

# What's coming in VIRTIO 1.2

New virtual I/O devices and  
features

Stefan Hajnoczi  
stefanha@redhat.com

# The VIRTIO standard

Open standard for virtual I/O devices

- ▶ Drivers available for Linux, BSD, macOS, Windows, and more
- ▶ Devices supported by multiple hypervisors and emulators
- ▶ Hardware implementations possible for many devices

New specification release every few years

- ▶ New features are often shipped in code once accepted into spec, you may already be using features from the upcoming spec release!

# New features in 1.2

## Core

- ▶ Per-virtqueue reset
- ▶ Shared memory resources
- ▶ Object resources for inter-device sharing
- ▶ Virtio-pci vendor-specific capabilities
- ▶ Virtio-pci queue\_notify\_data optimization

## virtio-blk

- ▶ **Secure erase**
- ▶ Lifetime metrics for eMMC-like media

## virtio-net

- ▶ **UDP Segmentation Offload**
- ▶ **Receive Side Scaling**
- ▶ **Per-packet hash reporting**
- ▶ Guest hdrln optimization
- ▶ Link speed and duplex reporting

## virtio-balloon

- ▶ Free page hinting & reporting
- ▶ Free page poisoning

## New devices in 1.2

net  
blk  
serial  
rng  
balloon  
scsi  
gpu  
input  
crypto  
vsock




iommu  
sound  
fs  
pmem  
mem  
rpmb  
scmi  
i2c  
gpio

Huge update with 9 new devices  
in the standard!

Even more device IDs reserved  
for future devices or  
non-standard applications.

We'll just look at the standard  
devices.

# sound




Sound card for multichannel PCM audio playback and capture with up to 384 kHz sample rate and 64-bit sample depth.

Reuses parts of High Definition Audio (HDA) spec.

Use cases:

- ▶ **Listening to music and movie/game audio.**
- ▶ **Voice calls.**
- ▶ **Pro audio/recording.**

## iommu




Simpler cross-arch VIRTIO alternative to emulated IOMMUs (Intel, AMD, ARM, etc).

Translates device memory accesses and enforces memory protection.

Use cases:

- ▶ **Linux VFIO apps (e.g. DPDK) inside guest.**
- ▶ **Isolating passthrough device DMA.**
- ▶ **Nested virtualization.**

# fs



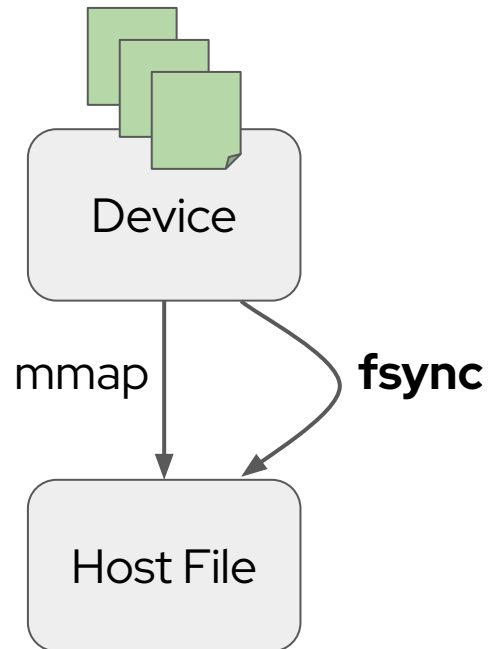
Shared host <-> guest file system with optional direct access to host page cache.

Based on Linux FUSE interface.

Use cases:

- ▶ **Secure container VMs like Kata Containers.**
- ▶ **Sharing source and build directories for development/testing in a VM.**
- ▶ **Providing files to VMs at install time.**

## pmem



Persistent memory (NVDIMM) device with flush mechanism.


Real NVDIMM devices may lack hypervisor-trappable mechanisms to flush written data to persistent storage.

Use cases:

- ▶ **Sharing host page cache to avoid copying data into all VMs.**
- ▶ **Persistent memory workloads where backing storage is not a real NVDIMM.**



# mem



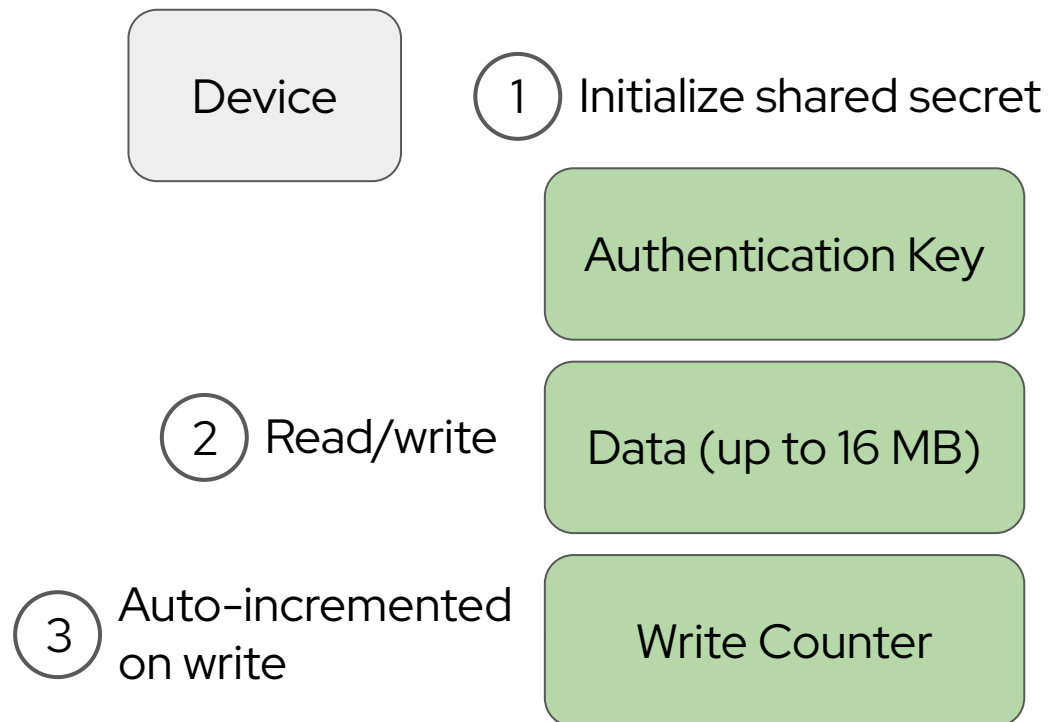
Cross-arch and NUMA-aware memory hotplug.

Supports fine-grained hotplug, not limited to arch-specific DIMM granularity.

Use cases:

- ▶ Adding/removing memory from guest at runtime.

# rpmb



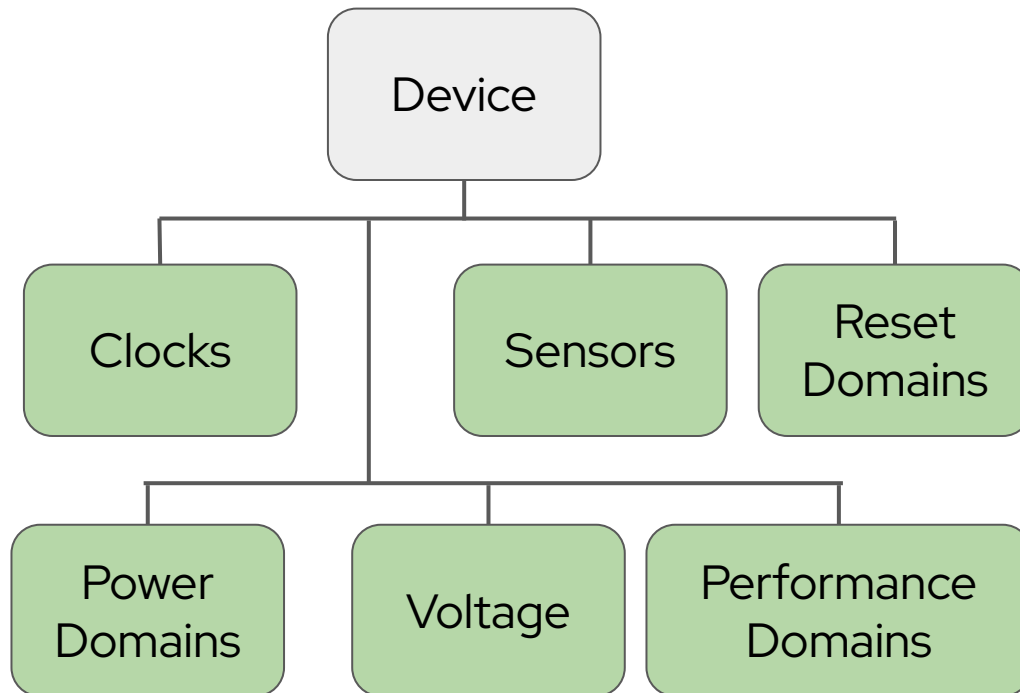
Replay Protected Memory Block device for tamper-resistant storage.

Based on feature from e.MMC and UFS flash devices.

Use cases:

- ▶ Protect system data (firmware or configuration) from modification.
- ▶ Storage for applications with sensitive data.

# scmi




Arm System Control and Management Interface device for power and performance management.

Use cases:

- ▶ Reusing bare metal Arm power and performance management code inside VMs.

## i2c



I2C adapter that gives access to emulated or passthrough I2C targets (EEPROMs, clocks, sensors, displays, etc).

Use cases:

- ▶ Emulation for applications that rely on I2C devices.
- ▶ Passthrough of physical I2C devices into VMs.

# gpio

Device

Line Name	Direction	Value
Power LED	Out	1
Reset Button	In	0
...	...	...

General Purpose Input/Output device with support for edge- and level-triggered interrupt lines.

Lines can have human-readable names.

Use cases:

- ▶ Interfacing external devices and logic signals with VMs.

## More about VIRTIO

Want to write a driver? Want to implement a device?

Discussion and questions:

[virtio-comment@lists.oasis-open.org](mailto:virtio-comment@lists.oasis-open.org)

Contribute to the specification:


<https://github.com/oasis-tcs/virtio-spec/blob/master/README.md>

# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 [linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

 [facebook.com/redhatinc](https://www.facebook.com/redhatinc)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

 [twitter.com/RedHat](https://twitter.com/RedHat)

