Lecture 5

Juni Kim

6.S913

January 29, 2026

Today's Lecture

- No essential new material, bit more relaxed pace for questions and lab hours
- Will be reviewing concepts from the last lectures
- Some other parts of a Linux System that you might use

Writing Good Shell Programs

- We want to write programs that behave predictably and are defensive to
  - Working directory
  - Environment variables
  - Filesystem contents
- Your program must behave predictably in hostile filesystem environments or error cleanly.

Critical Distinctions

- Host vs Target: We are building up a bootable image for a target system on a host machine. (Different characteristics)
- Source vs Artifact: Developers give us source code, the compiler emits artifacts. Out-of-tree builds separate these.
- Initramfs vs Root filesystem: Initramfs is temporary, it boots the root filesystem and runs `switch_root`.

The C Compiler and Friends

- Preprocessor, Compiler, Linker stages
- Corresponding flags are `CPPFLAGS`, `CFLAGS`, `LDFLAGS`
- You need to ensure the host libc does not contaminate your build.
- `./configure` → `make` → `make install` (Roughly, differs for different source)
- Please parallelize make with the `-j` option!

Building the Kernel

- You can configure the kernel with `make menuconfig`
- See handout for drivers you should not disable (you might not have a console or networking without these)
- Booting as an EFI Stub
- Configure default command line (Firmware won't give you anything)
- Embed initramfs into the kernel (you don't have a filesystem when booting as an EFI stub)

User Daemons and System Configuration

- Most configuration files will go in `/etc`, including declaring users.
- Busybox init reads `/etc/inittab` to see what it should start when.

Important Daemons init launches

- Getty - login screen
- Dhcpcd - IP address configuration + internet
- Chrony - Time
- Eudev - device discovery

Disk Partitions and Filesystems

- `mkfs` initializes filesystems
- `sgdisk` helps you assemble disk partitions together.
- We have two partitions:
  1. ESP - readable by firmware (`/dev/vda1`)
  2. Ext4 - actually useful as a root filesystem (`/dev/vda2`)
- GPT header and metadata tells firmware where the ESP is, the firmware will locate the kernel inside the ESP.

Some Debugging Pointers

- No output: serial console or kernel cmdline issue.
- Panic early: initramfs issue or missing /init.
- Boots but no login: getty, /etc/inittab, or permissions.
- Networking dead: dhcpcd user, device nodes, or kernel config.

Other things you might want to think about

OpenSSL

- Major library not covered in this assignment.
- Implements cryptographic protocols.
- Used for SSH, HTTPS, ...

SSH

- Protocol used for securely accessing servers
- You probably accessed most linux servers via ssh
- sshd is responsible for this

SUID Binaries

- Normally, programs run with the UID/GID of the invoking user.
- Some binaries are marked **SUID** (set-user-ID).
- When executed, they run with the *file owner's* UID instead.
- Common use: controlled privilege escalation (e.g. `sudo`, `doas`, `passwd`).
- Optional: try configuring doas on your operating system!

Timezones

- Linux systems use unix millis (time since Jan 1, 1970).
- Timezones are the responsibility of userspace
- Timezone info lives under `/usr/share/zoneinfo`.
- You generally symlink one of the zoneinfo files to `/etc/localtime`

Filesystem Mounts and `/etc/fstab`

- `/etc/fstab` declares what should be mounted and where (instead of editing init which can be catastrophic).
- Entries describe:
  - device or filesystem
  - mount point
  - filesystem type
- `mount -a` will consult fstab.

System Hostname

- The kernel maintains a system-wide hostname (default set in kernel config).
- Exposed via `/proc/sys/kernel/hostname`.
- You can write to this file to change the hostname on statup.
- Used by shells, networking tools, and system services.

What do Package Managers Do?

- The process of building your own packages is quite mechanical
- It is very easy to screw up
- What if there was some automated way to do all of this in Linux...
- Package managers are the main difference between most Linux Systems.

Virtualization

- You used some of these tools for your dev environment!
- Linux has cgroups and namespaces that are used as primitives for building up containers
- CGroups: limit memory, cpu usage, ...
- Namespaces: process gets own view of filesystem, PID (so it might think it is true init), networking, ...

GUI on Linux

Display:

- The kernel exposes drm for displays (/dev/dri), userspace libraries for interacting with that
- X11 display protocol (+ x.org server)
- Window managers on top of the x.org server
- Wayland (replacing xorg on desktop)

Sound:

- Kernel exposes the sound card via ALSA (/dev/snd)
- Standard solutions are pulseaudio and pipewire.

Device registration gets handled by both the kernel and udev.

Up Next

- Final Lab Hours
- Submit the lab by Friday 11:59PM
- Thank you for participating!