



# **KVM Forum 2015 - Keynote**

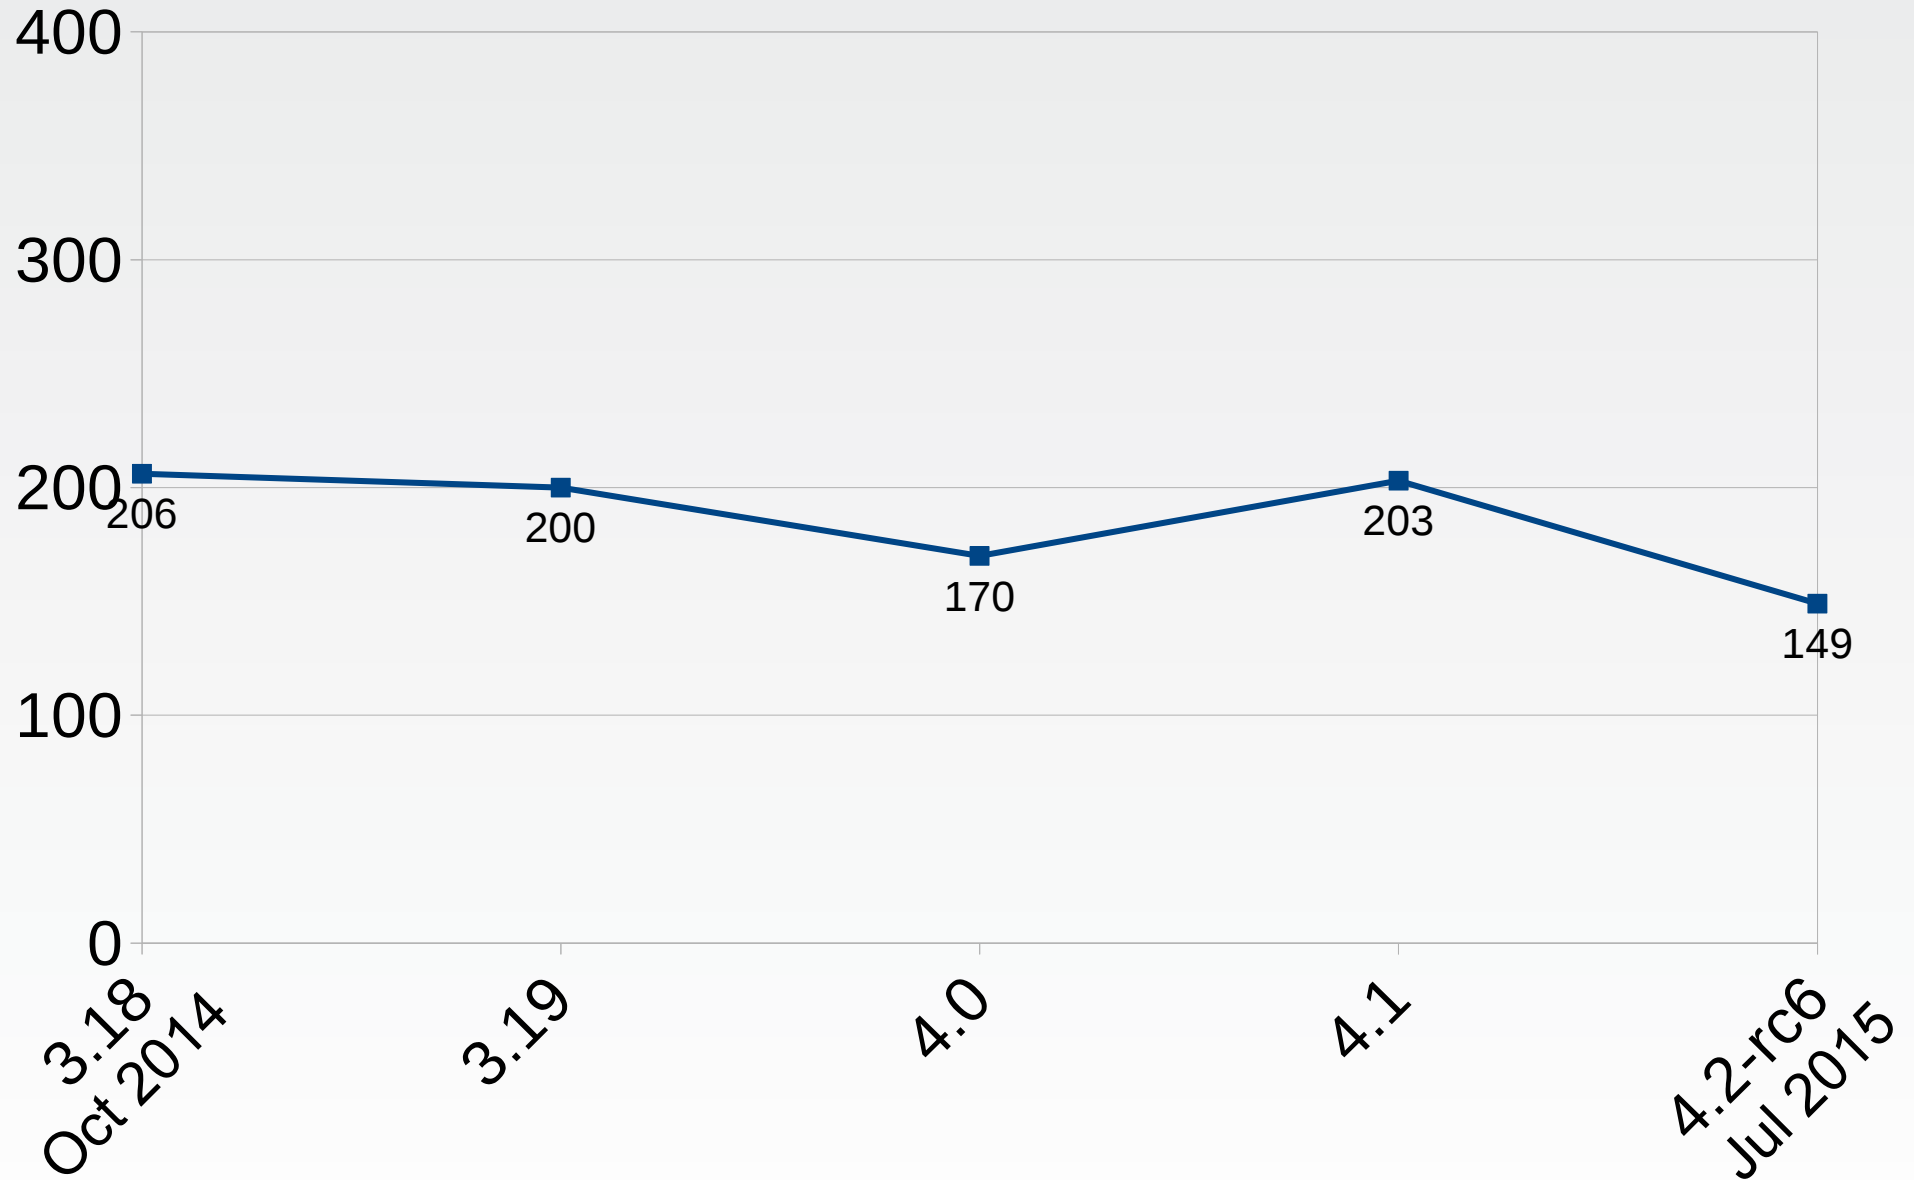
Paolo Bonzini  
Red Hat, Inc.  
KVM Forum 2015

# State of KVM

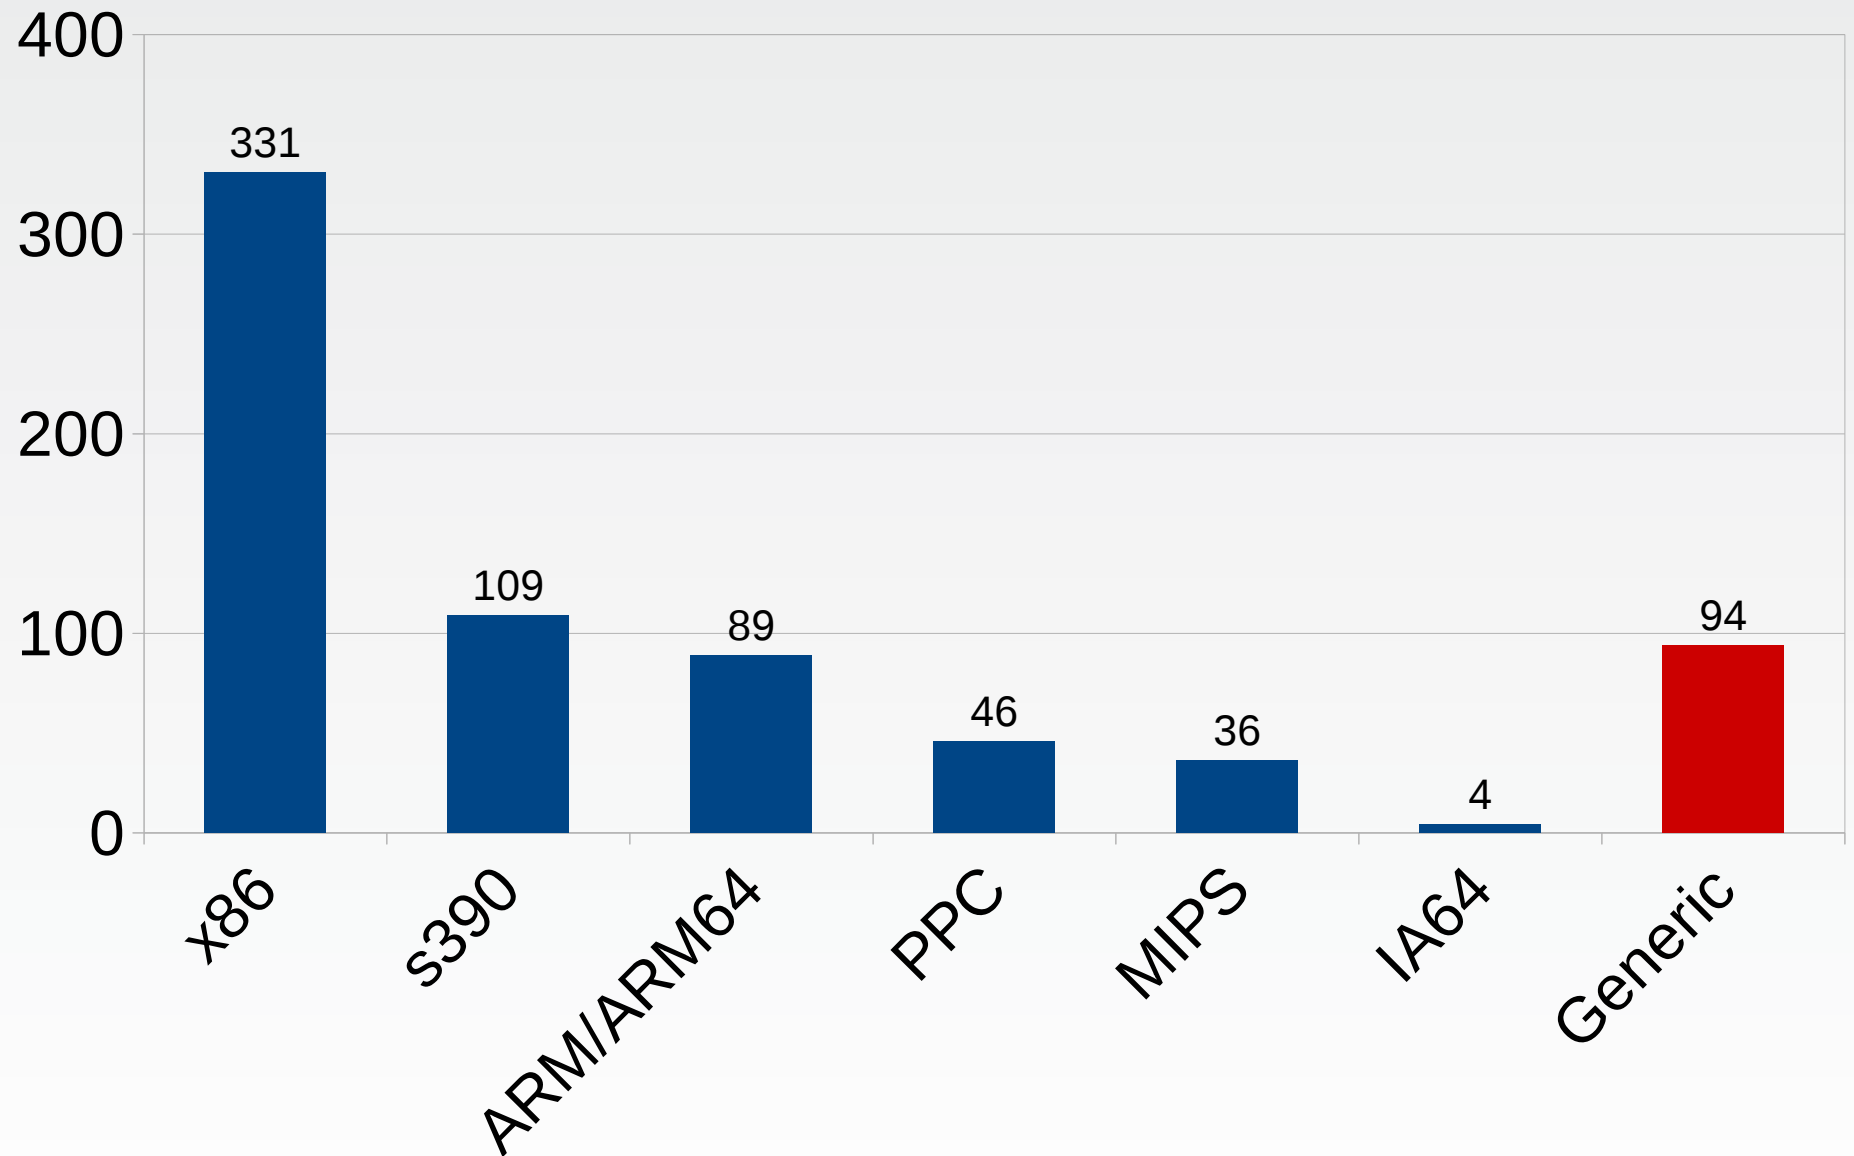
- 6 in-tree architectures
  - Active: ARM/ARM64, MIPS, PPC, x86, s390
  - ~~Dead: ia64~~
  - Out-of-tree: Tiler, MIPS hardware virt
- 2 main userspace implementations
  - All architectures ~~except ia64~~ supported in QEMU
  - ARM/ARM64, MIPS, PPC, x86 in kvmtool
- Since 3.18: 5 releases, 702 commits, ~25 companies



# Commits in each release



# Commits by architecture




# Main contributors

- Red Hat: maintenance, x86
- IBM: PPC, s390
- Linaro+ARM: ARM/ARM64
- Imagination Technologies: MIPS
- Intel: x86 hardware enablement & nested virt
- AMD, Google, Huawei, Samsung, Siemens, SuSE,...



# Commits by employer



# Highlights

- Spring cleaning: IA64, hardware-assisted virtualization on PPC970
- x86: real-time, nested APICv, PML, improved MTRR support, SMM, AMD PMU, XSAVES
- PPC: performance improvements, bug fixes
- s390: improved userspace access, SIMD, 2GB pages inside guests
- ARM: irqfd/ioeventfd (VFIO, vhost, dataplane)
- MIPS: FPU, SIMD



# Performance


- Still dominating SPECvirt, but we can do better!
- Low-latency hlt (halt\_poll\_ns)
- Eager FPU (x86 specific)
- Improved latency for real-time workloads





# KVM-RT latency jitter plot

cyclictest -m -n -N -q -v -p95 -h60 -i 200 -D 1h



# Realtime KVM patches (as of June 2015)

- sched: 45
- nohz: 12
- vmstat: 6
- timer: 4
- workqueue: 4
- **kvm: 4**
- **kernel-rt: 4**
- cpusets, isolcpus: 2
- irqbalance: 3
- libvirt: 6



# In the pipeline

- Performance improvements for guests spanning multiple NUMA nodes
- Split irqchip (x86, KVM Forum 2014)
- ARM debugging support
- ARM MSI support
- More kvm-unit-tests (ARM, PPC)



Have a great time!

