



Linux
Plumbers
Conference | Richmond, VA | Nov. 13-15, 2023

Embedded Linux BOF

November 2023

Linux Plumbers Conference

Tim Bird

Principal Software Engineer, Sony
Electronics



Outline

Technology Areas

Linux Kernel

High Level Issues List

Community

Discussions



Historical Embedded Linux Focus areas

- System size
- Boot time
- Power management
- Realtime
- Security
- Audio drivers
- Video Drivers
- Flash filesystems (MTD)
- Support for processors, SOC's, and boards
- Arch support and drivers



Technology Areas

- Architectures
- Bootloaders
- Boot time
- Filesystems
- Networking
- Security
- Testing
- Tools
- Toolchains
- Tracing
- System Size
- Build Tools and Distros



Technology Area Poll

What embedded Linux technology area are you most interested in?

- A) Architectures
- B) Bootloaders
- C) Boot time
- D) Filesystems
- E) Networking
- F) Security
- G) Testing
- H) Tracing
- I) System Size
- J) Build Tools and Distros



Architectures

- LoongArch support recently added (v5.19 and 6.x kernels)
- Support for many old ARM boards was recently removed
- Proposed removal of old Architectures
 - Super-H
 - Itanium



Bootloaders

- U-Boot
 - Now supports loading images over HTTP
 - Previously only supported the UDP protocol
 - Could only use NFS or TFTP as servers
 - Now can download kernel and other images (dtb, initrd, etc.) from a web server
 - See <https://www.linaro.org/blog/http-now-supported-in-u-boot/>
- Snagboot
 - See next page



snagboot

- Snagboot is a set of tools that can help boot and install images on boards that fail to boot
- Produced by Bootlin
- Consists of
 - snagrecover – to initialize memory and run your bootloader
 - snagflash – to flash a working system image using either DFU, USB Mass Storage, or fastboot
- Works on several different boards, and replaces proprietary tools
- See <https://bootlin.com/blog/releasing-snagboot-a-cross-vendor-recovery-tool-for-embedded-platforms/>



Boot Time

- Resurrection of Ureadahead
 - by Steve Rostedt
 - Presentation: https://elinux.org/images/5/5b/Ureadahead_resurrection.pdf
 - Video - <https://youtu.be/HwdWKMxM83E>
- What else is going on?



Boot-time Tool poll

What tool do you use to measure boot time?

- A) Stopwatch
- B) Logic Probe
- C) Printk times / grabserial (output timing)
- D) Bootchart
- E) Kernel Function Tracer (ftrace)
- F) Something else



Filesystems and I/O

- Some work on MTD spi-nor
 - Enhanced locking to support reads while writing
- EROFS enhancements
 - caching and speed improvements
- Lots of tweaks to existing filesystems and drivers
 - Not much specific to embedded



Filesystem Poll

What filesystem are you using in your project?

- A) ext4
- B) fat/exfat
- C) f2fs
- D) squashfs / pramfs / cramfs
- E) Something else





Realtime Poll

What realtime system are you using?

A) PREEMPT_RT

B) Xenomai

C) I'm not using RT in my project



Real-Time

- PREEMPT_RT status
 - Sleeping locks was mainlined (v5.16)
 - Patches have been going in continuously – through 6.6
 - Although some stalled (like the printk refactor)





PREEMPT_RT - What's left

- What's left in PREEMPT_RT patches out of mainline:
 - last year (2022): About 1300 lines of code, affecting 92 files (in 51 patches)
 - this year (2023) (patches-6.6-rc6-rt10): About 5934 lines of code, affecting 162 files (in 169 patches)
 - Some changes to the printk, 8250 serial driver, the core scheduler, some locking and timer tweaks, and a few other places.
 - New 'nbcon' – non-blocking console driver
- See <https://mirrors.edge.kernel.org/pub/linux/kernel/projects/rt/6.6/>



System Size

- SLOB memory allocator removed from kernel (v6.4)
 - Not enough people using it (or reporting that they use it)
 - Was deprecated in v6.2
- Patches for SLAB removal are already submitted
- Bloaty McBloatface size profiler tool
 - Tool by Google to analyze size of binary images, open sourced in 2006
 - Does deep inspection of ELF binaries
 - Can show comparison between two binaries
 - See <https://github.com/google/bloaty>



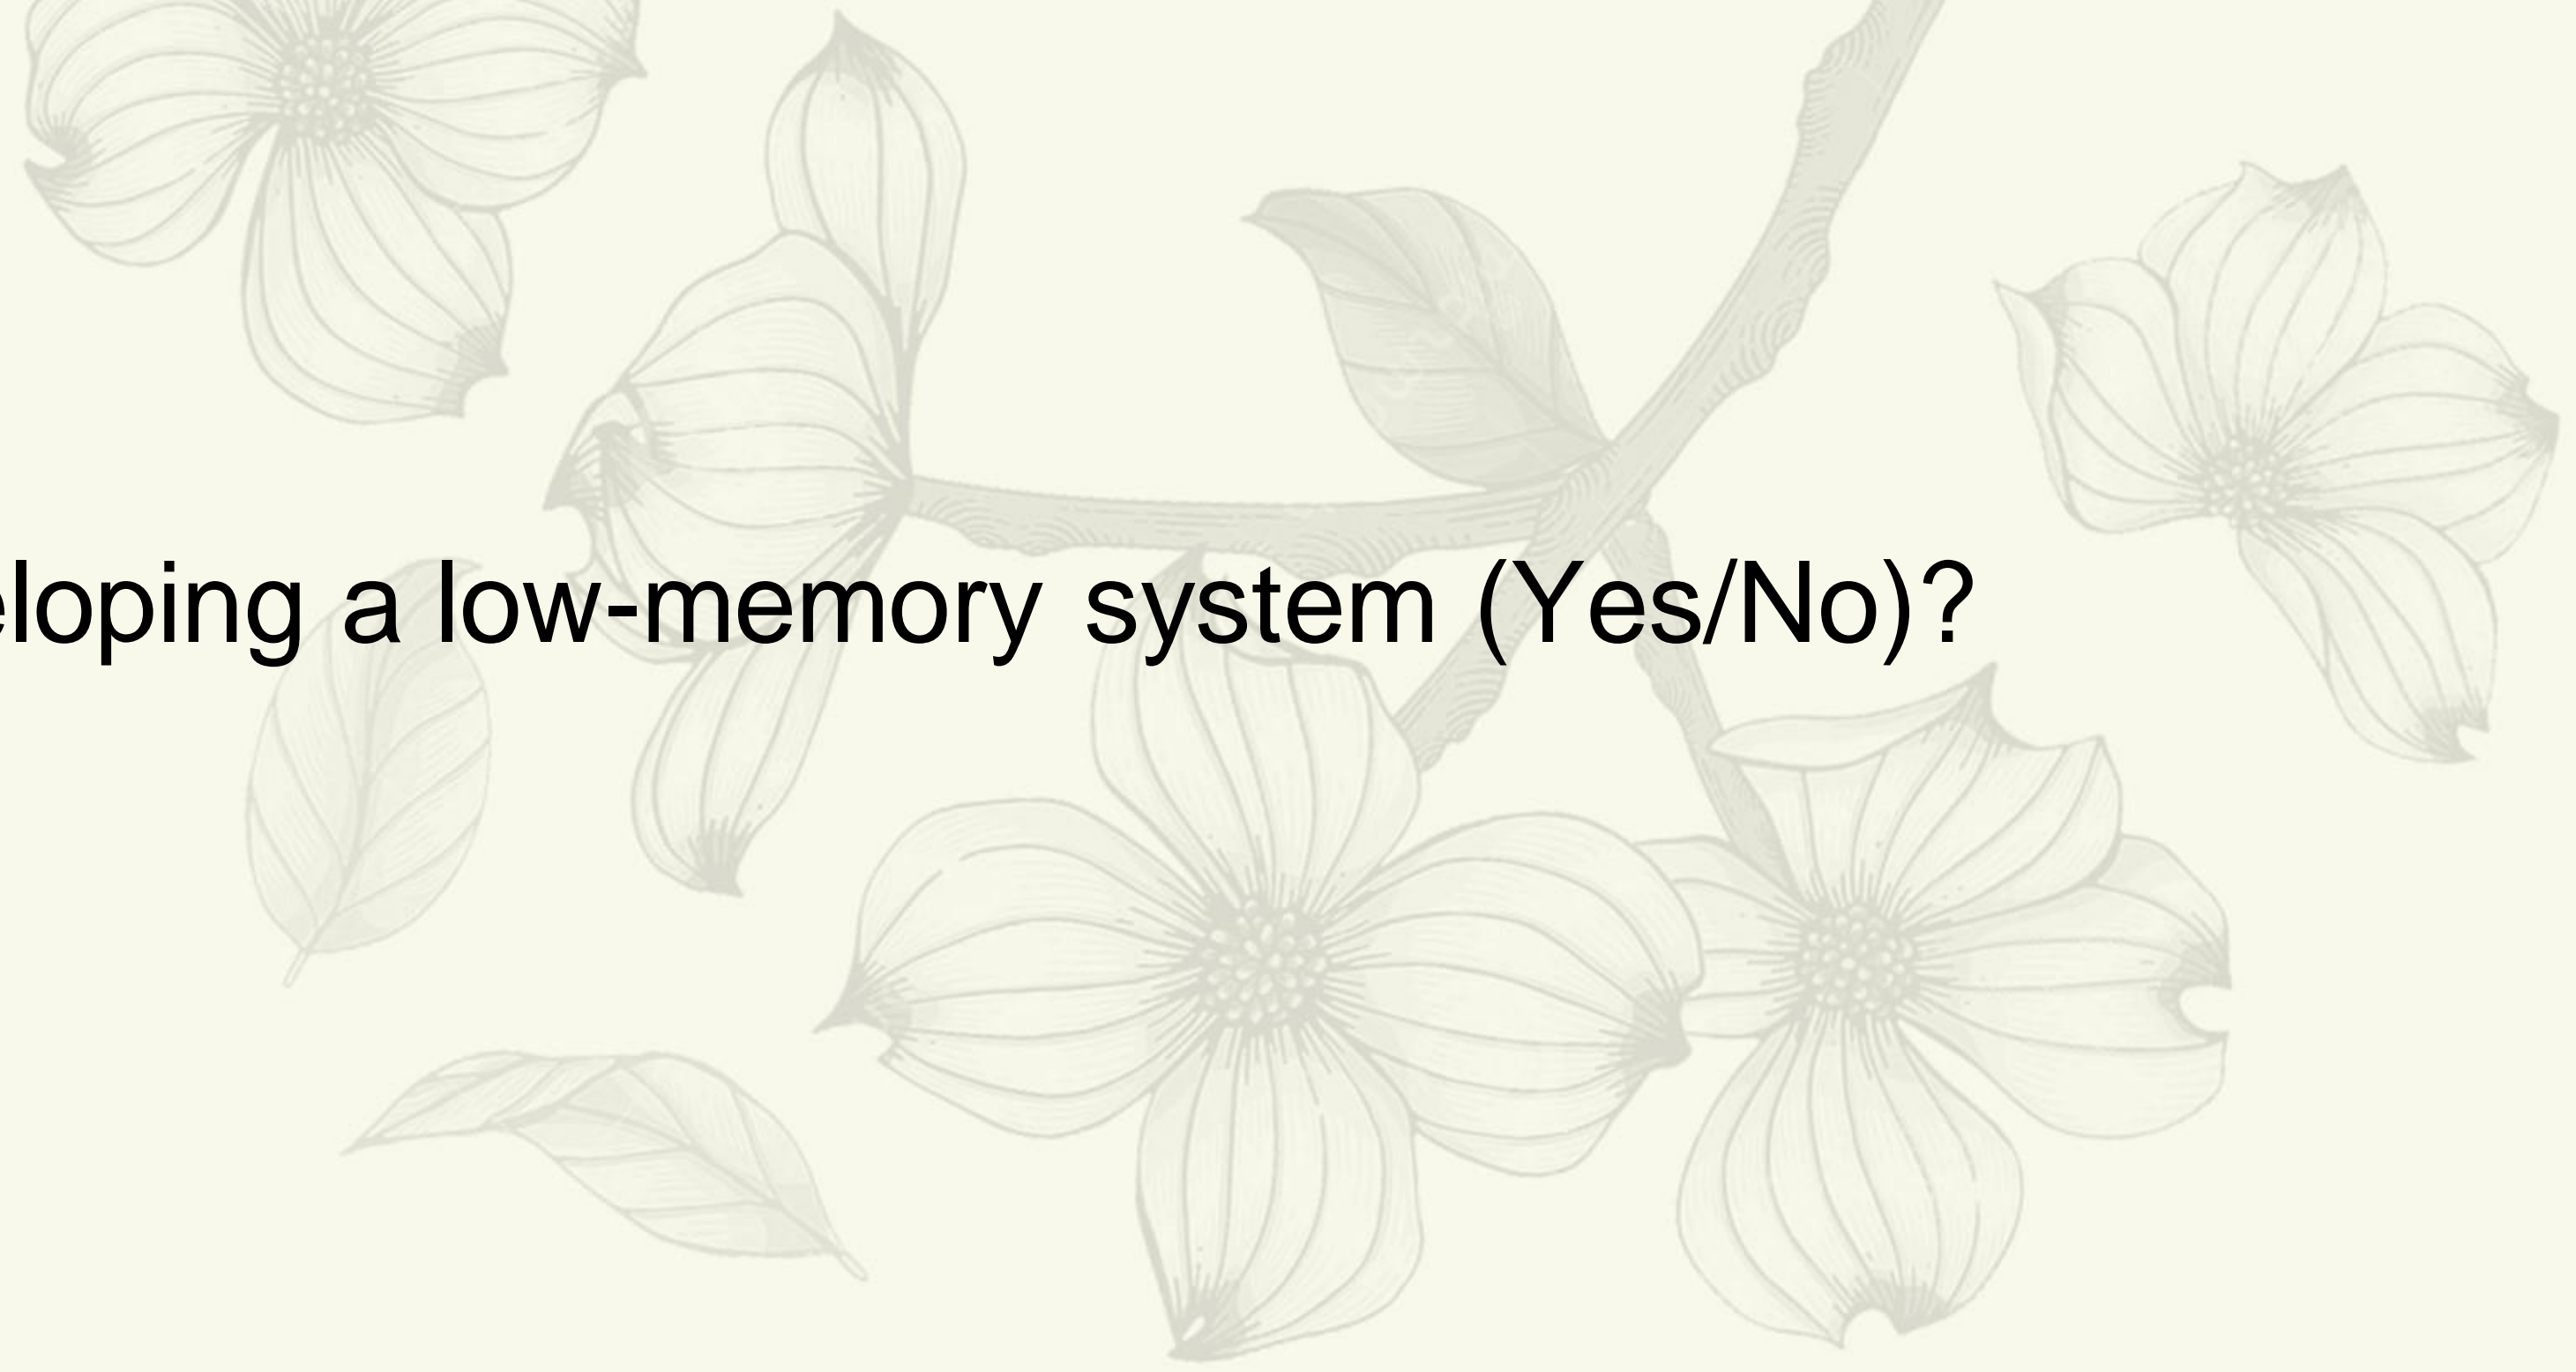
nolibc

- Is a minimal C-library replacement shipped with the kernel
 - Originally created for kernel testing
 - Consists of a set of .h files, which create macros to replace common C library functions
- Allows creation of a very small static binary
 - For some binaries, the statically linked binary using nolibc is smaller than the binary using libc (using dynamic libraries)
 - e.g. hello-libc size=1159 bytes, hello-nolibc = 1100 bytes
- Lots of issues to look out for, but an interesting approach to create minimal binaries in Linux
- See <https://lwn.net/Articles/920158/>



nolibc poll

Have you used nolibc for developing a low-memory system (Yes/No)?





System Size Tool Poll

What tool do you use for system size reduction?

A) Measuring binary size with 'ls -l' or readelf or objdump

B) size

C) top / ps / free / procfs

D) bloaty mcbloatface



Build Tools and Distros

- Yocto Project
 - Latest version = 4.2 (Mickeldore), released May, 2023
 - Kernel 6.1 and 350+ recipe upgrades
 - Improved memory and disk usage
 - Improved recipe parsing time
 - Now requires Python 3.8
 - Bunch of other changes
 - See release notes:
 - <https://docs.yoctoproject.org/next/migration-guides/release-notes-4.2.html>



Outline

Technology Areas

Linux Kernel

High Level Issues List

Community

Discussions



bootconfig

- Allows passing arguments to kernel as structured key/value pairs
- bootconfig uses a simple text file syntax
- Can be attached to an initrd image, or embedded in the kernel image
- This is useful for passing a large number of tracing options, which would normally make the boot command line too long and complex:
 - e.g. ftrace options, events, filters, actions, probes, fields, etc.
 - Can set up kernel boot-time tracing more easily
- See <https://docs.kernel.org/admin-guide/bootconfig.html>



Linux v6.2 (February 2023)

- More Rust infrastructure code was added
- Improvements to Squashfs
- New 'rv' tool added to control operation of the runtime verification subsystem



Linux v6.3 (April 2023)

- Lots of old unused Arm board files were removed
- Kernel can be configured with a built-in Dhrystone test
- Default “make V=0” option has been removed
- Minor change to Developer’s Certificate of Origin clarifies that nicknames can be used for contributions
 - As long as your identity is known
 - See <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=d4563201f33a>



Linux v6.4 (June 2023)

- The SLOB memory allocator was removed
- Some nice documentation added for building the kernel
 - <https://docs.kernel.org/admin-guide/quickly-build-trimmed-linux.html>
- MODULE_LICENSE() declarations were removed, for code that cannot be built as a module
 - See <https://lwn.net/Articles/927569/>
- User trace events (API fixes) was merged
 - See <https://lwn.net/Articles/927595/>



Most-Active Organizations, by Subsystem

- Here are the most-active organizations, by kernel area:

Kernel Area	Organizations with over 5% total contributions from 5.18 to 6.3 (approx. 1 year)
Core kernel	Google, Oracle, Huawei, Red Hat, Meta, Intel
Arch	Linaro, Google, IBM, (Unknown), Intel
Drivers	Intel, AMD, (Unknown), Linaro, (None)
Filesystem & block layer	Red Hat, SUSE, Oracle, Meta, Huawei, (Consultant)
Networking	Intel, Red Hat, Google, Meta, (Unknown), NVIDIA
Documentation	(Unknown), (None), Google, Intel, Red Hat, Meta, Loongson

Table data:



Contributions by embedded Linux companies (to kernel)

Company	company commits since July 2022	Top contributor	Work area(s) (of top contributor)
Baylibre	216	Corentin Labbe	rockchip, crypto
Bootlin	356	Miguel Raynal	network, mtd spi-nor, nvmem
Collabora	757	AngeloGioacchino Del Regno	mediatek processor support
Ideas On Board	485	Laurent Pinchart	camera, media drivers
Igalia	103	Guilherme G. Piccoli	drm AMD GPU driver, pstore
Konsulko	6	Matt Ranostay	USB to I2c driver fixups
Linutronix	346	Thomas Gleixner	timers, x86, preempt_rt
Pengutronix	1980	Uwe Kleine-König	i2c, iio, leds, driver cleanups, mtd
Toradex	199	Marcel Ziswiler	dts work on ARM platforms
Wind River	22	Paul Gortmaker	powerpc cleanups



Outline

Technology Areas

Linux Kernel

High Level Issues List

Community

Discussions



Issues From Embedded Linux Leadership Summit 2023

- Security is an ongoing concern
- OTA updates
- Long-term support
- Heterogeneous core support
- Shared testing
- Safety critical / certification
- Standardization of HW components
- AOSP competition
- Lack of investment in infrastructure/
tools/ upstreaming
 - (not enough contributions)
- RT marketing
- Upstream AI/ML accelerators
 - Support for AI processing units (APU)



Outline

Technology Areas

Linux Kernel

High Level Issues List

Community

Discussions





Embedded Linux Community

- Organizations and Projects
- Events
- General Resources
 - email lists, elinux wiki





Embedded Linux email lists

- For kernel code:
 - linux-embedded@vger.kernel.org
 - Status: not used very much, never suggested by get_maintainers.pl
 - The last substantive discussion was about the MAINTAINERS file, in 2021
- For strategy and community-building discussions:
 - elinux-discuss@lists.elinux.org
 - Status: brand new



Organizations and Projects

- Linaro – does ARM upstreaming and projects
- Linux Foundation – (see next page for projects)
- Robot Operating System (ROS) – handles robotics vertical
- Android Open Source Project (AOSP) – handles mobile phone vertical
 - And there are others (e.g. LineageOS, PinePhone, etc.)
- Yocto Project, Buildroot, OpenWRT– build systems for embedded Linux and routers



Linux Foundation projects

- Linux Foundation
 - Core Infrastructure Project (CIP) – handles support longevity
 - ELISA – handled issues with safety certification and standards
 - OpenChain – handles issues with supply chain
 - SPDX – Deals with licensing issues and SBOMs
 - Automotive Grade Linux (AGL) – handles automotive vertical
 - KernelCI – handles automated testing (for upstream)
 - Yocto Project – build system for embedded OSS (not just Linux)
 - DroneCode – handles drone vertical
 - Core Embedded Linux Project – is shutting down



Events - ELC

- Embedded Linux Conference
 - Is now only once a year
 - Alternates between Europe and North America
 - So, effectively only once per region every two years
 - May decide to fill in with regional events or Plumbers
 - Sometimes in Embedded Open Source Summit, and sometimes in Open Source Summit (North America)
 - Goal = Consistently in the spring (April/May)
 - next one: 2024 = April 15-19 in Seattle, USA





Events - Others

- Embedded Recipes
 - September 28-29, 2023 in Paris
- Linux Plumbers
 - This one: November 13-15, 2023 in Richmond, Virginia, USA
- FOSDEM – embedded track
 - February 3,4 2024 (not confirmed) in Brussels
- Regional, topic-specific:
 - Japan Technical Jamboree, ...





Elinux wiki

- Site is still used for:
 - Materials for embedded Linux development boards
 - Some academics use it for coursework
 - Event materials: slides and links to videos for ELC
- Some areas of the site are out-of-date
- I feel like the site is underutilized for sharing information
- What to do next....? (Any volunteers?)
- Setting up LFX Crowdfunding account to continue funding administrator
- Please https://elinux.org/ELinuxWiki:Mailing_List



eLinux mailing lists

- See https://elinux.org/ELinuxWiki:Mailing_List
- elinux-discuss@lists.elinux.org: for discussion of ecosystem and strategy
- elinux-wiki-dev@lists.elinux.org: for discussing development of the elinux wiki
 - Join this if you want to help out with maintaining the elinux wiki site



Outline

Technology Areas

Linux Kernel

High Level Issues List

Community

Discussions



Projects in new verticals

- Linux in Space
 - cube sats, Ingenuity, Perseverance
- Other verticals





Automating kernel specialization and tuning

- Embedded Linux often requires specialization
 - subtractive engineering (system size, boot time, power management)
- Kernel specialization is often not possible by non-kernel-experts
 - e.g. Linux in space – very few of the developers are kernel experts
- Kernel specialization is hard and time-consuming, even for kernel experts
 - Specialization can't be done manually – there's too much to do
 - size reduction, realtime, bootup time, power management can be 'death-by-a-thousand-cuts'
- Propose to find ways to automate specialization
 - e.g config bisect for static and dynamic kernel size



Discussions


- What is the kernel lacking for embedded projects?
- What do you think is missing to support embedded development?
- How can we build a stronger embedded Linux community?





Linux
Plumbers
Conference | Richmond, VA | Nov. 13-15, 2023

Thanks!



Linux Plumbers Conference

Richmond, Virginia | November 13-15, 2023

