

Xen and the linux console

why xencons={tty,ttyS,xvc} will go away.
Gerd Hoffmann <kraxel@redhat.com>

What is the Linux console?

- Answering this question is the point of this talk ...
- For most users it probably is the screen they are sitting in front of.
- The linux kernel has two different (but related) subsystems.
- Both have a bunch of `CONFIG_*_CONSOLE` config options

Virtual Terminals (CONFIG_VT)

- Using Alt-Fx users can switch between different terminals on the same physical screen.
- Each terminal has its own device: `/dev/tty<nr>`.
- Most Linux distros have getty's running on `/dev/tty{1-6}`, the X-Server uses `/dev/tty7`.
- `/dev/tty0`: The terminal which is visible at the moment.
- The VT subsystem needs a hardware specific drivers for actual output.

VT drivers (incomplete)

- `CONFIG_VGA_CONSOLE`
 - Drives VGA card in text mode.
- `CONFIG_FRAMEBUFFER_CONSOLE`
 - Provides text screens on top of a graphical display.
 - The graphical display is handled by some fb driver: `vesafb` (generic), `matroxfb`, `pvfb` (xen), ...
- `CONFIG_DUMMY_CONSOLE`
 - Used when no other driver is present.
 - Solves initialization order issues and usually runs for a short time at boot only.

Introducing /dev/console

- /dev/console is the linux console device.
- The kernel messages (printk) will go to the console device.
- The kernel boots /sbin/init with /dev/console as terminal.
- Often linked to /dev/tty0.
 - One reason for the confusion ...
 - That is only one option though.
 - Especially there is no strong connection between virtual terminals and /dev/console.

About Console drivers

- There are two kinds of console drivers:
 - Some can just print messages.
 - The others are associated with a full-featured terminal.
- The earlyprintk code is just a console driver too.
 - quite limited, print-only, no mem alloc for early boot operation.

Console driver list (#1)

- CONFIG_VT_CONSOLE
 - Console on virtual terminals. Most popular one.
- CONFIG_SERIAL_*_CONSOLE
 - Console on a serial line.
- CONFIG_LP_CONSOLE
 - Print your kernel messages on paper.
- CONFIG_NETCONSOLE
 - Send messages to a log server.

Console driver list (#2)

- CONFIG_HVC_CONSOLE
 - Hypervisor console infrastructure.
 - Two simple functions to send and receive characters are needed
 - You'll get a full-featured terminal device.
 - Created by the powerpc folks.
- There are alot more console drivers, most of them are architecture-specific.

XenLinux console, sparse tree (#1)

- It is a terminal driver, registers as console (aka `/dev/console`).
- It can hijack the console (aka VT subsystem) major/minor number range (`xencons=tty`).
 - Default on guest domains.
 - Convenient, gives `getty` without extra config.
 - Using virtual terminals at the same time is impossible. Showstopper for `pvfb`.
 - Code must deal with hijack fallout (emulate `/dev/tty{2..8}`).

XenLinux console, sparse tree (#2)

- It can also hijack the serial line major/minor numbers (xencons=ttyS)
 - Default on the control domain.
 - Makes it impossible for the dom0 kernel to use the serial line directly.
- Recent (3.0.4+) drivers can use the xen console major/minor number (xencons=xvc).
 - YES! Make it the default NOW!
 - Needs some manual config (stay tuned) until distros do it for you.
 - No conflicts with other drivers.

XenLinux console, paravirt_ops

- Remember CONFIG_HVC_CONSOLE?
- The xen console code in the paravirt_ops patch queue just uses that.
- Very few lines of code, the infrastructure in hvc_console.c handles the interfacing with other parts of the kernel for us.
- xencons= is gone. The console device is /dev/hvc0.

Linux console config, background

- The Linux kernel can have (and usually has) multiple console devices compiled in.
 - Who registers first is the default console.
 - Initialization order is link order, drivers/Makefile decides which console driver becomes default.
 - Can be overwritten on the kernel command line using `console=name,options`.
 - Multiple consoles can be specified, `printk` goes to all consoles, `/dev/console` is linked to the last one.

Linux console config, examples

- `console=ttyS0,115200n8`
 - Classic way to setup a serial console.
- `xencons=xvc console=xvc0`
 - 3.0.4+ sparse kernels, without virtual framebuffer.
- `xencons=xvc console=xvc0 console=tty1`
 - 3.0.4+ sparse kernels, when using the framebuffer.
- `console=hvc0`
 - `paravirt_ops` kernels.

Linux console config, userspace

- Start a getty on `/dev/hvc0` (or `/dev/xvc0`) for login.
 - Edit `/etc/inittab` (usually has a commented serial console sample line)
`h0:12345:respawn:/sbin/agetty -L 9600 hvc0 screen`
 - Add `hvc0` to `/etc/securetty` for root logins.
- `/dev/hvc0` is present even when not configured as console.
 - No boot messages then of course.
 - But getty's friendly login prompt will show nevertheless.

That's it. Questions?