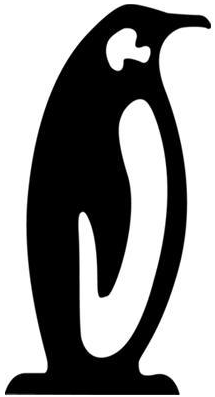

東京 2025 LINUX PLUMBERS CONFERENCE

TOKYO, JAPAN / DECEMBER 11-13, 2025

Highmem deprecation planning

Arnd Bergmann <arnd@linaro.org>

東京 2025
LINUX PLUMBERS CONFERENCE
TOKYO, JAPAN / DEC. 11-13, 2025

Highmem

- Allows more than 768MB RAM easily
- up to 16GB total on 32-bit Arm

Highmem

- Allows more than 768MB RAM easily
- up to 16GB total on 32-bit Arm
- Needs to die

Highmem

- Allows more than 768MB RAM easily
- up to 16GB total on 32-bit Arm
- Needs to die (eventually)

Refresher: 32-bit virtual memory layout

PowerPC, Arm, x86

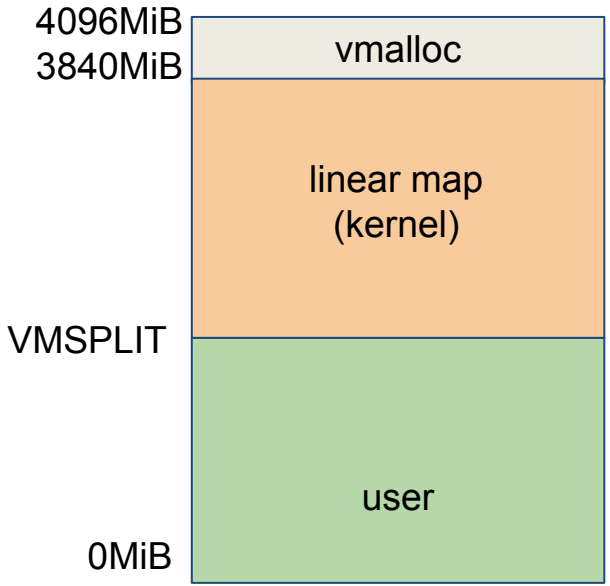

Kconfig option	User size	Linear size
VMSPLIT_3G	**3072MiB**	**768MiB**
VMSPLIT_3G_OPT	2816MiB	1024MiB
VMSPLIT_2G	2048MiB	1792MiB
VMSPLIT_2G_OPT	1792MiB	2048MiB
VMSPLIT_1G	1024MiB	2816MiB

Motivation for removing highmem

Embedded Linux maintainers

- Avoid regressions on kernel updates
- Keep 32-bit Linux viable

Linux mm maintainers

- Avoid regressions on embedded systems
- Simplify code base
- Prevent ideas of 64-bit highmem

Motivation for keeping highmem

Cost of change

- exposing kernel driver bugs
- recertification
- Userspace changes

Hard requirements

- Systems over 2GiB
- Applications with large address use

Impact of Highmem removal, 32-bit embedded

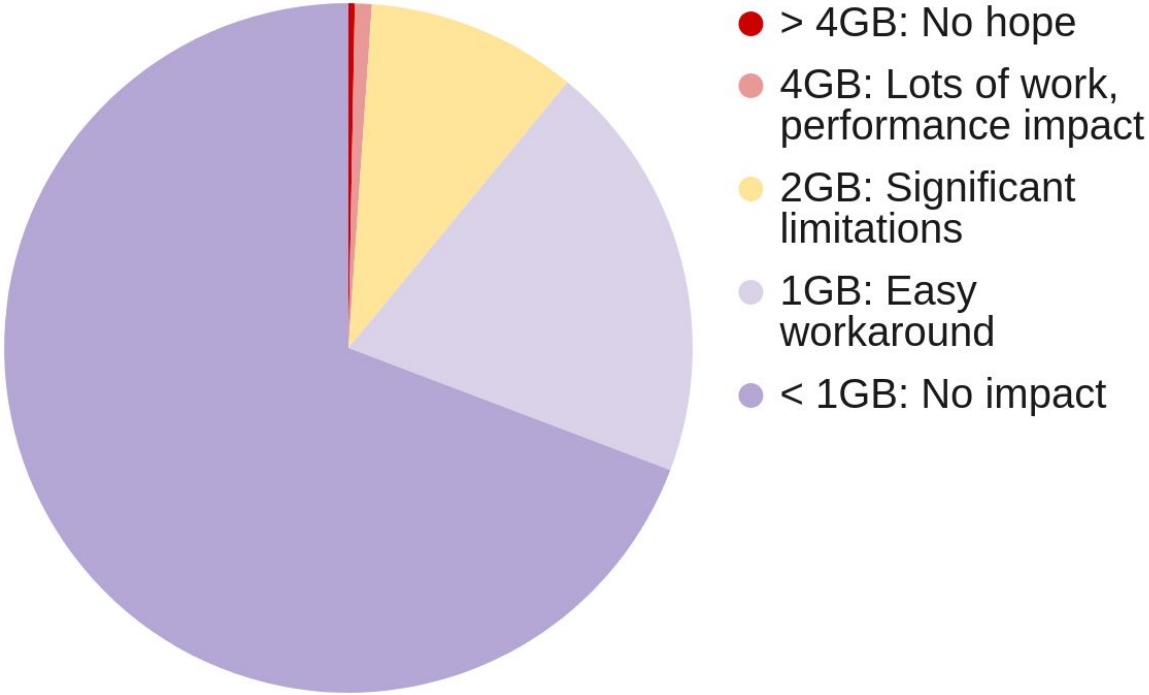

Embedded systems up to 2GB

1GB total memory

- Very common: ARMv7, 2*4Gbit DDR3
- Discontiguous physical addresses
- VMSPLIT_3G_OPT may need userspace changes (ex JVM)

2GB total memory

- Somewhat common on embedded arm/powerpc/x86
- VMSPLIT_2G_OPT (missing on Arm)
- 1.75GB user address too small for Firefox

Systems over 2GB

Known systems with 3-4GB

- x86/powerpc laptops (pre-2007)
- Arm Chromebook (2012-2013)
- In-flight entertainment
- Fire alarm system
- Digital signage

Known systems with >4GB

- Amazon/Annapurna Alpine
- Calxeda Midway
- HiSilicon HiP04
- Marvell Armada XP
- Intel/LSI axm55xx
- TI 66AK2Hxx

vmsplit across architectures

architecture	lowmem size	common physmem
arm	768MiB - 2816MiB	256MiB - 2048MiB
mips	448MiB - 512MiB	64MiB - 512MiB
powerpc	768MiB - 2816MiB	256MiB - 2048MiB
sparc (leon)	192MiB	64MiB - 2048MiB
x86	768MiB - 2816MiB	512MiB - 2048MiB
arc	2048MiB	512MiB - 2048MiB
csky	1536MiB - 2048MiB	???
µblaze	768MiB - 2816MiB	???
xtensa	128MiB - 512MiB	64MiB - 1024MiB

Proposals

Densemem

- Up to 2GB
- Needed for replacing sparsemem
- Variable vmsplit reduces user addresses

VMSPLIT_4G

- Up to 3.75GB
- Switching user vs lowmem
- Under discussion for years
- not much progress
- performance impact

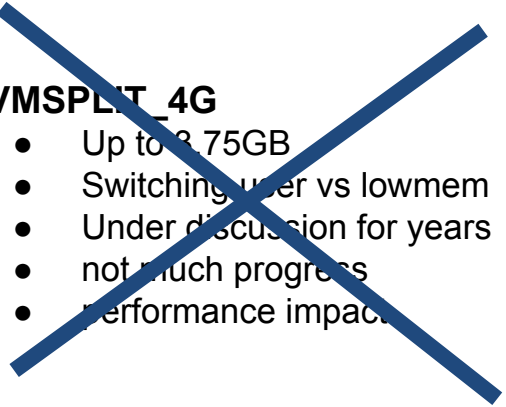

reduced-feature HIGHMEM

- Reduced complexity
- dropping HIGHMEM from page cache, highpte, drivers, ...
- reduces impact

Proposal: reduce highmem users

Still needed

- file backed mappings
- anonymous mappings
- zram
- zswap

Unclear

- hugetlb
- drbd
- dm-crypt
- percpu
- userfaultfd
- kexec
- pipe

Removal candidates

- page tables (highpte)
- fs metadata
- dmabuf
- drm (i915, amd, msm nouveau, etnaviv, ...)
- dma/iommu
- vmalloc
- kvm
- binder
- firmware
- ipu3
- infiniband
- rds
- balloon
- fuse
- uprobes

Proposal: policy changes

- Use GFP_HIGHMEM only where required, not where possible
- Phase out 32-bit desktop (3GiB/4GiB) use cases, minimize impact for embedded
- Make HIGHMEM depend on CONFIG_EXPERT
- Separate Kconfig options, accounting per highmem user
- Keep highmem page cache for 5+ more years

Proposal: possible timeline

2026 reduce highmem feature set
2026 default to VMSPLIT_2G_OPT on arm/powerpc/x86
2027 remove highmem on arc/microblaze/mips/sparc/xtensa
2027 add armv7 densemem support?
203x remove highmem page cache?
204x remove last 32-bit architectures?

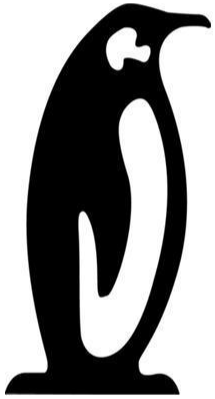

東京 2025
LINUX PLUMBERS CONFERENCE

TOKYO, JAPAN / DECEMBER 11-13, 2025