

SECURITY

- Why bother at all?
 - The Internet is not a secure place. Many people try to crack systems, and the network infrastructure is inherently insecure.
 - Any computer (and toasters too!) are nowadays connected to the Internet.
- How secure is secure?
 - No computer can be ever totally secure.
 - Security needs vary from case to case (e.g., your home computer vs. your bank's).
 - The more secure your system is, the more intrusive security becomes. You need to decide when your system will still be usable, and yet secure for your purposes.

UNPRIVILEGED SERVERS

- Vulnerabilities are greatly minimized if your server runs unprivileged.
 - A program inherits not only the open descriptors, but the user it belongs to.
 - However, a server is usually launched by the init script, which is run for obvious reasons as `root` (user ID 0).
 - Ergo, as soon as you can, you should **drop root privileges**, i.e., change the user ID your program runs under to something else than 0:

```
setuid(non-privileged-uid) ;
```

- Group privileges are also important, and thus they should be dropped too:

```
setgid(non-privileged-gid) ;
```

- This is arguably the biggest security improvement of them all.
- Typically, servers launch as `root` but then switch to special user IDs, created just for them and which have the minimum amount of privileges.

CONFINED SERVERS

- Servers should change the current working directory to a safe directory.
- Even so, nothing prevents them to write to any other directory: all they have to do is to provide full paths to the files they want to access.
- Sometimes you cannot do anything about it (whenever the server **must** access files all over the place).
- But sometimes your server needs files that are all located in a specific subtree of your file system.
 - If this is the case, then you should confine your server to that subtree.

```
chroot ( dir );
```

- The effect: *dir* becomes the **root directory** of your server.
 - * For instance, after your program does `chroot("/var/lib/bbserv")`, it will view the file `"/var/lib/bbserv/bblog"` as `"/bblog"`
 - * A file which is someplace else (say, `"/etc/passwd"`) is simply inaccessible.

WHY RUN AS ROOT IN THE FIRST PLACE?

- Some system calls have no effect if run unprivileged. They include `chroot`, `setgid`, and `setuid`. Also a non-root program cannot open a port below 1024.
- So your server must run as root at the very beginning, just to issue these calls and/or open its ports.
- As a consequence, once your server drops root privileges, it cannot get them back.
- In other words, the proper sequence of calls is:

```
chdir("/var/lib/bbserv");
chroot("/var/lib/bbserv");
// open master socket on port below 1024
setgid(99);
setuid(99);
```

- Of course, there are cases when you have to run your server as root all the way.
 - Then the potential of harm is huge.
 - You should be extra careful when programming such a server.

VALIDATE ALL INPUT

- Some inputs are from untrustable users, so those inputs must be validated (filtered) before being used.
 - You should **determine what is legal** and reject anything that does not match that definition.
 - Example of illegal strings: ". . .", anything starting with /, control characters (too small ASCII values) and/or characters with the high bit set (too large ASCII values).
 - But validate, **do not do the reverse** (do not identify what is illegal and write code to reject those cases)!

Strings. Identify the legal characters or legal patterns and reject anything not matching that form.

- A character sequence may have special meaning to the program's internal storage format (e.g., a slash in the name of a file). Check for these.

Numbers. Limit all numbers to the minimum (often zero) and maximum allowed values.

VALIDATE ALL INPUT (CONT'D)

- Input includes but is not limited to command line arguments, environment variables, and things received from the network.
 - Use text input as much as you can (easier to check).
- Limit the maximum character length (and minimum length if appropriate).
 - Be sure to not lose control when such lengths are exceeded.
- Tests should usually be centralized in one place so that the validity tests can be easily examined for correctness later.
- Make sure that your validity test is actually correct
 - this is particularly a problem when checking input that will be used by another program
 - these tests may have subtle errors, producing the “deputy problem” (the checking program makes different assumptions than the program that actually uses the data)

VALIDATE ALL INPUT (CONT'D)

- While parsing user input, it's a good idea to temporarily drop all privileges, or even create separate processes. This is especially true if the parsing task is complex, or if the programming language doesn't protect against buffer overflows (e.g., C and C++).
- **Validate command line arguments**
 - Attackers can send just about any kind of data through a command line (through calls such as `execve`).
 - You must definitely validate the command line inputs.
 - * In particular, never trust the name of the program reported by `argv[0]` (an attacker can set it to any value including NULL).
- **Validate file descriptors**
 - Do not assume that any file descriptor is opened and points to anything in particular.
 - Better close them all and reopen what is needed (a matter of resource management but also of security!).

VALIDATE ALL INPUT (CONT'D)

- **Validate file names**
 - Reject “globing” characters (*, ?) whenever possible.
 - * If you must glob, do so in a separate process, with limits on resources.
 - Filter dangerous file names, including:
 - * Names beginning with a dash
 - * Names with control characters (especially newlines) in them
 - * Names containing spaces
 - * Names containing characters with special meaning to the system and the programming language (e.g., <, ", ;, etc.)
- **Validate file content**
 - If a program takes directions from a file, the file must be considered suspect unless only trusted users can control its content (meaning: untrusted users cannot modify the file, its directory, or any of its ancestor directories).
 - If the file is suspect, make sure that the inputs from the file are protected as described in other places (taking data from a file is not an excuse).

AVOID BUFFER OVERFLOWS

- This is a very common and very dangerous security flaw.
- When allocating data (e.g., an array), validate the size.
 - It should be **positive**.
- When accessing data in an array, validate the index.
 - It should be within the array size, **and** positive.
- When copying stuff, check for bounds **and** for the format of the output. Especially important for strings.
 - Use “safe” functions (e.g., `snprintf` instead of `sprintf`, `strncpy` instead of `strcpy`).
 - **But** do not forget that you may thus lose the terminating null byte!
- Avoid dangling pointers at all cost.
 - Set deleted pointers to 0, and check before accessing the content of any pointer.

PERRILS OF BUFFER OVERFLOWS: A REAL-LIFE EXAMPLE

- Sendmail debug flags: `-dflag,value`
 - Sendmail `d8,100` sets flag number 8 to value 100.
- Name of config file (`/etc/sendmail.cf`) also stored in memory (before the flags).
 - `/etc/sendmail.cf` gives the path to `/bin/mail`
- Sendmail checked for maximum flag numbers, but not for positiveness.
- Integer larger than 2^{31} considered negative by C on 32-bit machines.
- `sendmail -d4294967269,117 -d4294967270,110 -d4294967271,113` changed “etc” to “tmp” in the name of the config file.

t m p

...	/	e	t	c	/	s	e	n	d	m	a	i	l	.	c	f	\0	...	flag0	...
-----	---	--------------	--------------	--------------	---	---	---	---	---	---	---	---	---	---	---	---	----	-----	-------	-----

- Attacker then creates `/tmp/sendmail.cf` which claims local mailer is `/bin/sh`.
 - debug call gives root shell!

FOLLOW GOOD PRINCIPLES FOR SECURE PROGRAMS

Least privilege. Each user and program should operate using the **fewest privileges possible**, thus limiting the damage from an accident, error, or attack.

Economy of mechanism/Simplicity. The design of the protection system should be simple and small as possible. Interfaces should be **minimal, narrow, and non-bypassable**. Trust should be minimized.

Open design. The protection mechanism must not depend on attacker ignorance. The mechanism should be public, depending on the secrecy of relatively few (and easily changeable) items like passwords or private keys.

Complete mediation. Every access attempt must be checked; position the mechanism so it cannot be subverted. For instance, in a client-server model the **server** must do all access checking.

Fail-safe defaults. The default should be denial of service.

Separation of privilege. Ideally, access to objects should depend on more than one condition, so that defeating one protection system won't enable complete access.

Least common mechanism. Minimize the amount and use of shared mechanisms (e.g. use of the `/tmp` or `/var/tmp` directories).

Psychological acceptability/Easy to use. The human interface must be designed for ease of use so users will routinely and automatically use the protection mechanisms correctly.

ONLY AN OVERVIEW

- This is just a brief incursion into security issues.
- Other things that have strong impact on security:
 - Environment variables (they are dangerous, really)
 - Race conditions
 - Calls to library functions
 - Random number generators
- For more details about secure programming, take a look at <http://www.linuxsecurity.com/docs/LDP/Secure-Programs-HOWTO/> and the references therein.
- An instructive tutorial on buffer overflow exploitation: <http://doc.bughunter.net/buffer-overflow/smash-stack.html>