



Contribution ID: 4

Type: **not specified**

linux/arch MC

Historically, the code in arch/ was developed for one architecture and then copied and adjusted by others. This created a lot of duplicated or almost duplicated code with subtle differences which prevents easy refactoring and consolidation.

The linux/arch microconference aims to bring architecture maintainers in one room to discuss how the code in arch/ can be improved, consolidated and generalized, at least where it makes sense.

The discussion at the previous linux/arch microconference in 2020 led to updates in RISC-V kprobes implementation [1], removal of DISCOTIGMEM memory model [2] and enablement of generic entry on s390 [3].

[1] <https://git.kernel.org/torvalds/c/c22b0bcb1dd0>

[2] <https://git.kernel.org/torvalds/c/bb1c50d3967f>

[3] <https://git.kernel.org/torvalds/c/56e62a737028>

The possible topics could be:

- reducing code duplication and generalizing the common code in arch/
- making headers in include/asm consistent
- on-boarding more architectures to use common entry code
- devicetree (unless they have their own microconf)
- identifying old machine support that may be either still in active use vs only in hobbyist/retro-computing vs completely obsolete and broken

MC Leads:

Mike Rapoport <rppt@kernel.org>, Arnd Bergmann <arnd@arndb.de>

I agree to abide by the anti-harassment policy

Primary authors: BERGMANN, Arnd (Linaro); RAPOPORT, Mike (IBM)

Presenters: BERGMANN, Arnd (Linaro); RAPOPORT, Mike (IBM)

Track Classification: LPC Microconference Track (CLOSED)