

## Call for Participation: OASIS Virtual I/O Device (VIRTIO) Technical Committee

Submitted by censign on Thu, 2013-06-20 21:47

**Type:**

Call for Participation

A new OASIS technical committee is being formed. The OASIS Virtual I/O Device (VIRTIO) Technical Committee has been proposed by the members of OASIS listed in the charter below. The TC name, statement of purpose, scope, list of deliverables, audience, IPR mode and language specified in the proposal will constitute the TC's official charter. Submissions of technology for consideration by the TC, and the beginning of technical discussions, may occur no sooner than the TC's first meeting.

The eligibility requirements for becoming a participant in the TC at the first meeting are:

(a) you must be an employee or designee of an OASIS member organization or an individual member of OASIS, and?

(b) you must join the Technical Committee, which members may do by using the Roster "join group" link on the TC's web page at [a].

To be considered a voting member at the first meeting:

(a) you must join the Technical Committee at least 7 days prior to the first meeting (on or before 23 July 2013); and?

(b) you must attend the first meeting of the TC, at the time and date fixed below (30 July 2013).

Participants also may join the TC at a later time. OASIS and the TC welcomes all interested parties.

Non-OASIS members who wish to participate may contact us about joining OASIS [b]. In addition, the public may access the information resources maintained for each TC: a mail list archive, document repository and public comments facility, which will be linked from the TC's public home page at [c].

Please feel free to forward this announcement to any other appropriate lists. OASIS is an open standards organization; we encourage your participation.

-----?

[a] <https://www.oasis-open.org/apps/org/workgroup/virtio/> [1]

[b] See <http://www.oasis-open.org/join/>?

[c] <http://www.oasis-open.org/committees/virtio/> [2]

=== TC Charter ===

The charter for this TC is as follows.

(1) Charter for OASIS Virtual I/O Device (VIRTIO) Technical Committee

(1)(a) TC Name

OASIS Virtual I/O Device (VIRTIO) Technical Committee

(1)(b) Statement of Purpose

Background:

Hardware virtualization allows multiple operating systems (?guests?) to share the same hardware (?host?) managed by host software (?hypervisor?).

These guests need networks, storage, consoles and similar but non-virtualization-aware standard devices cannot be shared, or guests may not be permitted to access host devices at all. The simplest solution is to emulate a device expected by the guest operating system, but this can be slow and/or complicated. As most operating systems have facilities for adding drivers for new physical hardware, we can use the same facilities to add drivers for devices which are easier and/or more efficient to implement in software.

As every hypervisor is different, they tend to implement hypervisor-specific devices, requiring every guest to support a new device for that environment. For example, Linux currently supports completely separate drivers for eight different virtualization platforms, with most drivers being sub-optimal. In 2007, an attempt was made to implement a hypervisor and OS-agnostic device model in Linux guests and the KVM hypervisor over the standard PCI bus. This is now also supported by the VirtualBox hypervisor (2010) and FreeBSD guests (2011).

A Draft Specification

In 2009, as interest accumulated, the ?Virtio PCI Card Specification? was published, with appendices for network, block storage and console devices. The emphasis was that virtual devices should be simple, look like driver authors expect physical devices to look, should be extensible, and that they should perform well.

Even at the time, there were implementations of virtio devices over non-PCI transports, and in 2011 the simplified ?mmio? transport was added as an Appendix, as well as a ?remote processor message? device which is actually used to communicate to a separate, physical CPU, rather than a virtual guest.

The years of experience have highlighted some of the implementation and design mistakes: enhancements have worked around many of them, but at cost of simplicity. Implementation bugs have also caused occasional anguish.

A 1.0 Specification

Our goal is to keep the good, discard the bad, and make the ugly optional. In particular, we will try not to break too much, and ensure it's possible for devices to support both 1.0 compliant and legacy guest drivers.

#### (1)(c) Scope

The ?Virtio PCI Card Specification? 0.9.5, minus the Appendix H, shall be used as a starting point (referred to as ?legacy?). Note that we expect the final output of this TC to be incompatible with that specification, though it will be possible for virtual devices to support both legacy and modern guest drivers.

The TC will produce an OASIS Standard by refining and documenting existing implementations and practice. After publication of the initial OASIS Standard, the TC may choose to develop a Version 2.0 standard that builds on lessons learned and identified trends and that may result in a significantly different architecture and approach. Until then, the TC shall not throw out the baby with the bathwater, boil the ocean, or embark on a PhD research topic.

The TC will consider one or more buses for virtual devices, including PCI. It will consider various kinds of devices, including network devices. Each part of the OASIS Standard will be considered in terms of portability, simplicity, least-surprise for driver authors, extensibility and performance. In particular, the effect of future radical extensions (such as layout changes) will be considered.

#### (1)(d) Deliverables

1. Specification of feature negotiation, configuration and queues, from both driver and device points of view.
  - 1a. Specifically for virtio over PCI.
  - 1b. Specifically for virtio over mmio.
  - 1c. Specifically for other transports as decided by TC.
2. Specification of device-specific configuration.
  - 2a. For network devices.
  - 2b. For block storage devices.
  - 2c. For entropy devices.
  - 2d. For console devices.
  - 2e. For memory ballooning devices.
  - 2f. For SCSI endpoint devices.
  - 2e. Non-normative advice for designing new device types.
  - 2f. List of reserved device numbers for non-specified device types. (This section is likely to see ongoing maintenance before 2.0).
  - 2g. List of reserved feature numbers for non-specified features for each device type. (This section is likely to see ongoing maintenance before 2.0).
3. Non-normative code examples for operation of guest/host side of buffers.
4. Non-normative guide for creating devices which also support previous mode(s).

We anticipate deliverables (1) and (2) to be a single work, and (3) and (4) to be separate works. The projected delivery dates are 12 to 16 months after the first meeting.

#### (1)(e) IPR Mode

## Non-Assertion Mode

### (1)(f) Audience

Developers of hypervisors and device driver authors.

### (1)(g) Language

US English.

## Section 2: Additional Information

### (2)(a) Identification of Similar Work

<http://www.pcisig.com/specifications/iov/> [3]

The PCI SIG I/O Virtualization Specifications allow a physical host device to be shared by multiple guests. This is different from a virtual device, which doesn't have any physical existence.

VMWare have a `?vmxnet?` device and Microsoft have `?vmbus?` devices for which there are open source drivers, but I'm unaware of any open specification (particularly from the device side, which is implemented within their proprietary products).

### (2)(b) First TC Meeting

IBM sponsored conference call:

UTC: Tuesday, 30 July 2013 at 16:00:00

Adelaide: Wed 1:30 AM

Austin: Tue 11:00 AM

London: Tue 5:00 PM

San Francisco: Tue 9:00 AM

Callin numbers (more available on request):

USA/Canada: 1-719-325-2507

Sydney: +61 (0) 2 8524 4805

Israel (toll free): 1809 213 262

India (toll free) : 000 800 100 8486

UK, London : +44 (0) 20 3481 0142

PASSCODE: 804 295 2715

### (2)(c) Ongoing Meeting Schedule

2-weekly, same concall number. Expect most work done via ml, this only needed for status and votes.

### (2)(d) TC Proposers

Rusty Russell (IBM)

Michael S. Tsirkin (Red Hat)

Pawel Moll (ARM)

Anthony Liguori (IBM)

Sasha Levin (Oracle)

Stefan Hajnoczi (Red Hat)

Amit Shah (Red Hat)

(2)(e) Primary Representatives' Support

Mark Little : Thu, 11 Apr 2013 12:42:42 +0100

As Red Hat's Primary Representative to OASIS, I would like to state that Red Hat approves the creation of the virtio TC Charter. We will also be actively involved in the TC once created and we endorse all Red Hat proposers.

Dave Ings : Thu, 11 Apr 2013 09:49:23 -0400

IBM approves our participation as co-proposers for this TC.

Martin Chapman Fri, 03 May 2013 05:03:50 -0700

As Oracle's Primary Representative to OASIS, I approve the VIRTIO TC Charter, and endorse all Oracle proposes listed in (2)(d).

Ian Ampleford Fri, 10 May 2013 03:44:35 +0100

As ARM's Primary Representative for OASIS, I approve Pawel as an initiating member of the VirtIO Technical Committee.

(2)(f) TC Convenor

Dave Ings (IBM)

(2)(g) OASIS Member Section

None.

(2)(h) Anticipated Contributions

<https://github.com/rustyruessell/virtio-spec/tree/v0.9.5> [4] (minus Appendix H)

(2)(i) FAQ Document

None at this stage.

(2)(j) Work Product Titles and Acronyms

Virtio Device Specification (VIRTIO).

**Associated TC:**

Virtual I/O Device (VIRTIO)

**Deadline:**

Thu, 2013-06-20

---

**Links:**

[1] <https://www.oasis-open.org/apps/org/workgroup/virtio/>

[2] <http://www.oasis-open.org/committees/virtio/>

[3] <http://www.pcisig.com/specifications/iov/>

[4] <https://github.com/rustyruessell/virtio-spec/tree/v0.9.5>