

CSci 4271W
Development of Secure Software Systems
Day 12: OS Attacks and Protection

Stephen McCamant
University of Minnesota, Computer Science & Engineering

Outline

- Shell code injection and related threats
- Announcements intermission
- Race conditions and related threats
- Secure OS interaction
- OS: protection and isolation

Two kinds of privilege escalation

- Local exploit: give higher privilege to a regular user
 - E.g., caused by bug in `setuid` program or OS kernel
- Remote exploit: give access to an external user who doesn't even have an account
 - E.g., caused by bug in network-facing server or client

Shell code injection

- The command shell is convenient to use, especially in scripts
 - In C: `system`, `popen`
- But it is bad to expose the shell's power to an attacker
- Key pitfall: assembling shell commands as strings

Shell code injection example

- Benign: `system("cp $arg1 $arg2"), arg1 = "file1.txt"`
- Attack: `arg1 = "a b; echo Gotcha"`
- Command: `"cp a b; echo Gotcha file2.txt"`

The structure problem

- What went wrong here?
- Basic mistake: assuming string concatenation will respect language grammar
 - E.g., that attacker supplied "filename" will be interpreted that way

Best fix: avoiding the shell

- Avoid letting untrusted data get near a shell
- For instance, call external programs with lower-level interfaces
 - E.g., `fork` and `exec` instead of `system`
- May constitute a security/flexibility trade-off

Less reliable: text processing

- Allow-list: known-good characters are allowed, others prohibited
 - E.g., username consists only of letters
 - Safest, but potential functionality cost
- Deny-list: known-bad characters are prohibited, others allowed
 - Easy to miss some bad scenarios
- "Sanitization": transform bad characters into good
 - Same problem as deny-list, plus extra complexity

Terminology note

- Historically the most common terms for allow-list and deny-list have been "whitelist" and "blacklist" respectively
- These terms have been criticized for a problematic "white=good", "black=bad" association
- The push to avoid the terms got significant additional attention in summer 2020, but is still somewhat political and in flux

Different shells and multiple interpretation

- Complex Unix systems include shells at multiple levels, making these issues more complex
 - Frequent example: `scp` runs a shell on the server, so filenames with whitespace need double escaping
- Other shell-like programs also have caveats with levels of interpretation
 - Tcl before version 9 interpreted leading zeros as octal

Related local dangers

- File names might contain any character except / or the null character
- The `PATH` environment variable is user-controllable, so `cp` may not be the program you expect
- Environment variables controlling the dynamic loader cause other code to be loaded

IFS and why it was a problem

- In Unix, splitting a command line into words is the shell's job
 - String \rightarrow argv array
 - `grep a b c` VS. `grep 'a b' c`
- Choice of separator characters (default space, tab, newline) is configurable
- Exploit `system("/bin/uname")`
- In modern shells, improved by not taking from environment

Outline

Shell code injection and related threats

Announcements intermission

Race conditions and related threats

Secure OS interaction

OS: protection and isolation

Upcoming activities

- Problem set 1 is posted, and due Friday October 21st
- Midterm exam 1 is on Tuesday in class
- Notes from demos and other studying resources on Piazza

Outline

Shell code injection and related threats

Announcements intermission

Race conditions and related threats

Secure OS interaction

OS: protection and isolation

Bad/missing error handling

- Under what circumstances could each system call fail?
- Careful about rolling back after an error in the middle of a complex operation
- Fail to drop privileges \Rightarrow run untrusted code anyway
- Update file when disk full \Rightarrow truncate

Race conditions

- Two actions in parallel; result depends on which happens first
- Usually attacker racing with you
 - Write secret data to file
 - Restrict read permissions on file
- Many other examples

Classic races: files in /tmp

- Temp filenames must already be unique
- But "unguessable" is a stronger requirement
- Unsafe design (`mktemp(3)`): function to return unused name
- Must use `O_EXCL` for real atomicity

TOCTTOU gaps

- Time-of-check (to) time-of-use races
 - Check it's OK to write to file
 - Write to file
- Attacker changes the file between steps 1 and 2
- Just get lucky, or use tricks to slow you down

Read It Twice (WOOT'12)

- Smart TV (running Linux) only accepts signed apps on USB sticks
 - Check signature on file
 - Install file
- Malicious USB device replaces app between steps
- TV "rooted"/"jailbroken"

TOCTTOU example

```
int safe_open_file(char *path) {
    int fd = -1;
    struct stat s;
    stat(path, &s)
    if (!S_ISREG(s.st_mode))
        error("only regular files allowed");
    else fd = open(path, O_RDONLY);
    return fd;
}
```

TOCTTOU example

```
int safe_open_file(char *path) {
    int fd = -1, res;
    struct stat s;
    res = stat(path, &s)
    if (res || !S_ISREG(s.st_mode))
        error("only regular files allowed");
    else fd = open(path, O_RDONLY);
    return fd;
}
```

TOCTTOU example

```
int safe_open_file(char *path) {
    int fd = -1, res;
    struct stat s;
    res = stat(path, &s)
    if (res || !S_ISREG(s.st_mode))
        error("only regular files allowed");
    else fd = open(path, O_RDONLY);
    return fd;
}
```

Changing file references

- With symbolic links
- With hard links
- With changing parent directories

Directory traversal with ..

- Program argument specifies file, found in directory files
- What about `files/../../../../etc/passwd`?

Outline

- Shell code injection and related threats
- Announcements intermission
- Race conditions and related threats
- Secure OS interaction
- OS: protection and isolation

Avoid special privileges

- Require users to have appropriate permissions
 - Rather than putting trust in programs
- Dangerous pattern 1: `setuid/setgid` program
- Dangerous pattern 2: privileged daemon
- But, sometimes unavoidable (e.g., email)

Prefer file descriptors

- Maintain references to files by keeping them open and using file descriptors, rather than by name
- References same contents despite file system changes
- Use `openat`, etc., variants to use FD instead of directory paths

Prefer absolute paths

- Use full paths (starting with `/`) for programs and files
- `$PATH` under local user control
- Initial working directory under local user control
 - But FD-like, so can be used in place of `openat` if missing

Prefer fully trusted paths

- Each directory component in a path must be write protected
- Read-only file in read-only directory can be changed if a parent directory is modified

Don't separate check from use

- Avoid pattern of e.g., `access` then `open`
- Instead, just handle failure of `open`
 - You have to do this anyway
- Multiple references allow races
 - And `access` also has a history of bugs

Be careful with temporary files

- Create files exclusively with tight permissions and never reopen them
 - See detailed recommendations in Wheeler (q.v.)
- Not quite good enough: reopen and check matching device and inode
 - Fails with sufficiently patient attack

Give up privileges

- Using appropriate combinations of `set*id` functions
 - Alas, details differ between Unix variants
- Best: give up permanently
- Second best: give up temporarily
- Detailed recommendations: Setuid Demystified (USENIX'02)

Allow-list environment variables

- Can change the behavior of called program in unexpected ways
- Decide which ones are necessary
 - As few as possible
- Save these, remove any others

For more details...

- The first external reading is chapters from a web-hosted book by David A. Wheeler
- Reading questions will be due one week after they are posted on Canvas

Outline

Shell code injection and related threats
Announcements intermission
Race conditions and related threats
Secure OS interaction
OS: protection and isolation

OS security topics

- Resource protection
- Process isolation
- User authentication (will cover later)
- Access control (already covered)

Protection and isolation

- Resource protection: prevent processes from accessing hardware
- Process isolation: prevent processes from interfering with each other
- Design: by default processes can do neither
- Must request access from operating system

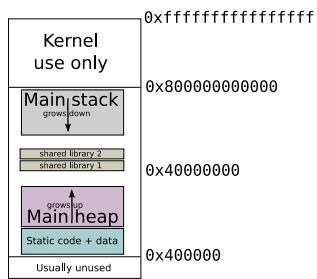
Reference monitor

- Complete mediation: all accesses are checked
- Tamperproof: the monitor is itself protected from modification
- Small enough to be thoroughly verified

Hardware basis: memory protection

- Historic: segments
- Modern: paging and page protection
 - Memory divided into pages (e.g. 4k)
 - Every process has own virtual to physical page table
 - Pages also have R/W/X permissions

Linux example



Hardware basis: supervisor bit

- Supervisor (kernel) mode: all instructions available
- User mode: no hardware or VM control instructions
- Only way to switch to kernel mode is specified entry point
- Also generalizes to multiple "rings"