Keynote / KVM Status Report 2020

Christian Bornträger
IBM

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.
If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

The following are trademarks or registered trademarks of other companies.

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

THE LINUX FOUNDATION

State of KVM (the Linux kernel part)

- 1 core maintainers: Paolo Bonzini
- Some changes in architecture support
  - ARM32 host support is gone
  - MIPS was unmaintained and got new maintainers
    - - James Hogan
    - + Huacai Chen
    - + Aleksandar Markovic
  - ARM64, PPC, x86, s390
    - Stable maintainers
  - Increasing numbers of reviewers listed in MAINTAINERS file
- RISC-V on the horizon
  - Code looks pretty good, fully reviewed, waiting for architecture stabilization
  - Very strict rules by the RISC-V kernel maintainers

Trends in KVM

- KVM is ubiquitous in the cloud
  - Amazon, Google, IBM, Alibaba, Huawei, Tencent Cloud, ByteDance, Yandex, Oracle, …
  - Software stack range from standard QEMU to highly customized
- KVM is also used in the container and kubernetes world
  - Kata containers for container
  - KubeVirt
- Trusted computing is coming
  - AMD SEV
  - IBM Power Protected Execution
  - IBM Z Secure Execution
  - ARM protected KVM efforts (see also Will Deacon’s talk)
  - Intel TDX (see Sean Christopherson‘s talk)
- I/O: Hardware passthrough is still a hot topic but co-exists with virtio
- Test coverage improves

KVM Commits in each release (long term)

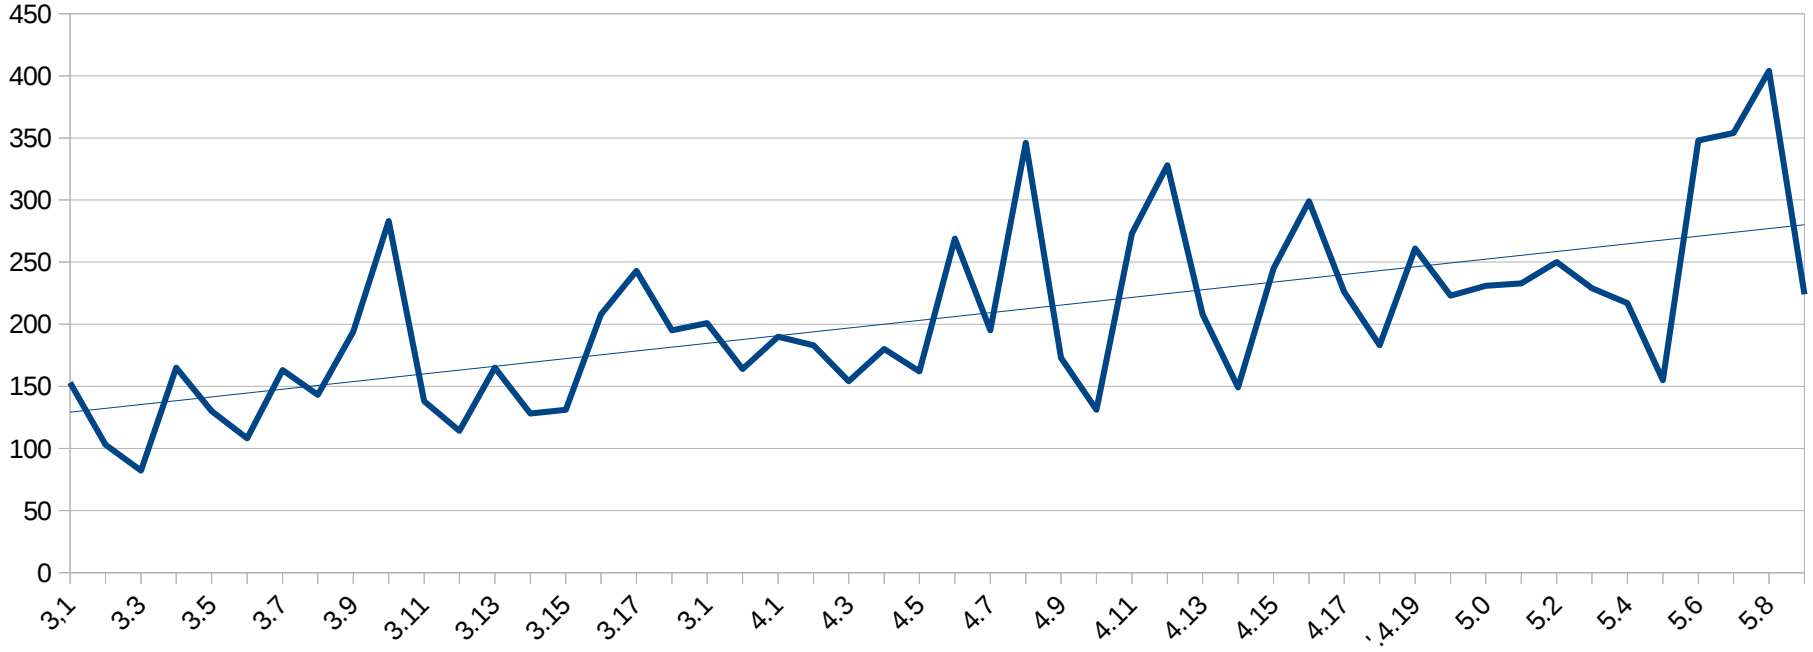

KVM change rate is still growing

Commits 5.4-rc1..5.9-rc1

- Last years numbers in ()
- 1549(1116) non-merge commits
- 159(117) merge commits for kvm files
- 579(319) commits have “Reviewed-by:”
- 101(87) commits have “Acked-by:”
- 243(222) commits have “Fixes:”
- 126(113)commits have “Cc: stable”

Top authors

419 Sean Christopherson
142 Paolo Bonzini
79 Marc Zyngier
58 Vitaly Kuznetsov
43 Miaohe Lin

Top reviewers

157 Vitaly Kuznetsov
66 Cornelia Huck
57 Jim Mattson
34 David Hildenbrand
27 Peter Shier

Top repairmen

73 Sean Christopherson
28 Paolo Bonzini
19 Marc Zyngier
13 Vitaly Kuznetsov
13 Marios Pomonis

Employers

- Commits from >20 companies
  - Redhat: Overall maintainer
  - Intel
  - IBM: Power, IBM Z
  - Google
  - Huawei
  - Marc Zyngier: Amazon/Google
  - ARM
  - Oracle
  - AMD
  - Tencent Cloud
  - Lemote,SUSE, Microsoft, Amazon, Alibaba, Canonical, Linutronix,….

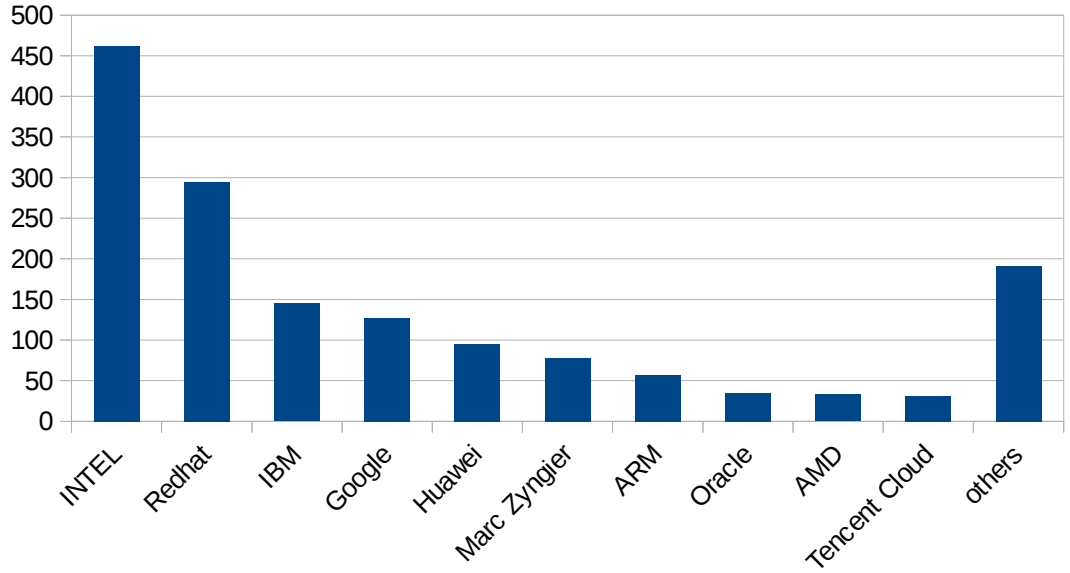

Highlights testing

- kvm-unit-tests used for testing non-KVM hypervisors
- Rehosting of kvm-unit-tests on gitlab
  - CI
- More and more selftests
- More and more kvm unit tests

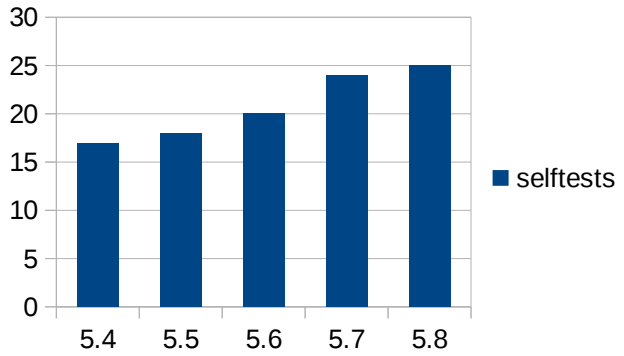

Highlights overall

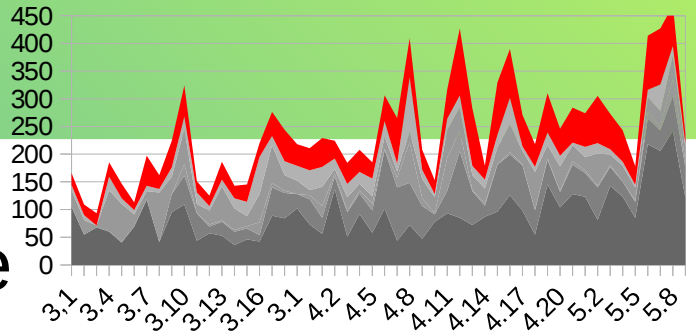

- Trusted Computing everywhere
- kvm_stat logging and CSV
- Unify shadow MMU cache data structures across architectures
- Fixes, cleanups etc.

Highlights x86

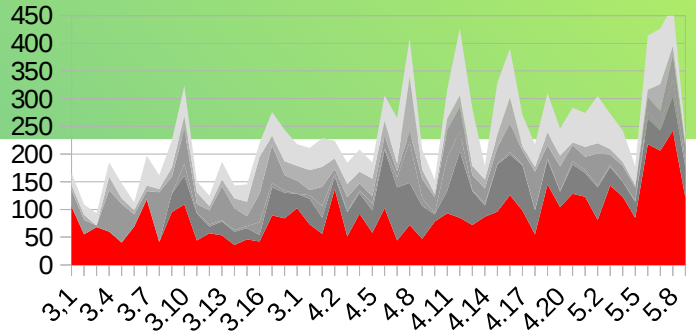

- async page fault rework
- Dirty bitmap optimization
- Refactoring and optimization (MMU, CPUID, pointer chasing, PMU, IOAPIC, TLB, event injection)
- more spectre-like work and optimization (retpoline, Spectre-v1/L1TF, TSX_CTRL)
- Nesting (5 level page tables, AMD improvements, PMU)
- SEV improvements (ASID allocation and flushing)
- Support for mapping DAX areas with large nested page table entries
- Fast path for IPI delivery and tsc deadline timer

Highlights ARM

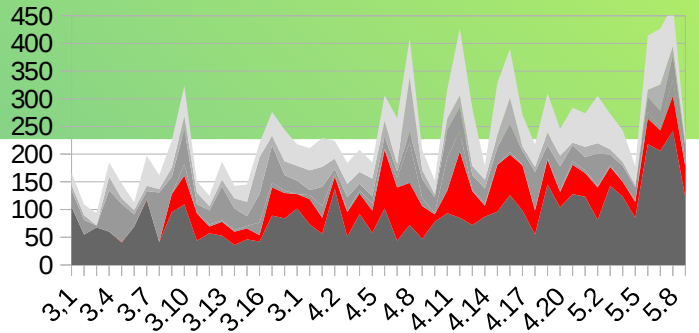

- Removal of 32 bit host support
- Interrupt controller improvements
- Split VHE and nVHE hypervisor code
- Pointer authentication for guests on nVHE
- steal time support
- data abort report and injection
- Level-based TLB invalidation support

Highlights MIPS

- Loongson support

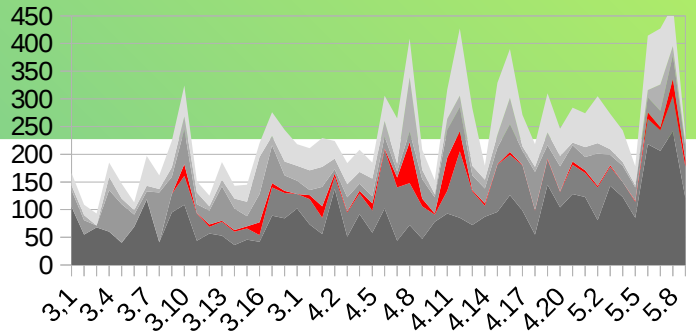

Highlights Power

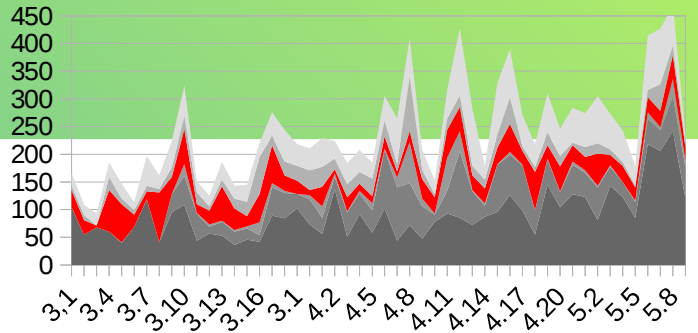

- PPC secure guest support
- Preliminary POWER10 support
- Up to 4094 guests for HV KVM
- Interrupt improvements
- Single-step capability indication

Highlights IBM Z (s390x)

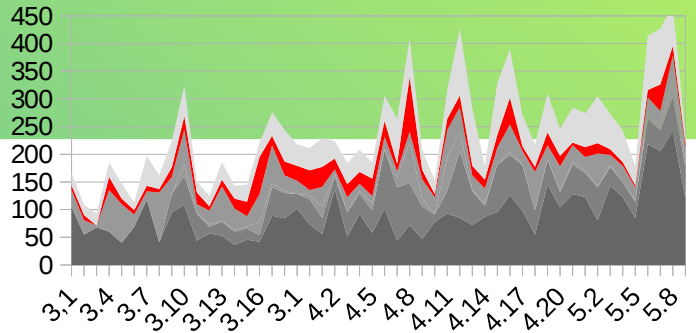

- Secure execution
- Selftests
- Yield improvements
- Nesting fixes
- Diagnose 318 handling

Forward looking (beyond 5.9)

- X86 new MMU for two-dimensional paging
- More trusted computing
  - AMD SEV enhancements (secure state and secure nested paging)
  - ARM Protected KVM?
  - Intel TDX
  - Power, s390x enhancements
- RISC-V

Have a great KVM Forum

Christian Borntraeger <borntraeger@de.ibm.com>

KVM Forum program committee <kvm-forum-2020-pc@redhat.com>

Have a great KVM Forum

Christian Borntraeger <borntraeger@de.ibm.com>

KVM Forum program committee <kvm-forum-2020-pc@redhat.com>


KVM
FORUM