SELinux news in Fedora 16

Miroslav Grepl

ABSTRACT

SELinux overview

File name transitions

Pre-built policy

Shrinking policy

Permissivedomains module

WHAT IS SELINUX

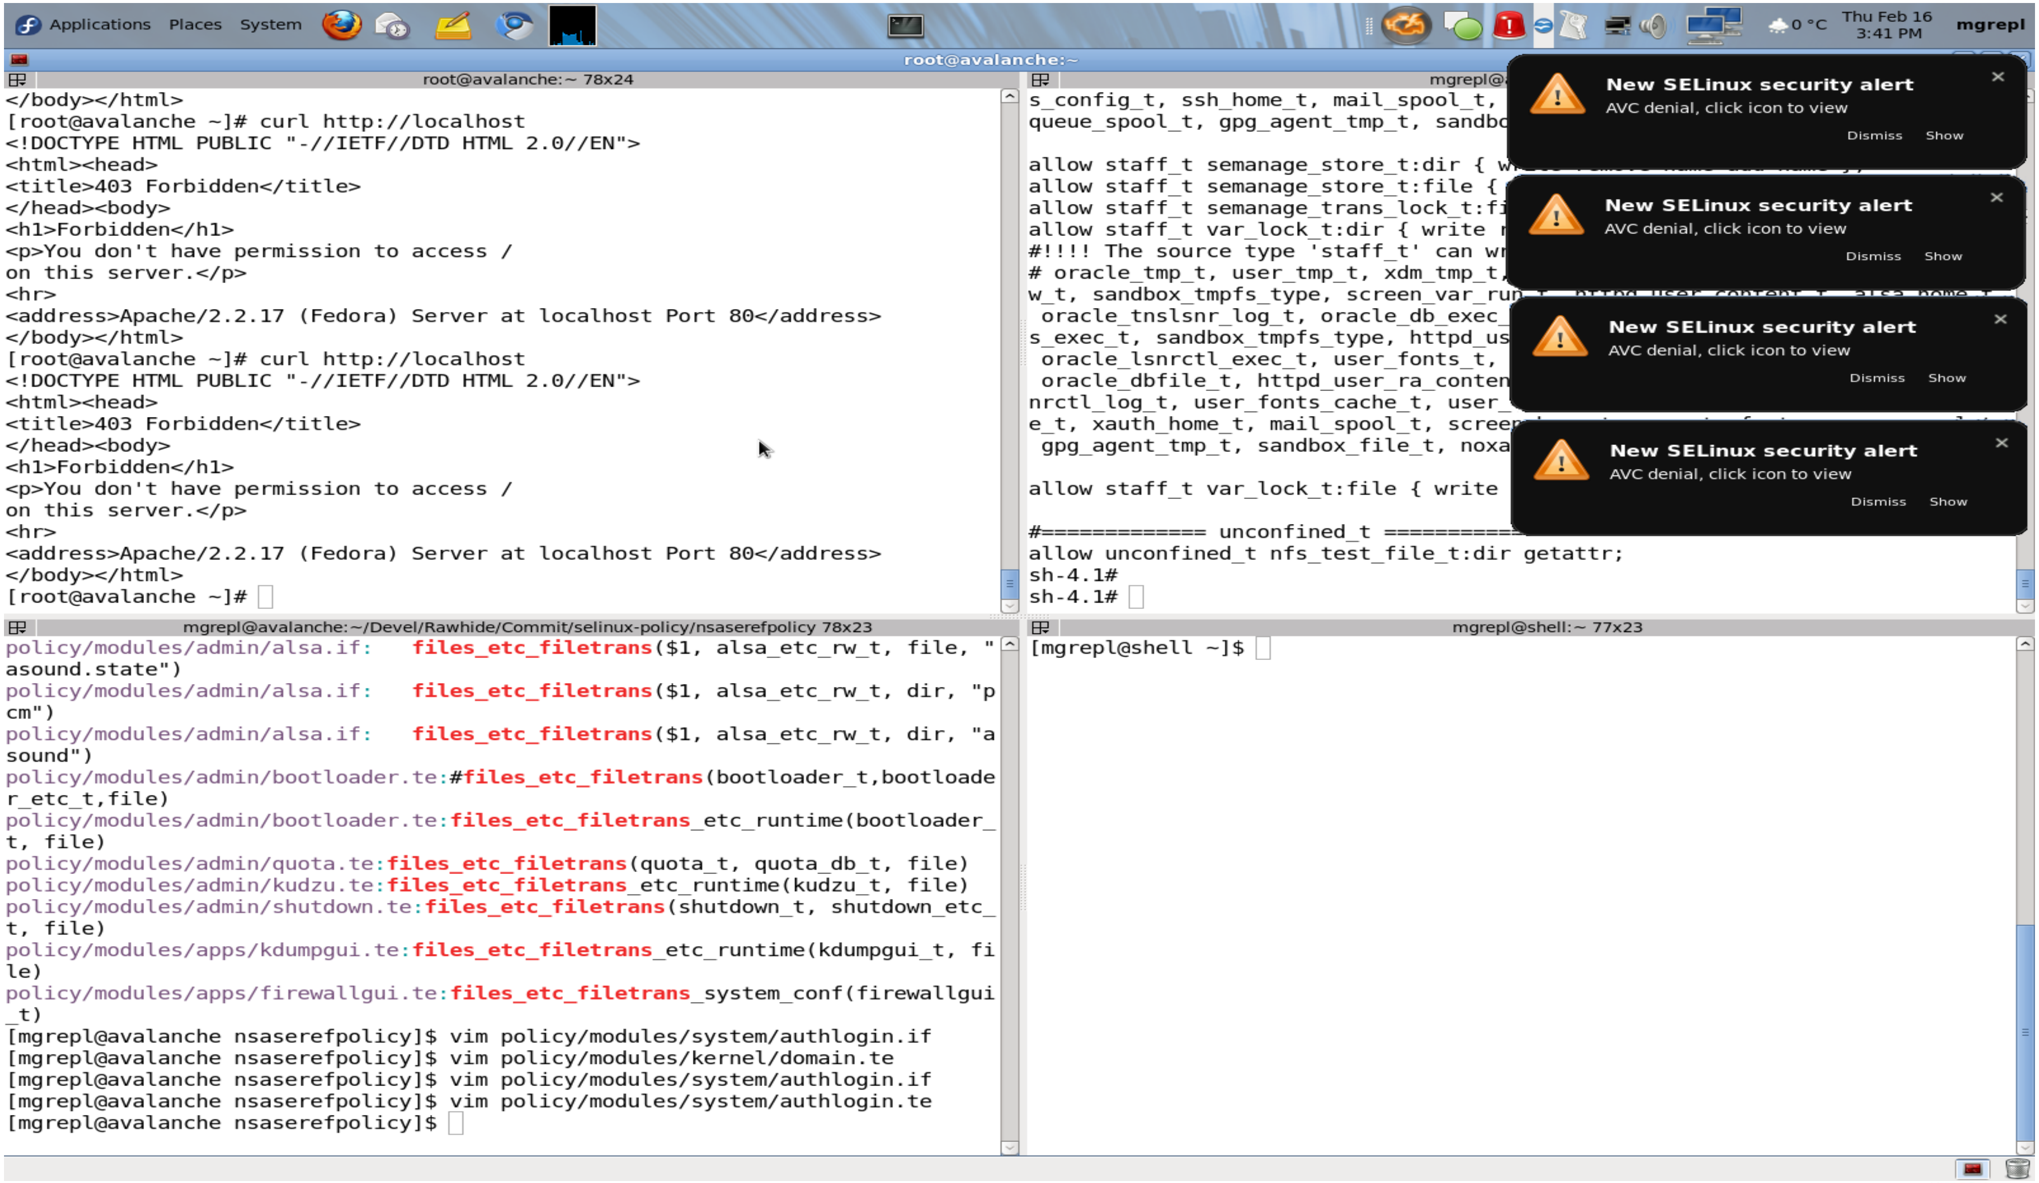

```
</body></html>
[root@avalanche ~]# curl http://localhost
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>403 Forbidden</title>
</head><body>
<h1>Forbidden</h1>
<p>You don't have permission to access /
on this server.</p>
<hr>
<address>Apache/2.2.17 (Fedora) Server at localhost Port 80</address>
</body></html>
[root@avalanche ~]# curl http://localhost
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>403 Forbidden</title>
</head><body>
<h1>Forbidden</h1>
<p>You don't have permission to access /
on this server.</p>
<hr>
<address>Apache/2.2.17 (Fedora) Server at localhost Port 80</address>
</body></html>
[root@avalanche ~]#
```

```
s_config_t, ssh_home_t, mail_spool_t,
queue_spool_t, gpg_agent_tmp_t, sandbo

allow staff_t semanage_store_t:dir { w
allow staff_t semanage_store_t:file {
allow staff_t semanage_trans_lock_t:fi
allow staff_t var_lock_t:dir { write r
#!!!! The source type 'staff_t' can wr
# oracle_tmp_t, user_tmp_t, xdm_tmp_t,
w_t, sandbox_tmpfs_type, screen_var_run
 oracle_tnslsnr_log_t, oracle_db_exec_
s_exec_t, sandbox_tmpfs_type, httpd_us
 oracle_lsnrctl_exec_t, user_fonts_t,
 oracle_dbfile_t, httpd_user_ra_conten
nrctl_log_t, user_fonts_cache_t, user_
e_t, xauth_home_t, mail_spool_t, screen
 gpg_agent_tmp_t, sandbox_file_t, noxa

allow staff_t var_lock_t:file { write

#============= unconfined_t ==========
allow unconfined_t nfs_test_file_t:dir getattr;
sh-4.1#
sh-4.1#
```

New SELinux security alert
AVC denial, click icon to view

New SELinux security alert
AVC denial, click icon to view

New SELinux security alert
AVC denial, click icon to view

New SELinux security alert
AVC denial, click icon to view

```
policy/modules/admin/alsa.if:    files_etc_filetrans($1, alsa_etc_rw_t, file, "
asound.state")
policy/modules/admin/alsa.if:    files_etc_filetrans($1, alsa_etc_rw_t, dir, "p
cm")
policy/modules/admin/alsa.if:    files_etc_filetrans($1, alsa_etc_rw_t, dir, "a
sound")
policy/modules/admin/bootloader.te:#files_etc_filetrans(bootloader_t,bootloade
r_etc_t,file)
policy/modules/admin/bootloader.te:files_etc_filetrans_etc_runtime(bootloader_
t, file)
policy/modules/admin/quota.te:files_etc_filetrans(quota_t, quota_db_t, file)
policy/modules/admin/kudzu.te:files_etc_filetrans_etc_runtime(kudzu_t, file)
policy/modules/admin/shutdown.te:files_etc_filetrans(shutdown_t, shutdown_etc_
t, file)
policy/modules/apps/kdumpgui.te:files_etc_filetrans_etc_runtime(kdumpgui_t, fi
le)
policy/modules/apps/firewallgui.te:files_etc_filetrans_system_conf(firewallgui
_t)
[mgrepl@avalanche nsaserefpolicy]$ vim policy/modules/system/authlogin.if
[mgrepl@avalanche nsaserefpolicy]$ vim policy/modules/kernel/domain.te
[mgrepl@avalanche nsaserefpolicy]$ vim policy/modules/system/authlogin.if
[mgrepl@avalanche nsaserefpolicy]$ vim policy/modules/system/authlogin.te
[mgrepl@avalanche nsaserefpolicy]$
```

```
[mgrepl@shell ~]$
```

WHAT IS SELINUX

SELinux knows if you do bad things

WHAT IS SELINUX

- Implementation of Mandatory Access Control (MAC)
  - which subject can access which object
  - subjects (processes, users) and objects (files, devices) have

security context == ***label***

system_u:system_r:**abrt_t**:s0

  - SELinux makes decisions based on these labels

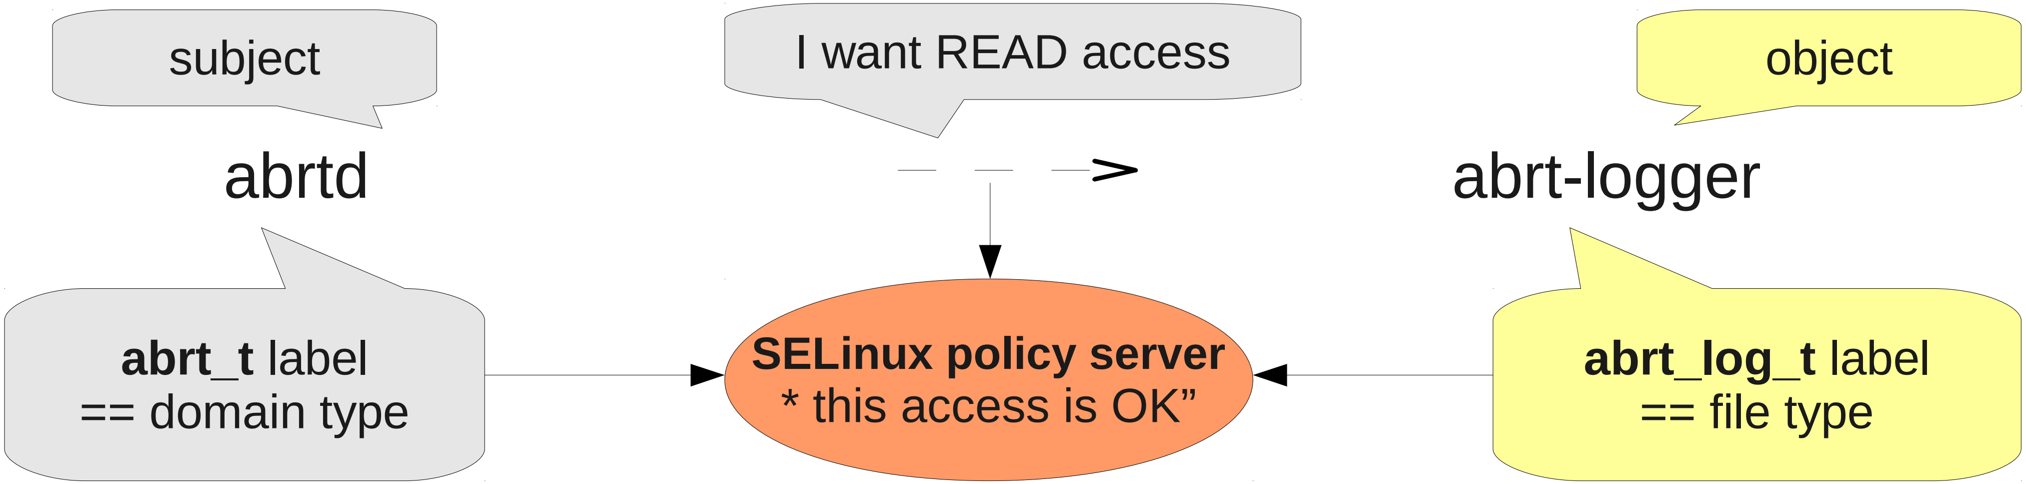

DON'T TURN OFF SELINUX

- your /etc/selinux/config **should not contain**

  SELINUX=disabled

- rather please use
  - **PERMISSIVE MODE –** you can do anything but SELinux reports Access Vector Cache (AVC) messages

SELINUX = permissive	/etc/selinux/config
enforcing = 0	as a kernel parameter
setenforce 0	on the command line

  - **PERMISSIVE DOMAINS –** SELinux allows a domain to do anything but reports AVC's

    ```
    semanage permissive -a DOMAIN
    ```

New features in Fedora 16

FILE NAME TRANSITIONS

- labeling files is now easier for users/administrators
- accidental mislabeling of file objects is now sanitized

Previously

```
$ mkdir /root/.ssh
$ ls -dZ /root/.ssh
system_u:object_r:admin_home_t:s0
$ matchpathcon /root/.ssh
/root/.ssh   system_u:object_r:ssh_home_t:s0
```

Now

```
$ mkdir /root/.ssh
$ ls -dZ /root/.ssh
system_u:object_r:ssh_home_t:s0
```

FILE NAME TRANSITIONS

- we can write a policy rule that states

  *“If the unconfined_t user process creates the ".ssh" directory in a directory labeled admin_home_t, then it will get created with the ssh_home_t label. *”*

  - example of a rule

```
filetrans_pattern(unconfined_t, admin_home_t, ssh_home_t, dir, ".ssh")

filetrans_pattern(unconfined_t, etc_t, passwd_file_t, file, "group")
```

- reduce many errors => **BIG STEP FORWARD**

```
$ sesearch  -T -c file | grep \" | wc -l
1384
```

PRE-BUILT POLICY

Previously

- SELinux policy has been always re-built in the post install

=> more time, more memory

Now

- selinux-policy-TYPE packages are shipped with a pre-built policy
- installation selinux-policy packages is faster

SHRINKING POLICY

- systemd output in Fedora 16 devel phase
  - part of boot message on boot

“I also added some basic profiling output for SELinux which unfortunately shows that SELinux costs around 5s on every boot on f16 (and that on my really fast machine!). Sad.”

- everyone could know how much time SELinux was costing them on boot
- the policy contained over 300 thousands rules => where did come from?

SHRINKING POLICY

- policy language uses attributes to reduce rules
  - attributes can cover more types

    port_type attribute => for all defined ports

    reserved_port_type attribute => for all defined reserved ports
  - we can define a single rule rather than many

```
allow domain_t dhcpc_port_t:tcp_socket name_bind
allow domain_t dns_port_t:tcp_socket name_bind
....
```

vs

```
allow domain_t reserved_port_type:tcp_socket name_bind
```

```
$ seinfo -axreserved_port_type
58
```

SHRINKING POLICY

- we can define a rule like

```
allow ssh_t { port_type -reserved_port_type }:tcp_socket name_bind
```

  - we ended with a rule for each type

```
allow ssh_t amqp_port_t:tcp_socket name_bind;
allow ssh_t asterisk_port_t:tcp_socket name_bind;
...
```

=> ***100's** of allow rules*

  - we changed the rule

```
allow ssh_t unreserved_port_type:tcp_socket name_bind;
```

=> *only **1** rule*

SHRINKING POLICY

Previously

- on a Fedora 15

```
$ seinfo
Allow:          282 444
Dontaudit:      184 516
```

Now

- on Fedora 16

```
$ seinfo
Allow:          88 242
Dontaudit:      11 302
```

=> tools use load policy run about 3 times as fast

=> policy takes less kernel memory

PERMISSIVEDOMAINS MODULE

Previously

- permissive flag was in individual policy modules

```
$ cat /etc/fedora-release
Fedora release 15 (Lovelock)

$ grep permissive PATHTO/abrt.te
permissive abrt_dump_oops_t;
permissive abrt_retrace_worker_exec_t;
permissive abrt_retrace_coredump_t;
```

=> permissive statement was permanent

PERMISSIVEDOMAINS MODULE

Now

- flags have been moved to a new policy module called *permissivedomains.pp*
- you can disable all permissive domains from the system

```
$ semanage permissive -l | wc -l
62
$ semodule -d permissivedomains.pp
```

PERMISSIVEDOMAINS MODULE

- **permissive domain** – can do everything and AVC messages are logged
- **unconfined domain** – can do everything but AVC messages are not logged

```
$ seinfo -xaunconfined_domain_type
unconfined_domain_type
  rpm_t
  anaconda_t
  rpm_script_t
  ...
```

- stricter policy in one step + confined users

```
$ semodule -d unconfined.pp permissivedomains.pp
```

REFERENCES

- http://danwalsh.livejournal.com
    - http://danwalsh.livejournal.com/43264.html
    - http://danwalsh.livejournal.com/46018.html
    - http://danwalsh.livejournal.com/46245.html
    - http://danwalsh.livejournal.com/46388.html
- http://mgrepl.wordpress.com/

QUESTIONS?