

Information and Organisational Security

Guides for Practical Classes

João Paulo Barraca and Vitor Cunha

Department of Electronics, Telecommunications and Informatics
University of Aveiro

2018–2019

Contents

10 Security Mechanisms of the Linux OS	10-1
10.1 Introduction	10-1
10.2 System Login	10-1
10.3 Identity	10-1
10.4 User <code>root</code>	10-2
10.4.1 <code>sudo</code> command	10-2
10.4.2 Set-UID mechanism	10-3
10.5 File Protection Mechanisms	10-4
10.5.1 Read protection	10-4
10.6 Confinement	10-5
10.6.1 <code>chroot</code>	10-7
10.6.2 Linux Security Modules: Apparmor	10-8
10.7 Bibliography	10-10

10

Security Mechanisms of the Linux OS

10.1 Introduction

This guide aims to study the security mechanisms of the Linux Operating System. In particular the methods for elevation or restriction of privileges. It should be executed in the Linux OS, preferably in the Virtual Machine used in this course.

10.2 System Login

To access a system it is required to start the process of login. You can repeat this process by using the key combination **CTRL-ALT-F1**, or executing the command `login`, by the super user.

As the result of the login process of a specific user, a session is started using the command interpreter (*shell*) associated with the user account.

10.3 Identity

It is possible to obtain information about the identity of the current user, as well as the groups to which he belongs. One of the possible commands is the following:

```
id
```

It is also possible to obtain lists of available users and groups, through the `getent` command. The following command lists the possible entries to be enumerated:

```
man getent
```

while the following command obtains all users:

```
getent passwd
```

The output of this command is obtained from several sources, depending on the system configuration. In this case, for users and groups, it should consist of `/etc/passwd` and `/etc/group`.

10.4 User root

The `root` is the Linux *super-user*. With the exception of Mandatory Access Control mechanisms, everything is always allowed. It is omnipotent and has the ID with value of 0.

10.4.1 sudo command

The syntax `sudo <command>` allows executing the command specified as argument, using the identity of the `root` user. In order for this action to be authorized, the user that executes this command needs to provide his credentials.

If you execute the following command, you will validate the user that is used for executing the `id` command:

```
sudo id
```

If you execute the command repeatedly, you will not require to introduce the credentials. Moreover, only a restricted set of users may effectively be authorized. The authorization is controlled by the content of the file (`/etc/sudoers`), which lists the users or groups that are allowed to elevate their privileges through this method.

Check the content using:

```
cat /etc/sudoers
```

This method is consistent with the general execution of actions, split in at least two parts: authentication (validate the credentials) and authorization (verification of the rules in the `/etc/sudoers` file).

10.4.2 Set-UID mechanism

Some commands should always be executed by a specific user, in particular the super-user. Examples of such commands are the `passwd`, `ping`, `traceroute`, or `chfn` commands. All manipulate files restricted to the super-user, or invoke actions limited to the super-user. Without a way of elevating the rights for specific actions it would be impossible to execute trivial and essential tasks such as validating connectivity or changing the user password.

The `sudo` program is also an example of a program that always requires the elevation of the privileges, so that one user can execute commands as another user, or as the super-user.

Therefore, this is a dilemma: users should be restricted to their privileges, but should also be able to execute action as another users.

The way of dealing with this is the Set-UID mechanism: a bit in the binary file storing a program, specifies that the process should be **always** started using an alternative user, instead of the user that is executing the process.

Analyze the permission of the `passwd` and `sudo`:

```
ls -la /usr/bin/passwd
ls -la /usr/bin/sudo
```

Then compare the permissions of these files with the ones of the `ls` command (which is not privileged):

```
ls -la /bin/ls
```

To see the differences of setting this flag, copy the `/usr/bin/id` file to the home directory of your unprivileged user, execute it, and register the output:

```
cd ~
cp /usr/bin/id .
./id
```

Now, set the Set-UID bit of the local file, and execute it again. Compare the results obtained:

```
sudo chmod u+s ./id
./id
```

You should notice the addition of a field named `euid`, the short of *Effective User ID*, and you should be able to see that the UID is different from the EUID. In the case of the Linux OS, the EUID is the value that effectively sets the user, and by consequence the permissions of the process.

10.5 File Protection Mechanisms

The files and folders in the filesystem have a basic protection allowing to indicate if they can be read, modified, or executed (**rx**) by their owner, their primary group or other users.

The `ls -la` command allows to inspect the permissions set to any file in the filesystem.

You can also create a folder and file, whose permissions can be inspected:

```
mkdir test_dir
cd test_dir
date > test_file
```

where `verb_dir` should be the name of the folder to create, and `test_file` the name of the file to create. In this case the file is created using the current date as its content.

10.5.1 Read protection

Check the permissions applied to the file `test_file` with the following command:

```
ls -la test_file
```

Check who is the file owner, the main group and the protections applied to this specific entities, as well as for all other users.

You can show the content of the file using the `cat` command:

```
cat test_file
```

Then, remove the permission for the owner to read the file:

```
chmod u-r test_file
```

Try to show the content of the file, which should not be possible. You can use the syntax `u+r` to enable the owner to read the file.

Correct the permissions and check if you can access the file content (you should be able to do it)

You can manipulate the remaining permissions in order to protect the content of the file, and limit the actions as adequate.

The permissions of the `test_dir` can also be manipulated. For this purpose, consider the operations of creating a new file, changing the content of an

existing file, reading a file and listing the contents.

Remove the write permissions for the folder:

```
chmod u-w test_dir
```

and check the impact for the before mentioned actions. Do the same for the remaining permissions (`read` and `execute`).

Can you explain the results obtained with these changes? In particular the restriction of the `execute` permission will also impact the visibility over the directory.

10.6 Confinement

Besides the permissions of folders and files, as well as groups a user belongs to, it is frequently required (or at least useful) to restrict a program to a subset of the resources and activities available. This is important to implement the basic security principle least privilege: services should only be provided the means required for them to execute their tasks.

If a process is started from the command line, without any other mechanism, this will not be trivial to enforce.

For the purpose of this section, consider the following program, which is also available in the course website:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <string.h>

#define DATADIR "/data"

int main(int argc, char** argv) {
    struct sockaddr_in saddr;
    int sockfd = socket(AF_INET, SOCK_DGRAM, 0);

    bzero((char *)& saddr, sizeof(saddr));
    saddr.sin_family = AF_INET;
    saddr.sin_addr.s_addr = htonl(INADDR_ANY);
```

```

saddr.sin_port = htons((unsigned short) 1234);

bind(sockfd, (struct sockaddr *)&saddr, sizeof(saddr));

while (1) {
    struct sockaddr_in caddr;
    socklen_t clen = sizeof(caddr);
    int n;
    char buf[1024];
    char fname[2048];
    bzero(fname, 2048);
    bzero(buf, 1024);

    int len = recvfrom(sockfd, buf, 1024, 0, \
                      (struct sockaddr *)&caddr, &clen);
    if (len <= 0)
        continue;

    sprintf(fname, "%s/%s", DATADIR, buf);
    FILE* fd = fopen(fname, "r");
    fprintf(stderr, "Serving file %s...", fname);
    if (fd == NULL) {
        fprintf(stderr, "Error\n");
        continue;
    }

    while (len > 0) {
        char data[1400];
        len = fread(data, 1, 1400, fd);
        if (len > 0)
            sendto(sockfd, data, len, 0, \
                  (struct sockaddr *)&caddr, clen);
    }
    fclose(fd);
}
}

```

It contains a very simple UDP server that outputs the content of any file requested in the payload of a UDP message. It should provide files that are contained in the DATADIR, but it has a vulnerability as it allows for directory transversal. Can you find the vulnerability?

Assuming the name is `server.c`, it can be compiled issuing the following command:

```
gcc -o server server.c
```

The server can be executed issuing:

```
./server
```

and **in another terminal**, you can send a command to the server by executing:

```
echo -n "t.txt" | nc -u -w 5 127.0.0.1 1234
```

The result should be the content of the `/data/t.txt` file. If this file doesn't exist, you can create it and test. The `nc` program can be stopped by pressing **CTRL-C**, or by waiting 5 seconds (due to the arguments `-w 5`).

You can request other files to test the program. An attacker would exploit this server by requesting files such as `../etc/shadow` (note the `..`), exploiting the directory transversal vulnerability.

Note: The following sub-sections can be executed with any executable service that uses network capabilities. The purpose of using such a simple program is due to the fact that it has a small amount of dependencies from configuration files, data files or other libraries, simplifying the execution of the confinement exercises.

10.6.1 chroot

It is frequent that services do not need to operate over all files of the filesystem, but only over a subset of the files. In the case of our test program, we would like to confine it to the directories required for its operation, plus the `DATADIR`.

The `chroot` command allows to confine the vision that a process has over the filesystem. It works by changing the notion the application has of its root folder. Therefore, a process will see an hierarchy starting in the root directory (as it always would), but this root directory can be manipulated, so that the process is confined to a sub-folder.

Creating the `|chroot|` implies constructing an environment with all files required for the execution of the program, which also includes the libraries required to load it from a binary file.

In order to create the environment for our server should start by creating a folder for the `chroot` environment:

```
sudo mkdir -p /opt/chroot
sudo mkdir -p /opt/chroot/bin
sudo mkdir -p /opt/chroot/lib
sudo mkdir -p /opt/chroot/lib64
sudo mkdir -p /opt/chroot/data
```

The `server` binary should be placed in the `/opt/chroot/bin` folder, and all the libraries it requires should be placed in the `/opt/chroot/lib` folder.

We can enumerate the libraries required using the `ldd` command:

```
ldd server
```

which should result in the following output:

```
linux-vdso.so.1 (0x00007ffc32f2f000)
libc.so.6 => /lib/x86_64-linux-gnu/libc.so.6 (0x00007f69b01d9000)
/lib64/ld-linux-x86-64.so.2 (0x00007f69b07cd000)
```

Therefore, copy this libraries to the appropriate sub-folder inside the `/opt/chroot` folder (e.g, `/opt/chroot/lib/x86_64-linux-gnu/libc.so.6`).

Afterwards, place some small text files in the `/opt/chroot/data` folder, as these files will be served by our server.

Finally, the server can be started in the `chroot` environment by issuing:

```
chroot /opt/chroot /bin/server
```

You can use `nc` to obtain the files in the `DATADIR`, and the problem will still be vulnerable to the previous vulnerability. However, the impact is reduced to the files inside the `chroot` environment.

10.6.2 Linux Security Modules: Apparmor

The Linux Security Modules allows the configuration of Mandatory Access Control, which are applied to processes, independently of the user that is executing the process. One implementation of these modules is `Apparmor`.

When considering our test server, its access is limited to the user permissions, and if executed with a restricted user, the impact of its vulnerability will be reduced to the files accessible by the user. As an example, an attacker would be able to access personal files of that user, but not of another user, or some system files.

What is appropriate is to restrict the service to network communications, and access to the DATADIR.

This can be achieved using Apparmor, and an appropriate profile for our server.

The first step is to generate a profile stating what is allowed or denied. The profile is in `/etc/apparmor.d` and assumes the server is in a specific location.

In our case, place the server in `/opt` and create a file name `/etc/apparmor.d/opt.server` with the following content:

```
/opt/server {  
  
    /lib/x86_64-linux-gnu/ld-*.so mr,  
    /lib/x86_64-linux-gnu/libc*.so mr,  
  
    /opt/server mr,  
  
    network inet udp,  
    /data/** r,  
}
```

This profile allows map and read access to the binary file as well as the libraries it requests, network access using UDP sockets. It also adds the authorization for reading the contents of the `/data` folder.

Activate the profile by issuing:

```
apparmor_parser -r /etc/apparmor.d/opt.server
```

Start the server and request one file that exists in the DATADIR. Then, request `../etc/passwd` and analyze what happened. You can have more information in `dmesg`.

The use of `apparmor` is very rich, and you can create audit records for the activities executed by each process. As an example, you can replace the last line of the profile with:

```
audit /data/** r,
```

This will specify that the actions towards those files should be audited, resulting in a log entry. After changing this line and activating the new profile, you can use `dmesg` to analyze the entries that are added when a file is transferred.

10.7 Bibliography

- User identifier, http://en.wikipedia.org/wiki/User_ID
- Group identifier, http://en.wikipedia.org/wiki/Group_identifier
- Super-user, <http://en.wikipedia.org/wiki/Superuser>
- Filesystem permissions, http://en.wikipedia.org/wiki/File_system_permissions
- Set-UID mechanism, <http://en.wikipedia.org/wiki/Setuid>
- sudo command , <http://en.wikipedia.org/wiki/Sudo>
- chroot mechanism, <http://en.wikipedia.org/wiki/Chroot>
- apparmor mechanism, <https://gitlab.com/apparmor>