Kernel Summit 2009
End User Panel
Pain Points for Using Linux

By
Tim Bird
Sony Corporation

Agenda

- Perspective of Linux Status
- Sony Pain Points
- Examples
- Ideas
- Issues for embedded

Linux status

- Linux has achieved (or will soon achieve) world domination in all major OS categories except desktop
- Linux is the new monopoly
  - Network effects create and sustain monopolies
  - Linux ecosystem has LOTS of network effects that entrench Linux as the dominant (open source) OS for the foreseeable future
  - Linux core developers, lacking a profit motive, are arguably more benevolent than previous monopoly holders

Pain Points in Using Linux

- "Version Gap"
  - The difference between what CE developers work on daily and mainline
- Good news: It used to be a "version chasm".
  - 3 years ago most companies were using 2.4 kernels for *new products*
  - They were at least 4 years behind mainline
  - Now, many companies have moved to at least 2.6.11

Causes of Version Gap

- Development cycle
  - Choose kernel version first, then customize, debug, test and ship
    - Development cycle is from 6 to 18 months
- Board support lag
  - CE vendors rely on 3rd parties to do a lot of board support
    - OS vendors, Semiconductor vendors
  - Usually 2 versions behind
- Custom drivers for SOC features
  - OK – sometimes binary, but always well behind mainline
  - SOC drivers come from chip vendors
  - Often, another 2 versions behind
- Version stickiness

Version Gap

- Example 1:
  - Sony is now selecting kernel version for 2011 TV sets
  - Version will likely be 2.6.29
  - Why?
    - We’ve used it before
    - We have drivers and support for major features we need
    - It has known characteristics

Sony patches to 2.6.29

- Sony maintains 1029 patches to kernel
  - 637 – features developed external to Sony, which have not been mainlined
    - RT-preempt, lttng = 80% of patches
    - Others are ftrace backport, oprofile, axfs, kmemleak, linux-tiny
  - 164 – board support not mainlined
  - 93 – RT fixes and tuning for Sony platforms
  - 68 – Local features and bug-fixes
    - Exception monitor, memory analysis
  - 34 – patches specific to internal build system
  - 28 – fixes backported from later kernels

SONY

Sony patches to 2.6.29

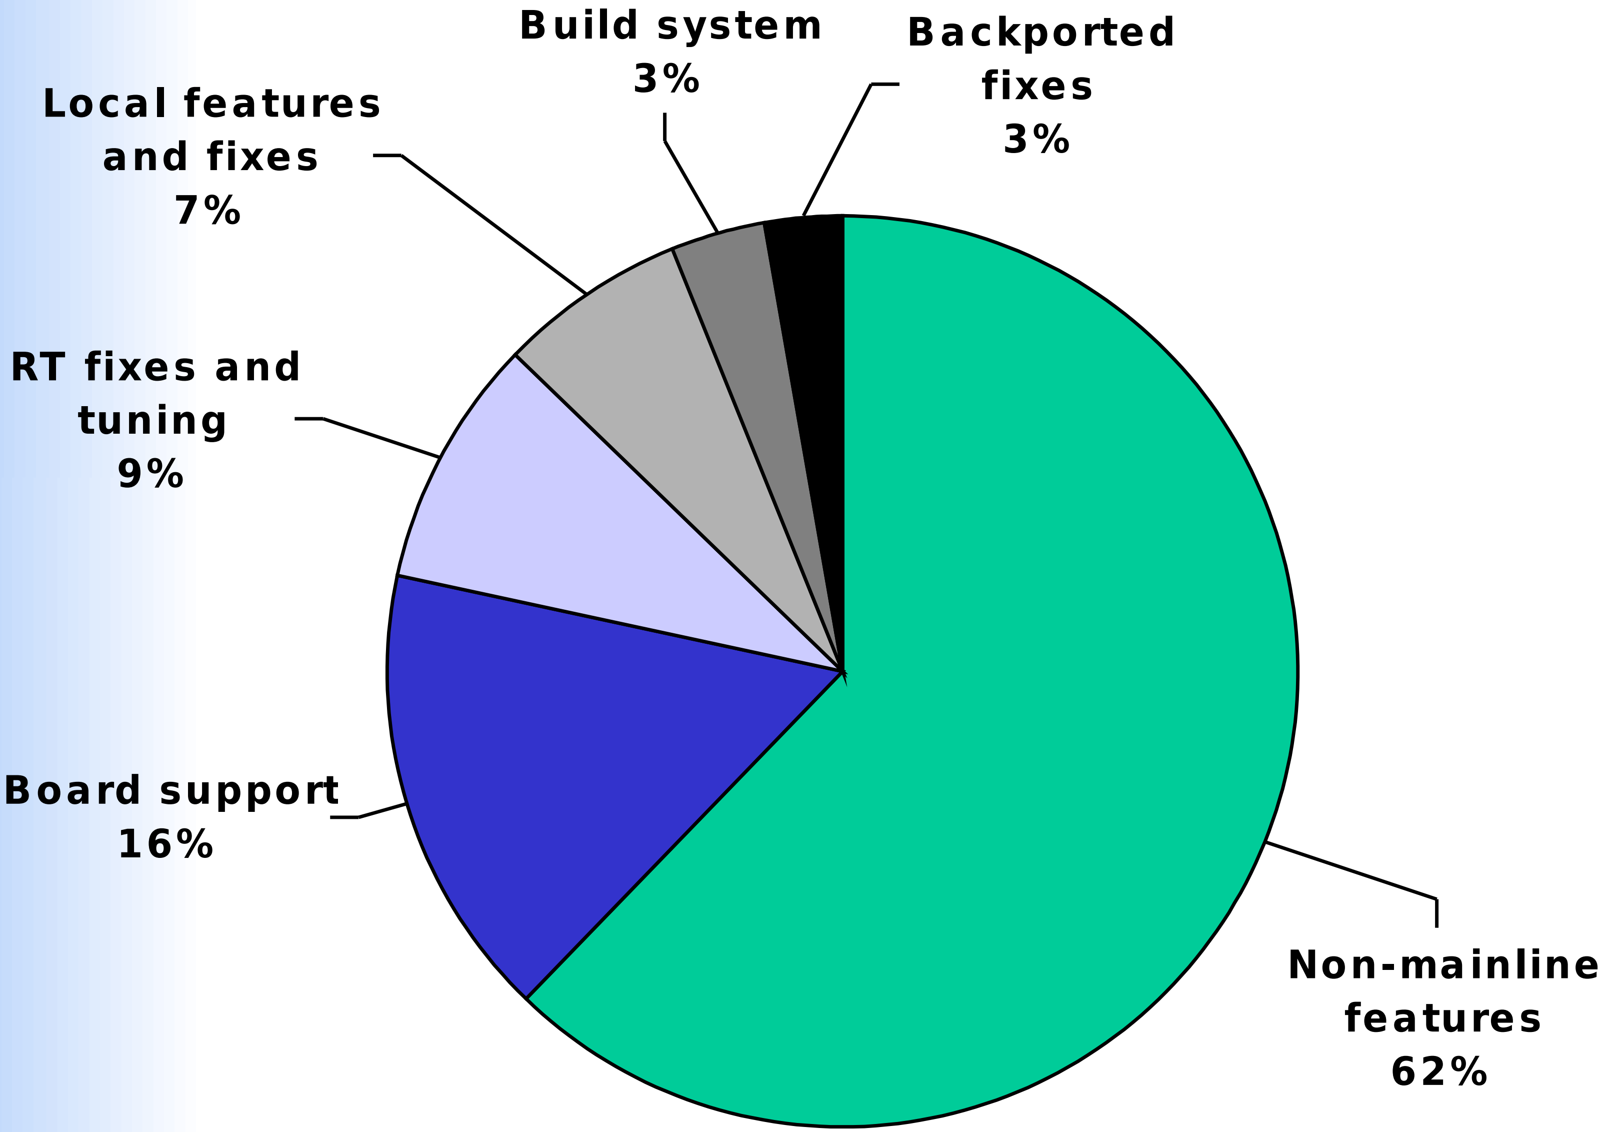

Mainlining patches

- Requests for extended scope, semi-related work
  - Outside scope of original work, and outside work experience of patch author
  - Ex: when “memory notifications for cgroups” was submitted, the author was requested to write a new generic event mechanism for cgroups
  - Ex: Adding ‘notrace’ to certain clock routines to avoid ftrace recursion resulted in request to re-write the clock system for ARM OMAP
    - Huh? I’m hanging on by my fingernails already here

Miscellaneous pain points

- Lack of contribution
  - Would like to contribute, but code developed for embedded is often not good enough to contribute
- Hearing about how crummy embedded developers are, for not contributing
  - Message has softened lately
- Developers who do this part time are very slow at it
  - Often don’t know git well, or other practices

Ideas

- Would be nice to have lower barrier to switching versions
  - Out-of-mainline patches are biggest problem
    - RT-preempt, LTTng = 44% of maintained patches
  - Second biggest problem is out-of-mainline board support
- Config bisect would be nice

CE Industry issues

- Size
- Bootup time
- File Systems (flash-aware)
- Power management
- Memory management for constrained devices
- Video/audio drivers
- Security

Issues Notes

- Size
  - Moore’s law saves us
  - Kernel is growing slower than rest of system
- Bootup time
  - Lots of improvements -- thanks!
  - Target = .5 second kernel boot
  - Mostly dependent on external factors now
    - USB, networking bringup, etc.

Issues Notes (2)

- File systems (flash aware)
  - Squashfs and UBIFS are nice
  - Need scalable boot time and good runtime performance (LogFS?, AXFS?, PCM?)
- Power Management
  - Need to support mostly-asleep systems
    - E.g. wakelocks
  - Support for runtime device PM
    - Clock management
    - Device suspend

Issues Notes(3)

- Memory management for constrained devices
  - OOM notifications so user-space can manage app lifecycle
- Video/audio drivers
  - RT is probably biggest deal here
  - DirectFB and GStreamer support in-kernel
- Security
  - SMACK and Tomoyo and Android security wonkiness give us plenty to play with for a while – Thanks!

Thanks for your time