

# System Security

Alessandro Barenghi

Dipartimento di Elettronica e Informazione  
Politecnico di Milano

*barenghi - at - elet.polimi.it*

April 3, 2012

# Lesson contents

## Overview

- The \*nix user control and file permission model
- Linux Capabilites based permission control
- Secure programming practices

# How to train your users

## Overview

- Linux<sup>a</sup> are natively conceived as a multi-user operating system
- Multiple users will naturally perform different task, and have different needs on the machine
- To avoid abuses, the user access to the hardware should be somehow restricted
- It should also be possible to unite users in groups, in order to ease the control

---

<sup>a</sup>and more in general all the Unices

# Access to the system

## Login and credential storage

- The basic Linux authentication and user control method relies on assigning a unique numeric User Identifier to every user
- The UIDs, together with some information about the users are stored in `/etc/passwd`
- For safety reasons, the actual hashes of the user passwords<sup>a</sup> are preserved in `/etc/shadow`
- At login, the system checks if the user exists, hashes the passwords and matches it against the correct hash and executes the preferred command interpreter for the user
- For large system it is possible to use a database backend to store the credentials together with a centralized login

---

<sup>a</sup>as noone, in his sane mind should store cleartext passwords

# Processes

## Overview

- Every process in modern Unices has three UID identifying it :
  - real UID: it is the UID of the actual process owner
  - effective UID: it is the UID used used for access control
  - saved UID: stores a previous UID in order to be able to change it back
- Every process in Linux also has an FSUID, which acts as a separate UID for file access
- Usually the FSUID follows the EUID, unless explicitly set
- Analogously to UIDs, 4 group IDs are attached to each process

# File permissions

## Overview

- As we recall from the basic system administration lessons, under Unix everything is a file
- It is thus natural to have the access permissions set on every file
- Every file has both an owner and a group to which it belongs
- The basic permissions to act on a file are stored for both the owner and the group

# File permissions

## the User-Group-Other model

- Three access check are made under Unix : the permission to read (r), write (w) and execute a file x
- The common permission are usually represented visually by 9 characters `rw-rw-rw-`
- The first 3 (`rw-rw-rw-`) are the permissions for the owner
- The second 3 (`rw-rw-rw-`) are the permissions for the members of the same group which owns the file
- The last 3 (`rw-rw-rw-`) are the permissions for the others
- Permissions can be changed via the `chmod` command

# File permissions

## Changing permissions and ownerships

- File Permissions can be changed via the `chmod` command
- `chmod` accepts either a 4-digit octal representation of the permissions...
- or a list of the permissions to be added-revoked
- Revoking the execute permission on a directory forbids its traversal
- A special permission, the *sticky bit* allows only the owner of a file to remove it (even if others have full access)
- The `chown` command allows to change the owner/group of a file



# The administrator

## Got root?

- Among all the users, one is special: the system administrator
- The system administrator, by default named **root** has a reserved UID of 0
- *Any* permission check will be skipped if the issuer of an action has UID 0
- This implies that any process with EUID 0 can read/write/run anything
- Having an EUID of 0 allows a process to modify any access credentials without knowing the previous ones

# Saved IDs

## Set user ID and set group ID

- Executable programs can be endowed with the set user ID (suid) bit via `chmod`
- These programs are executed with an EUID equal to the owner of the file and not the caller
- This allows to execute programs as someone else, even without his/her credentials
- Analogously, the set group ID permission changes the EGID value
- Suid root binaries act exactly as if the system administrator was running them
- You can find all of them via `find / -perm +6000 -type f -user root` (suid and sgid are expressed as an octal digit)

# A programmer point of view

## Linux API for UNIX permissions

- The `getuid` and `geteuid` function return respectively the real and effective UID
- The `setuid` and `seteuid` try to set the respective UIDs if permitted
- The `chmod` function acts exactly as the command
- For `chmod` to work, either the EUID matches the owner of the file, or the program runs as `root`
- The same concept applies to the `chown` function

# Linux Capabilities

## When root is too much

- The power held by the administrator in UNIX is absolute, yielding it in a single block may be too much
- Typical example : why must my web server run as root when it just needs to open a privileged port?
- Solution : POSIX Capabilities : partition the administrative rights in a set of capabilities
- Affix some of the capabilities to an executable file, instead of setting it suid root

# Linux Capabilities

## Common capabilities

- The list of all system capabilities is available with `man capability`
- Useful capabilities are :
  - `CAP_NET_RAW` : allows to use raw sockets
  - `CAP_NET_ADMIN` : allows to change routing tables
  - `CAP_KILL` : unlocks signal sending to everyone
  - `CAP_SYS_NICE` : allows to renice with negative values
- You can retrieve the list of capabilities of a file with `getcap <filename>`
- Symmetrically, you can set the capabilities `setcap <capability>[+|-|=]ep <filename>`

# Key Ideas

## Do not take anything for granted

- Secure programming implies checking that the inputs and the environment of a program are safe
- Error messages should be explanatory, but not too much
- Memory management errors are a possible source of endless issues
- The `printf` case
- Capabilities or extra privileges granted should be discarded when not needed

# Hostile Environments

## Mind user input

- User input may not just be “rotten”, but actually intentionally poisoned
- It is possible to deviate the common behaviour of a program if it does not take any care in verifying that the input matches a sane format
- We will see two examples :
  - the `system` call which runs a command in a new instance of the command interpreter
  - the `exec` call family, which runs a command directly, replacing the code of the running process
- Remember that, when a call to `fork` or `exec` is issued, the value of EUID and EGID is **preserved**

# Hostile Environments

## Error notifications

- Our good nature pushes us towards using meaningful error messages to help the user
- However, the error messages have historically been a source of information leakage
- Every time an error message is displayed, some information about the state of the program leaks
- Sample issue: what if I help the user outputting the wrong line of a config file?
- Best practice : employ verbose error messages only enclosed in `DEBUG` macros and remove them afterwards



# Hostile Environments

## What's in a SEGFALT? Wouldn't it be as dangerous as...

- Bad memory management and boundary checking looks as just a safety issue
- It turns out that performing memory copy operations without checking boundaries may allow the attacker to write anywhere in the memory segment of the program
- Since the stack is a commonly writeable area, it may happen that a part of the stack is overwritten during a `strcpy` operation
- ... but on the stack there's the **return pointer** of the actual call!

# Hostile Environments

## Proper memory management

- Arbitrary memory rewriting allows control flow hijacking: it's feasible and it's being done since 1996
- Always allocate enough space when you are copying strings (`#char+1`)
- Use only boundary checking copy functions (`strncpy`)
- Always check if memory allocations succeed, copying into a NULL pointer is also dangerous<sup>a</sup>

---

<sup>a</sup>NULL is defined as `(void*) 0` , so if someone mmaps things to 0...

# Hostile Environments

## Format string issues

- At a first glance, the `printf` function looks *mostly* harmless
- However, the number of arguments it treats is implicitly indicated in the format string
- What if a `printf` does not have any parameters, and the user can specify a part of the string being printed?
- Reminder: among the format options allowed by C there is `%n`, which writes the number of printed character in a variable...

# Dropping privileges

## Typical suid privilege dropping

The correct practice if you have a suid root program to be run is to drop the privileges as soon as you do not need them anymore :

State	Action	EUID	Real UID	Set UID
Startup		0	user	0
temp. drop	seteuid(getuid())	user	user	0
restore	seteuid(0)	0	user	0
Perm. drop	setuid(getuid())	user	user	user
restore	seteuid(0)	user	user	user

# Dropping privileges

## Typical cap privilege dropping

- Active capabilities can be checked via the `cap_get_flag` function (returning true or false)
- Capabilities can be dropped easily via the `cap_set_flag` function
- Capability dropping is permanent for the current run of the program (no mechanism as saved caps)
- the whole capability set can be retrieved via a call to the `cap_get_proc` function