

# KVM Status Update

Oliver Upton


# Welcome!

# Statistics


## A Year(ish) in Review

- When we last met: 6.4-rc\*
- Now: 6.11, 6.12 merge window
- ~1700 commits
  - 90 cc'd to stable
  - ~20% selftests changes


## Contributions over time




## Contributions by Employer



## Contributions by Employer



## Contributions by area





# What's new

# Generic

- `guest_memfd` (!)
- FD-backed memslots
- Memory attributes
- Fixlets + cleanups

## arm64

- Even more nested code (traps, shadow MMU, CPU features)
- Configuration of CPU feature set via ID registers
- Normal-NC mappings for VFIO
- Host FF-A support in protected mode
- Performance improvements
  - MMU: Eager page splitting, TLBIRANGE
  - VGIC SGI and LPI routing

## x86

- Virtualizing new CPU features
- SEV-SNP support
- UAPI improvements for SEV/SEV-ES
- CPUID for max mappable GPA
- ioctl for pre-faulting secondary MMU
- A few hardware bugs

# RISC-V

- CPU features
- AIA irqchip
- New virtual SBI features, eg.
  - Stolen time
  - Debug console
  - PMU snapshot, overflow

Rest of the gang

Loongarch

- Architecture support
- PV IPI, steal time

powerpc

- 'NestedV2' ABI
- Bugfixes

s390

- PV crypto
- Secret UVCs

# What's next

## What's next

- Better support for non-refcounted pages
- More features based on guest\_memfd
  - mmap()
  - Hugepages
  - Unmapping from kernel
  - Test for poisoning
- Still interest in arch-independent MMU work?



# What's next

- “Passthrough” vPMU
  - perf/KVM interface across arches
- More TDX work
- Nested virtualization on arm64
- More pKVM upstreaming?

# Closing

# KVM Forum

- Grown a bit from last year!
- Back in Brno
- There's a social event :)
- Feedback for 2025

# Sponsors

## Platinum Sponsor



## Gold Sponsor



Google Cloud



**NVIDIA**®

# Organizers



## Local team:

- Jaroslav Suchanek
- Pavel Hrdina

## SFC:

- Rosanne DiMasio

Thanks!