

Replacing CONFIG_VT/linux-console

David Herrmann

FOSDEM 2013

What is CONFIG_VT?

- access control to graphics/input devices
- VT102 terminal emulator
- VGACON + FBCON

What is it used-for?

- session dispatcher
- lightweight UI with little hardware requirements
- emergency console
- kernel log (boot/shutdown-log, oops/panic-screen)
- backwards-compatibility

Linux-Console

- DEC “standards” VT100-VT520
- xterm de-facto standard
- more or less VT102 implementation
- lots of private extensions

Linux-Console Problems

- UI in kernel space
- poorly maintained
- bad keyboard handling
- bad font rendering
- missing modesetting/multihead support
- no multi-seat awareness
- limited hotplugging support
- limited to VT102
- blinking cursors

Solution?

- Move to user-space!
 - ▶ XKB, pango, OpenGL, ... limited to user-space
- reasons to keep the linux-console
 - ▶ oops/panic-screen
 - ▶ emergency-console

KMSCON

- started November 2011, improved during GSoC 2012
- terminal emulator in user-space, but can do a lot more

- libuterm
- TSM
- kmscon

KMSCON Features

- highly modular
- libxkbcommon, libudev, libpixman, glibc, (libdrm)
 - ▶ everything else is optional
- seat/device hotplug support
- per-seat/device configuration
- session-manager
- TSM: xterm-compatible terminal emulator state-machine
- XKB
- Unicode/Internationalization
- hardware-acceleration

Replacing agetty

- `systemctl start kmsconvt@tty1`
- `ln -s /lib/systemd/system/kmsconvt@.service /etc/systemd/system/autovt@.service`
- starts `getty@.service` on failure

Virtual Terminals

- access-control to input/graphics devices
 - ▶ hard to compare to audio (pulseaudio), v4l2, ...
- input devices provide grabs
- DRM devices provide DRM-Master logic
- accessed via `/dev/tty<num>`
- only one VT active at a time
- TEXT-MODE: controlled via `read()`, `write()`, `ioctl()`
- GRAPHICS-MODE: setup via `ioctl()`, otherwise mostly ignored

Problems with VTs

- ugly API
 - ▶ signal-based
 - ▶ unrelated to input/graphics devices
 - ▶ unmaintained
- dead VT => dead system
- no VT_WHOAMI
 - ▶ fstat() to retrieve MINOR
- no multi-seat awareness
- *Replacement required!*

What do we want?

- compositor: xserver, weston, directfb, ...
 - emergency-console
 - standalone applications (e.g., normal-console)
 - legacy applications
 - any combinations
-
- multi-seat awareness
 - sane VT API
 - unified environment across VTs
 - forced VT-switches
 - device revocation

Extending CONFIG_VT

- multi-seat awareness (/dev/seat_<seat>/tty<num>)
- sane VT-API (/sys/class/tty/tty0/active)
- forced VT-switches
- device-revocation (revoke() syscall)

- still kernel-code
- ugly code-base
- no unified environment
- maintainer?

- Move to DBus?

System-Compositor

- rudimentary & lightweight compositor
- fullscreen-only windows
- needs yet to be written

- multi-seat aware
- sane VT-API
- unified environment
- forced VT-switches
- device revocation

drmStealMaster()

- steals DRM-Master from current master
- notification via `drmHandleEvent()` (and EACCES)
- used solely for emergency
- `systemctl isolate emergency.target`
 - ▶ destroys environment

KMSCON Sessions

- session-manager controlled by global keyboard-shorcuts
- many session types
 - ▶ dummy/blank session
 - ▶ terminal-session
 - ▶ (wayland-session)
 - ▶ (dbus-sessions)
 - ▶ cdev-sessions
- implemented via CUSE (based on FUSE)
- supports full VT API
- one device per seat (/dev/ttyF<seat>)
- each open file of the cdev is a different VT
 - ▶ can still be passed via fork() or UNIX-sockets
- no replacement for CONFIG_VT
- only for backwards compatibility
 - ▶ CONFIG_VT for all seats!

fblog + drmlog

- boot/shutdown-log, oops/panic-screen
- kernel-log on all connected monitors
- on/off setting per connector via sysfs
- no eye-candy
- split-screen (width $\geq 8 * 80 * 2 = 1280$)
- drmlog needs yet to be written (ongoing work)
- fblog-v3 available on linux-fbdev@vger.kernel.org
- kmscon-session available

High-Level Overview

- seat0 (CONFIG_VT=y)
 - ▶ CONFIG_VT is primary session-manager!
 - ▶ linux-console or kmscon as terminal-emulator on single VTs
 - ▶ compositors on single VTs
 - ★ or: system-compositor on single VT
 - ★ system-compositor is secondary session-manager
 - ★ compositors as clients of system-compositor

High-Level Overview

- `$seat != seat0` or `CONFIG_VT=n`
 - ▶ (single session mode)
 - ▶ or: system-compositor
 - ★ primary session manager
 - ★ compositors as clients
 - ★ kmscon as client
 - ▶ or: kmscon as `CONFIG_VT` replacement
 - ★ cdev-sessions can also be integrated into system-compositor

What do we get?

- compositor: xserver, weston, directfb, ...
- emergency-console
- standalone applications (e.g., normal-console)
- legacy applications
- boot/shutdown-log, oops/panic-screen
- any combinations

- multi-seat awareness
- sane VT API
- unified environment across VTs
- forced VT-switches
- device revocation

Related Projects

- libuterm
- dvbe/defi (vesafb.c & efifb.c as DRM drivers)
- DRM user-space documentation
- fblog/drmlog
- Wayland System-Compositor
- getting kmscon into distributions (ArchLinux already provides kmscon-6)