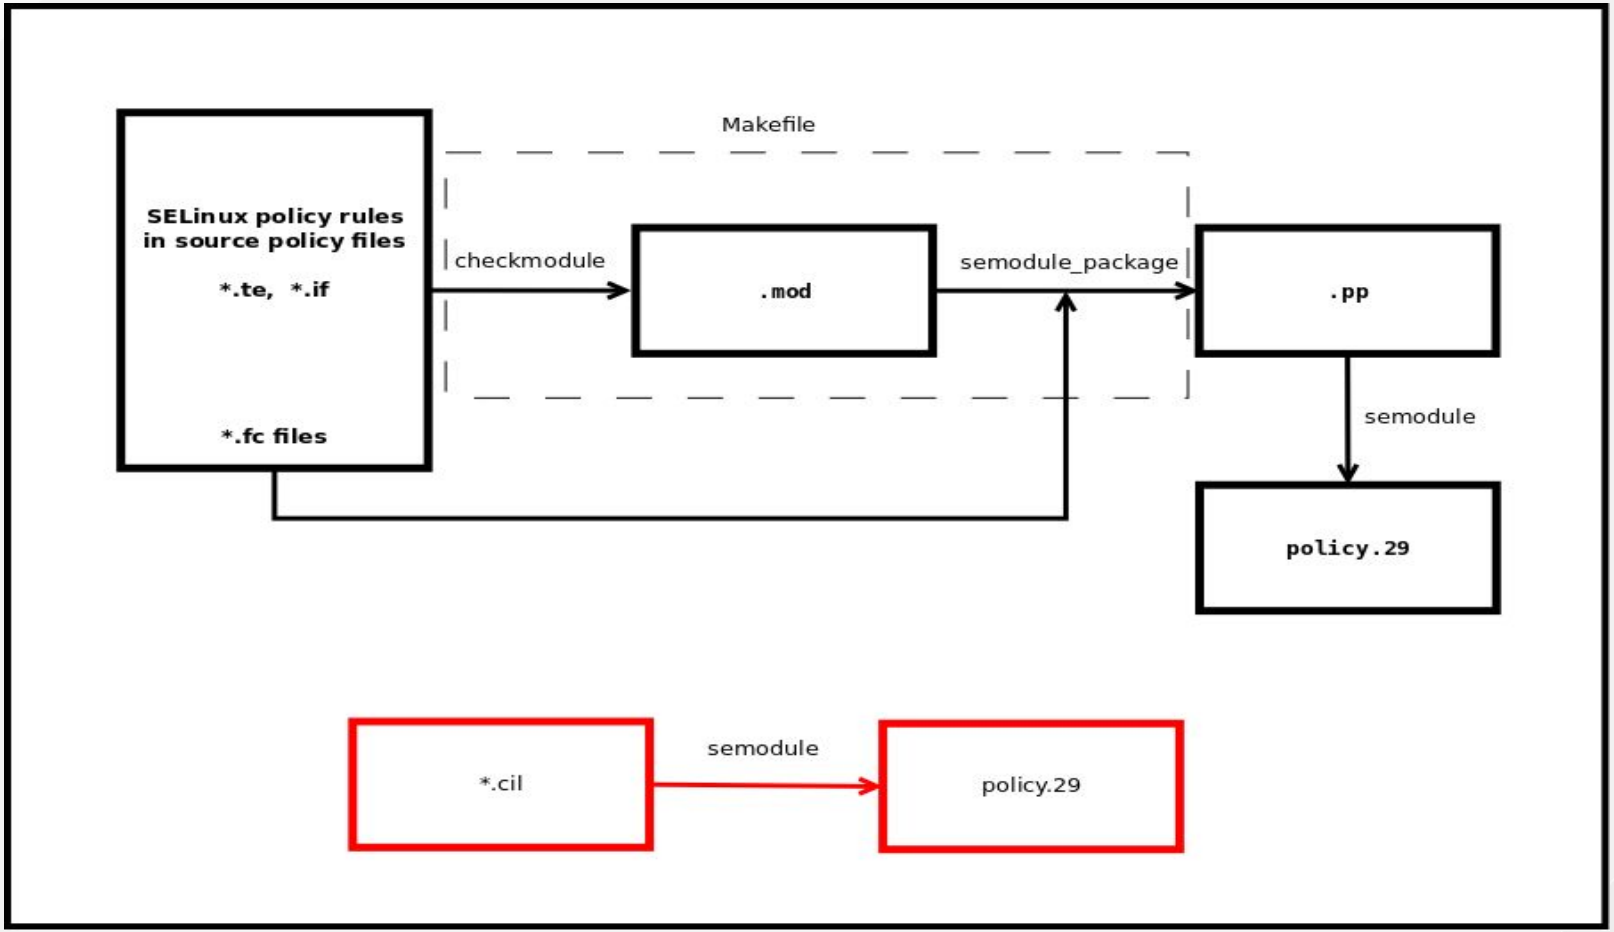
Makefile
SELinux policy rules
in source policy files
*.te, *.if
*.fc files
checkmodule
.mod
semodule_package
.pp
semodule
policy.29
*.cil
semodule
policy.29

redhat.

PERFORMANCE IMPROVEMENTS

redhat.

PERFORMANCE IMPROVEMENTS

NEW POSSIBILITY FOR HLL

redhat.

PERFORMANCE IMPROVEMENTS

NEW POSSIBILITY FOR HLL

USABILITY

redhat.

LOCAL POLICY IN TWO STEPS

redhat.

"I have an apache process that needs to access its log file. I would like to add SELinux policy rule reflecting the following interaction so that I am able to read important info from the apache logging file."

redhat.

```
# cat myapache.cil
(allow httpd_t httpd_log_t (file (open read
getattr)))
```


redhat.

```
# semodule -i myapache.cil
```


redhat.

HOW DO WE DO IT WITH M4 + COMPILATION?

redhat.

```
# cat myapache.te
require {
 type httpd_t;
 type httpd_log_t;
}
allow httpd_t httpd_log_t:file { open read
getattr };
```


redhat.

```
# make -f /usr/share/selinux/devel/Makefile
# semodule -i myapache.pp
```


redhat.

FUTURE

redhat.

redhat.

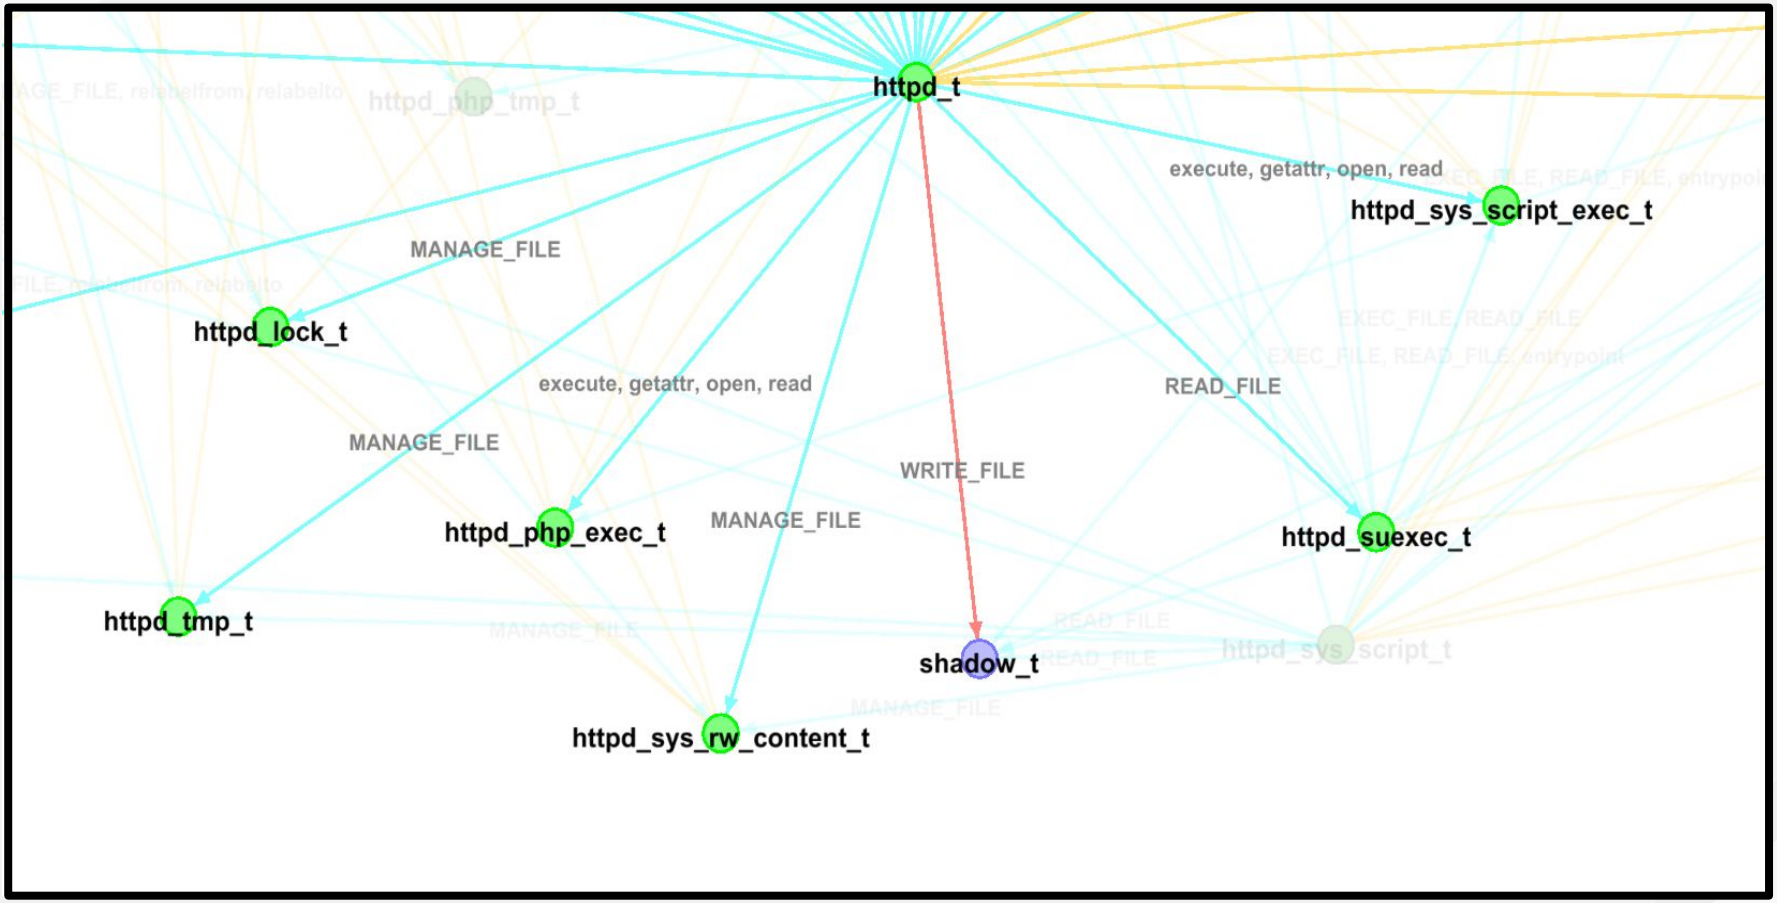
httpd_t
execute, getattr, open, read
httpd_sys_script_exec_t
MANAGE_FILE
httpd_lock_t
execute, getattr, open, read
READ_FILE
MANAGE_FILE
WRITE_FILE
MANAGE_FILE
httpd_php_exec_t
httpd_suexec_t
httpd_tmp_t
shadow_t
httpd_sys_rw_content_t

redhat.

SELinux Awareness

redhat.

SELINUX ENHANCING TECHNOLOGIES

redhat.

SYSTEMD

redhat.

SYSTEMD WORKS AS AN SELINUX ACCESS MANAGER

redhat.

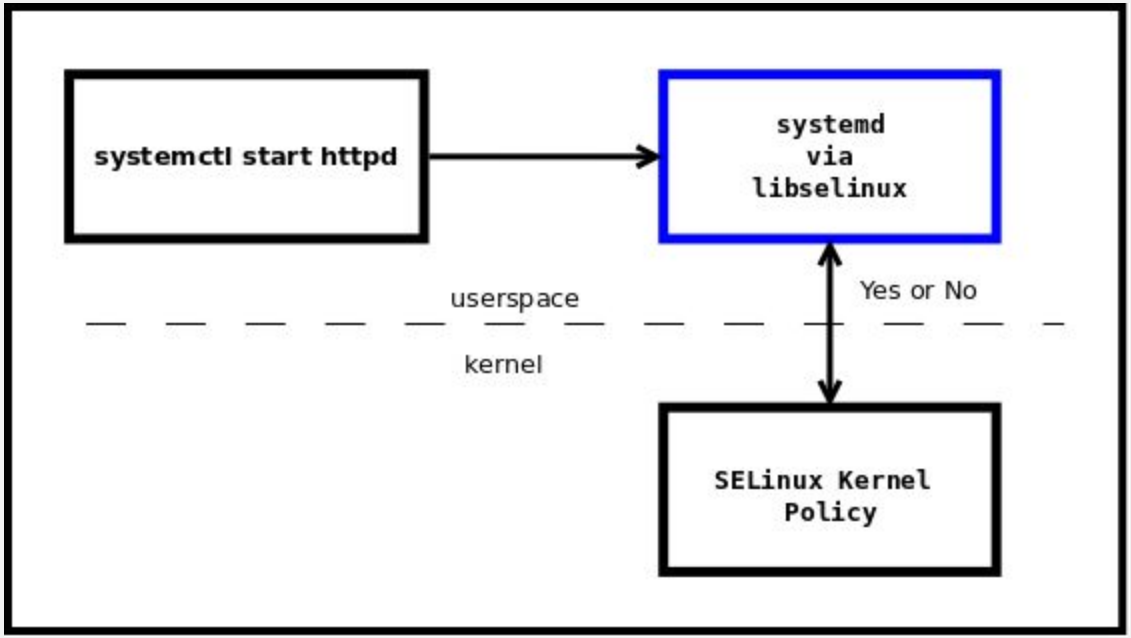
systemctl start httpd
systemd
via
libselinux
Yes or No
userspace
kernel
SELinux Kernel
Policy

redhat.

DO YOU REMEMBER SELINUX ALLOW RULE SYNTAX WITH **TYPE ENFORCEMENT**?

redhat.

```
ALLOW TYPE1 TYPE2:OBJECT_CLASS
PERMISSION;
```


redhat.

```
ALLOW HTTPD_T HTTPD_UNIT_FILE_T:SERVICE
START;
```


redhat.

SYSTEMD

SVIRT

redhat.

APPLIES MAC TO IMPROVE SECURITY WHEN USING VIRTUAL MACHINES

redhat.

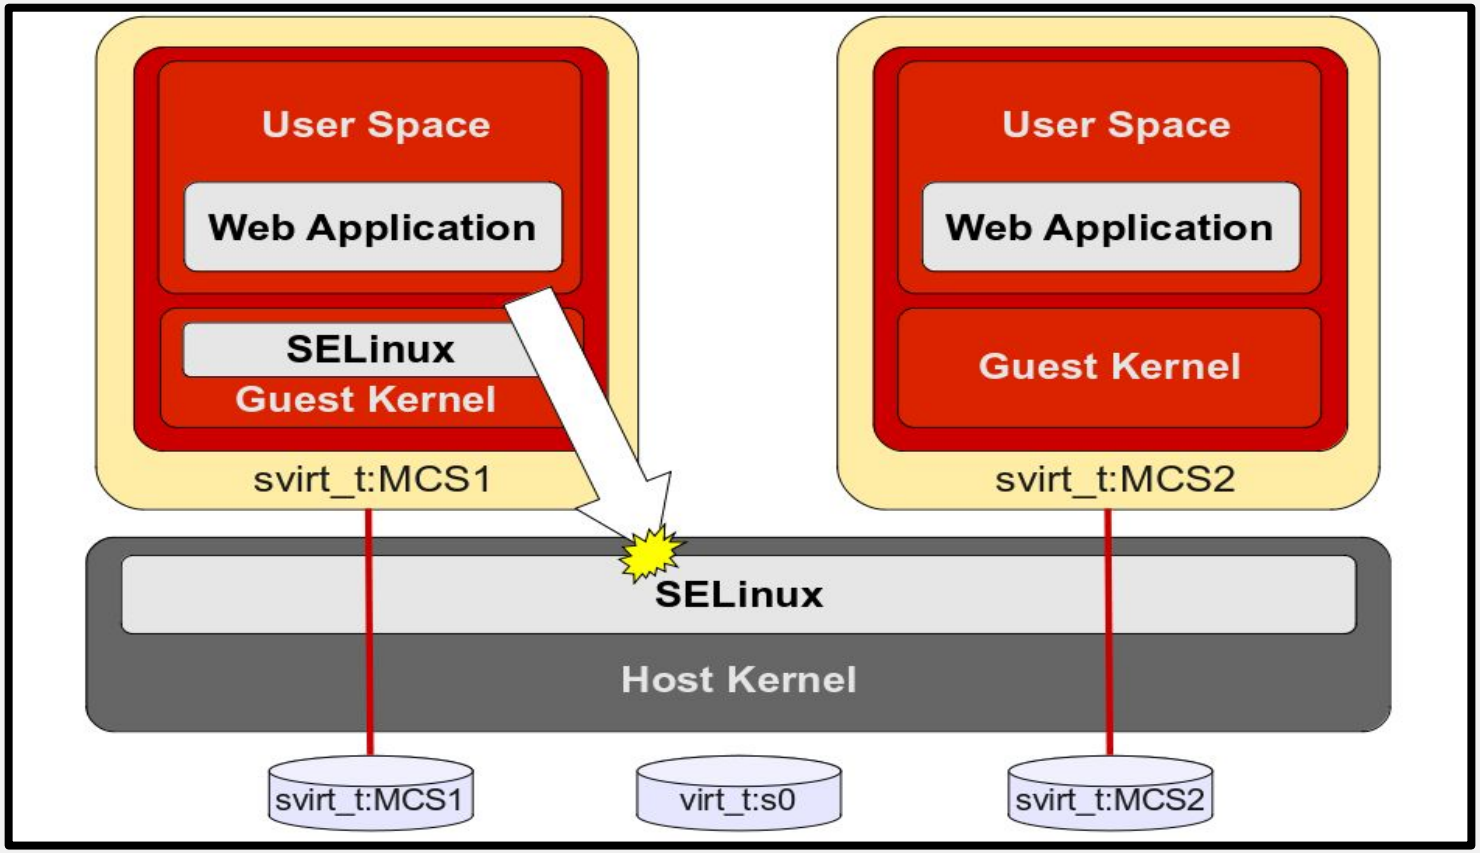
User Space
Web Application
SELinux
Guest Kernel
svirt_t:MCS1
User Space
Web Application
Guest Kernel
svirt_t:MCS2
SELinux
Host Kernel
svirt_t:MCS1
virt_t:s0
svirt_t:MCS2

redhat.

SELinux user:SELinux role:SELinux type:SELinux category

redhat.

```
SELinux user:SELinux role:SELinux type:SELinux category

system_u:object_r:svirt_t:c306,c536
```


redhat.

SELinux user:SELinux role:SELinux type:SELinux category

system_u:object_r:svirt_t:c306,c536

system_u:object_r:svirt_t:c206,c636

redhat.

SYSTEMD

SVIRT

CONTAINERS

redhat.

SELINUX KEEPS YOUR CONTAINER IN ITS OWN SPACE

redhat.

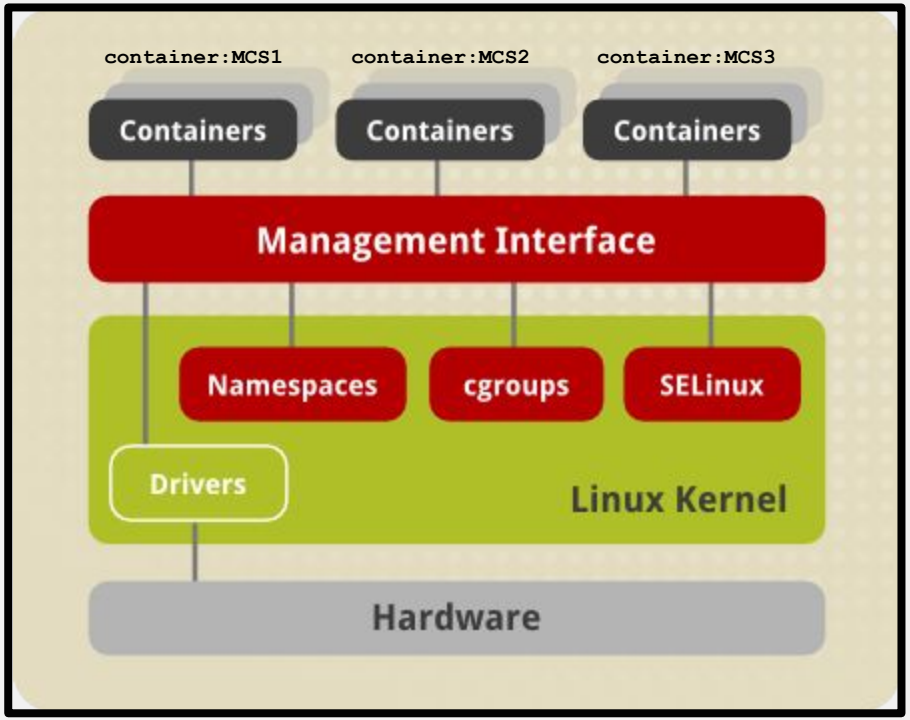
container:MCS1
container:MCS2
container:MCS3
Containers
Containers
Containers
Management Interface
Namespaces
cgroups
SELinux
Drivers
Linux Kernel
Hardware

redhat.

SELinux user:SELinux role:SELinux type:SELinux category

redhat.

```
SELinux user:SELinux role:SELinux type:SELinux category

system_u:object_r:container_t:c306,c536
```


redhat.

```
SELinux user:SELinux role:SELinux type:SELinux category
      system_u:object_r:container_t:c306,c536
      system_u:object_r:container_t:c206,c636
```


redhat.

```
SELinux user:SELinux role:SELinux type:SELinux category
      system_u:object_r:container_t:c306,c536
      system_u:object_r:container_t:c206,c636
      system_u:object_r:container_t:c406,c736
```


redhat.

SELINUX MANAGEMENT

redhat.

SELINUX TROUBLESHOOTING IN COCKPIT

redhat.

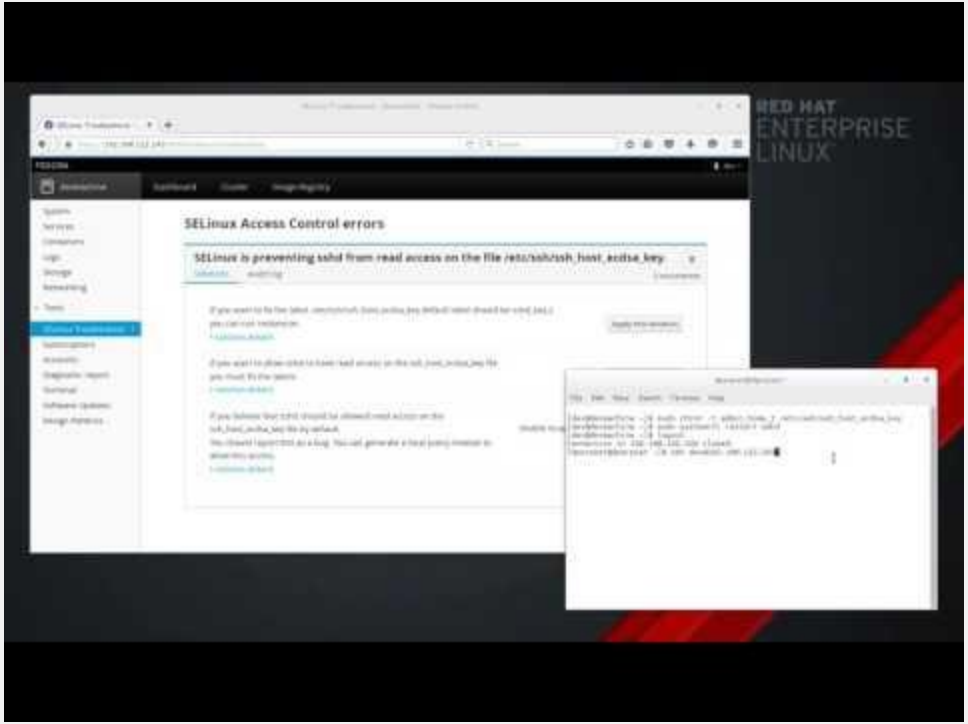
RED HAT
ENTERPRISE
LINUX
SELinux Access Control errors
SELinux is preventing sshd from read access on the file /etc/ssh/ssh_host_ecdsa_key

redhat.

SELINUX TROUBLESHOOTING IN COCKPIT

SELINUX USERSPACE INSIDE CONTAINER

redhat.

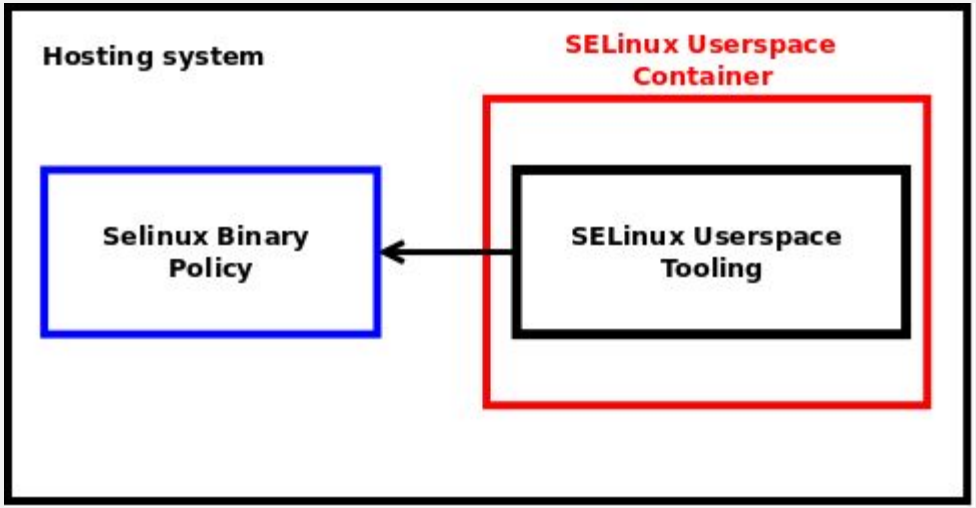
Hosting system
SELinux Userspace
Container
Selinux Binary
Policy
SELinux Userspace
Tooling

redhat.

```
$ getenforce
Permissive

$ sudo docker run --privileged -v /sys/fs/selinux:/sys/fs/selinux:rw
-v/etc/selinux:/etc/selinux:rw -v /etc/selinux:/var/lib/selinux:rw -it
selinux-container:latest setenforce 1

$ getenforce
Enforcing
```


redhat.

SELINUX TROUBLESHOOTING COCKPIT

SELINUX USERSPACE INSIDE CONTAINER

SELINUX WITH ANSIBLE

redhat.

https://github.com/cockpit-project/system-api-roles/tree/master/roles/SELinux

redhat.

DEMO PLAYBOOK

redhat.

```
$ cat > demo-playbook.yml <<EOF
---
- hosts: all
  remote_user: root
  vars:
    SELinux_type: targeted
    SELinux_mode: enforcing
    SELinux_change_running: 1
    SELinux_booleans:
      - { name: 'samba_enable_home_dirs', state: 'on' }
      - { name: 'ssh_sysadm_login', state: 'on', persistent: 'yes' }
    SELinux_file_contexts:
      - { target: '/tmp/test_dir(/.*)?', setype: 'user_home_dir_t', ftype: 'd' }
    SELinux_restore_dirs:
      - /tmp/test_dir
  roles:
    - SELinux
EOF
```


redhat.

RUN THE PLAYBOOK USING SELINUX ROLES

redhat.

```
$ ansible-playbook -i
mgrepl-rhel-73.virt,mgrepl-rhel-6.virt,mgrepl-fedora-25.virt
, demo-playbook.yml
```


redhat.

Writing SELinux Policy

redhat.

SELINUX MODES

redhat.

ENFORCING

redhat.

ENFORCING

SELINUX SECURITY POLICY IS ENFORCED BY KERNEL

redhat.

PERMISSIVE

redhat.

PERMISSIVE

SELINUX SECURITY POLICY IS NOT ENFORCED BY KERNEL

redhat.

PERMISSIVE

SELINUX SECURITY POLICY IS NOT ENFORCED BY KERNEL

ACCESSES ARE LOGGED

redhat.

LET’S SWITCH SELINUX TO PERMISSIVE TO COLLECT MORE AVC MESSAGES

redhat.

AVC MESSAGES

redhat.

WHERE CAN WE FIND LOGS?

redhat.

```
# cat /var/log/audit/audit.log
```


redhat.

```
# cat /var/log/audit/audit.log
# ausearch -m AVC
```


redhat.

```
type=AVC msg=audit(1226882925.714:136): avc:  denied
  { read } for  pid=2512 comm="httpd" name="file1"
                dev=dm-0 ino=284133
     scontext=unconfined_u:system_r:httpd_t:s0
     tcontext=unconfined_u:object_r:shadow_t:s0
                tclass=file
```


redhat.

HOW TO PARSE AVC MESSAGES?

redhat.

```
# sesearch
```


redhat.

```
# sesearch
# audit2allow
```


redhat.

```
$ ausearch -m AVC -ts recent

type=AVC msg=audit(1226882925.714:136): avc:  denied  { read } for
pid=2512 comm="httpd" name="shadow" dev=dm-0 ino=284133
scontext=unconfined_u:system_r:httpd_t:s0
tcontext=unconfined_u:object_r:shadow_t:s0 tclass=file

$ ausearch -m AVC -ts recent | audit2allow

#============= httpd_t ==============

allow httpd_t shadow_t:file read;
```


redhat.

ALL NECESSARY FILES FOR WRITING CUSTOM POLICY CAN BE FOUND IN **/USR/SHARE/SELINUX/DEVEL** DIRECTORY

redhat.

rhsummit.service

redhat.

DUMMY LINUX DAEMON FOR TESTING PURPOSE

redhat.

DUMMY LINUX DAEMON FOR TESTING PURPOSE

WE WILL WRITE POLICY FOR IT

redhat.

DUMMY LINUX DAEMON FOR TESTING PURPOSE

WE WILL WRITE POLICY FOR IT

ACTIONS

redhat.

Connecting on port 80 tcp on lvrabec-selinux.rhcloud.com

Logging messages into journal

Creating pid file

Reading /proc/meminfo

redhat.

```
# systemctl status rhsummit
● rhsummit.service - Testing SELinux for Red Hat Summit 2017
   Loaded: loaded (/usr/lib/systemd/system/rhsummit.service; disabled; vendor preset: disabled)
   Active: inactive (dead)
```


redhat.

```
# systemctl start rhsummit

# systemctl status rhsummit
● rhsummit.service - Testing SELinux for Red Hat
Summit 2017
   Loaded: loaded
(/usr/lib/systemd/system/rhsummit.service; disabled;
vendor preset: disabled)
   Active: active (running) since Tue 2017-04-11
16:16:47 CEST; 9s ago
  Process: 2827 ExecStart=/usr/bin/rhsummit
(code=exited, status=0/SUCCESS)
```


redhat.

NO SELINUX POLICY FOR THE *RHSUMMIT* SERVICE
MEANS THE SERVICE IS UNCONFINED SERVICE
FROM SELINUX POINT OF VIEW!

redhat.

```
# ps -efZ | grep rhsummit
system_u:system_r:unconfined_service_t:s0 root 2828 1
0 16:16 ?    00:00:00 /usr/bin/rhsummit
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
root 2849 2344  0 16:33 pts/0 00:00:00 grep
--color=auto rhsummit
```


redhat.

```
# sepolicy generate --init /usr/bin/rhsummit
Loaded plugins: product-id
Created the following files:
/root/code/policy/rhsummit.te # Type Enforcement file
/root/code/policy/rhsummit.if # Interface file
/root/code/policy/rhsummit.fc # File Contexts file
/root/code/policy/rhsummit_selinux.spec # Spec file
/root/code/policy/rhsummit.sh # Setup Script

# ./rhsummit.sh
```


redhat.

HOW TO CHECK AN SELINUX STATUS FOR OUR SUMMIT SERVICE?

redhat.

```
# systemctl stop rhsummit
# systemctl start rhsummit
# ps -efZ | grep summit

# ps -efZ | grep summit
system_u:system_r:rhsummit_t:s0 root      3049    1  0 17:03
?       00:00:00 /usr/bin/rhsummit
unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023 root
3051 2344  0 17:03 pts/0 00:00:00 grep --color=auto summit
```


redhat.

RHSUMMIT SERVICE RUNS WITH **RHSUMMIT** LABEL
WHICH MEANS THAT THE SERVICE IS **CONFINED**!

redhat.

IMPORTANT SELINUX POLICY FILES

redhat.

rhsummit.fc file contains SELinux security context for all service objects (files, dirs, lnk_files, sockets, ...)

rhsummit.te file contains rules for rhsummit_t domain

rhsummit.if file contains interfaces to access rhsummit_t types

rhsummit.sh bash script for compiling and installing SELinux policy for rhsummit service

redhat.

WHY THE RHSUMMIT SERVICE RUNS AS RHSUMMIT_T DOMAIN?

redhat.

SELINUX PROCESS TRANSITION RULES!

redhat.

```
# sesearch -T -s init_t -t rhsummit_exec_t
Found 1 semantic te rules:
   type_transition init_t rhsummit_exec_t : process
rhsummit_t;
```


redhat.

AVCS IN AUSEARCH OUTPUTS

redhat.

```
# ausearch -m AVC -ts recent
…
…
...

# ausearch -m AVC -ts recent > ~/avc_file
```


redhat.

AVCS RELATED TO CREATING PID FILES

redhat.

/var/run/rhsummit.pid should have custom label **rhsummit_pid_t** instead of **var_run_t**

```
type=AVC msg=audit(1491928595.618:342): avc:  denied  { write } for
pid=3924 comm="rhsummit" name="rhsummit.pid" dev="tmpfs" ino=35478
                scontext=system_u:system_r:rhsummit_t:s0
      tcontext=unconfined_u:object_r:var_run_t:s0 tclass=file
```


redhat.

rhsummit.te

```
+ type rhsummit_pid_t;

+ files_pid_file(rhsummit_pid_t)

+ manage_files_pattern(rhsummit_t, rhsummit_pid_t, rhsummit_pid_t)

+ files_pid_filetrans(rhsummit_t, rhsummit_pid_t, { file })
```

rhsummit.fc

```
+ /var/run/rhsummit.* -- gen_context(system_u:object_r:rhsummit_pid_t,s0)
```

AVCS RELATED TO READING /PROC

redhat.

rhsummit process reads /proc/meminfo file, this access should be allowed

```
type=AVC msg=audit(1491928595.618:343): avc:  denied  { read }
   for  pid=3924 comm="rhsummit" name="meminfo" dev="proc"
   ino=4026532028 scontext=system_u:system_r:rhsummit_t:s0
      tcontext=system_u:object_r:proc_t:s0 tclass=file
```


redhat.

rhsummit.te

```
+ kernel_read_system_state(rhsummit_t)
```

Interface can be found in

```
/usr/share/selinux/devel/include/kernel/kernel.if
```


redhat.

AVCS RELATED TO NETWORK ACCESS

redhat.

service connects to “lvrabec-selinux.rhcloud.com”, this should be allowed

```
type=AVC msg=audit(1491939758.424:387): avc:  denied  {
 name_connect } for  pid=4430 comm="rhsummit" dest=80
      scontext=system_u:system_r:rhsummit_t:s0
tcontext=system_u:object_r:http_port_t:s0 tclass=tcp_socket
```


redhat.

rhsummit.te

```
+  corenet_tcp_connect_http_port(rhsummit_t)
```

Interface can be found in

```
/usr/share/selinux/devel/include/kernel/corenetwork.if
```


redhat.

Service resolves DNS records, so access to /etc/resolv.conf is needed

```
type=AVC msg=audit(1491942725.855:436): avc:  denied  { read }
 for  pid=4874 comm="rhsummit" name="resolv.conf" dev="dm-0"
    ino=5051289 scontext=system_u:system_r:rhsummit_t:s0
    tcontext=system_u:object_r:net_conf_t:s0 tclass=file
```


redhat.

rhsummit.te

```
+  sysnet_read_config(rhsummit_t)
```

Interface can be found in

```
/usr/share/selinux/devel/include/system/sysnetwork.if
```


redhat.

AVCS RE-CHECK FOR THE *RHSUMMIT* SERVICE

redhat.

```
# ausearch -m AVC -ts recent

<no matches>
```


redhat.

SELINUX POLICY FOR THE *RHSUMMIT* SERVICE IS READY!

redhat.

Troubleshooting Existing Policy

redhat.

MISLABELED SYSTEM I.E LABELS ON OBJECTS ARE WRONG

redhat.

```
# restorecon -Rv /

or

# fixfiles onboot
# reboot
```


redhat.

SELINUX MODIFICATIONS OF THE DISTRO POLICY VIA SELINUX USERSPACE TOOLING

redhat.

SOMETIMES IT'S NOT NECESSARY TO CREATE CUSTOM SELINUX POLICY, LOCAL MODIFICATION CAN FIX IT.

redhat.

APACHE HTTP SERVER WITH CHANGES IN THE DEFAULT CONFIGURATION

redhat.

httpd service configured to listen on tcp port 3131 instead of port 80

redhat.

httpd service configured to listen on tcp port 3131 instead of port 80

document root will be /var/test_www/ instead of /var/www/

redhat.

Change in /etc/httpd/conf/httpd.conf

```
Listen 80 -> Listen 3131

DocumentRoot "/var/www/html" => DocumentRoot "/var/test_www/html"

<Directory "/var/www/html">  => <Directory "/var/test_www/html">
```


redhat.

Change in /etc/httpd/conf/httpd.conf

```
# sed -i 's_Listen 80_Listen 3131_' /etc/httpd/conf/httpd.conf

# sed -i 's_DocumentRoot "/var/www/html"_DocumentRoot
"/var/test\_www/html"_' /etc/httpd/conf/httpd.conf

# sed -i 's_<Directory "/var/www/html">_<Directory
"/var/test\_www/html">_' /etc/httpd/conf/httpd.conf
```


redhat.

APPLY CHANGES IN THE CONFIGURATION AND SEARCH FOR AVC DENIALS

redhat.

```
# systemctl restart httpd

# ausearch -m AVC -ts recent

....
```


redhat.

httpd service trying to bind on port 3131 instead of 80, this should be changed in SELinux policy

```
type=AVC msg=audit(1491948261.488:599): avc:  denied  { name_bind } for
 pid=5920 comm="httpd" src=3131 scontext=system_u:system_r:httpd_t:s0
  tcontext=system_u:object_r:unreserved_port_t:s0 tclass=tcp_socket
```


redhat.

```
# sesearch -A -s httpd_t -t http_port_t  -c tcp_socket  -p name_bind
Found 1 semantic av rules:
  allow httpd_t http_port_t : tcp_socket name_bind ;

# semanage port -a -t http_port_t -p tcp 3131
# semanage port -l | grep http_port_t
http_port_t                    tcp    3131, 80, 81, 443, 488, 8008,
8009, 8443, 9000
```


redhat.

httpd service DocumentRoot is in /var/test_www/html and this directory has wrong label

```
type=AVC msg=audit(1491949594.146:622): avc:  denied  { read }
   for  pid=6094 comm="httpd" name="index.html" dev="dm-0"
     ino=13485999 scontext=system_u:system_r:httpd_t:s0
     tcontext=unconfined_u:object_r:var_t:s0 tclass=file
```


redhat.

```
# matchpathcon /var/test_www/html/index/html
/var/test_www/html/index/html    system_u:object_r:var_t:s0

# matchpathcon /var/www/html/index.html
/var/www/html/index.html    system_u:object_r:httpd_sys_content_t:s0

# semanage fcontext -a -t httpd_sys_content_t "/var/test_www(/.*)?"
# restorecon -Rv /var/
```


redhat.

```
# semanage fcontext -l | grep httpd_sys_content_t | grep www
/var/www(/.*)?              all files
system_u:object_r:httpd_sys_content_t:s0
/var/test_www(/.*)?    all files
system_u:object_r:httpd_sys_content_t:s0
....
....
```


redhat.

HTTPD SERVICE SELINUX DENIALS ARE FIXED WITHOUT WRITING CUSTOM POLICY!

redhat.

LOCAL POLICY MODULE WRITTEN IN CIL

redhat.

When system interface cannot be used (e.g: there is no such interface), it's possible to create local module in CIL language.

For example we have following AVC:

```
type=AVC msg=audit(1491942725.855:436): avc:  denied  { read }
 for  pid=4874 comm="rhsummit" name="resolv.conf" dev="dm-0"
    ino=5051289 scontext=system_u:system_r:rhsummit_t:s0
    tcontext=system_u:object_r:net_conf_t:s0 tclass=file
```


redhat.

```
# cat module.cil
 (allow source_context target_type (class (permissions)))

 # cat local_module.cil
 (allow rhsummit_t net_conf_t (file (read open getattr)))

 # semodule -i local_module.cil
 # semodule -lfull | grep local_module
 400 local_module   cil
```


redhat.

QUESTIONS?

Miroslav Grepl's blog	https://mgrepl.wordpress.com/
Paul Moore's blog	http://www.paul-moore.com/
Lukas Vrabec's blog	https://lvrabec-selinux.rhcloud.com/
Dan Walsh's blog	http://danwalsh.livejournal.com/

redhat.

THANK YOU


plus.google.com/+RedHat


facebook.com/redhatinc


linkedin.com/company/red-hat


twitter.com/RedHatNews


youtube.com/user/RedHatVideos

LEARN. NETWORK.
EXPERIENCE
OPEN SOURCE.

redhat.