DPDK

DATA PLANE DEVELOPMENT KIT

THE LINUX FOUNDATION

DPDK | IT大咖说 知识分享平台

DPDK SUMMIT CHINA 2017

主办方： intel

参与方： 腾讯云 ZTE 美团云 Panabit® 太一星晨 Balance Your Networks UnitedStack有云 云杉网络 Yunshan Networks

协办方： SDNLAB 专注网络创新技术

视频支持方： IT大咖说 知识分享平台

DPDK

DATA PLANE DEVELOPMENT KIT

THE LINUX FOUNDATION

DPDK | IT大咖说 知识分享平台

Accelerate VM IO via SPDK Vhost Solution

Changpeng Liu, Intel

主办方： intel

参与方： 腾讯云 ZTE 美团云 Panabit® 太一星晨 Balance Your Networks UnitedStack 云杉网络 Yunshan Networks

协办方： SDNLAB 专注网络创新技术

视频支持方： IT大咖说 知识分享平台

LEGAL DISCLAIMER

- No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.
- Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.
- This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.
- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.
- © 2017 Intel Corporation. Intel, the Intel logo, Intel. Experience What's Inside, and the Intel. Experience What's Inside logo are trademarks of Intel. Corporation in the U.S. and/or other countries.
- *Other names and brands may be claimed as the property of others.
-

Agenda

- Introduction
- SPDK Vhost Architecture
- Usage Cases
- Benchmarks
- Plans

Introduction

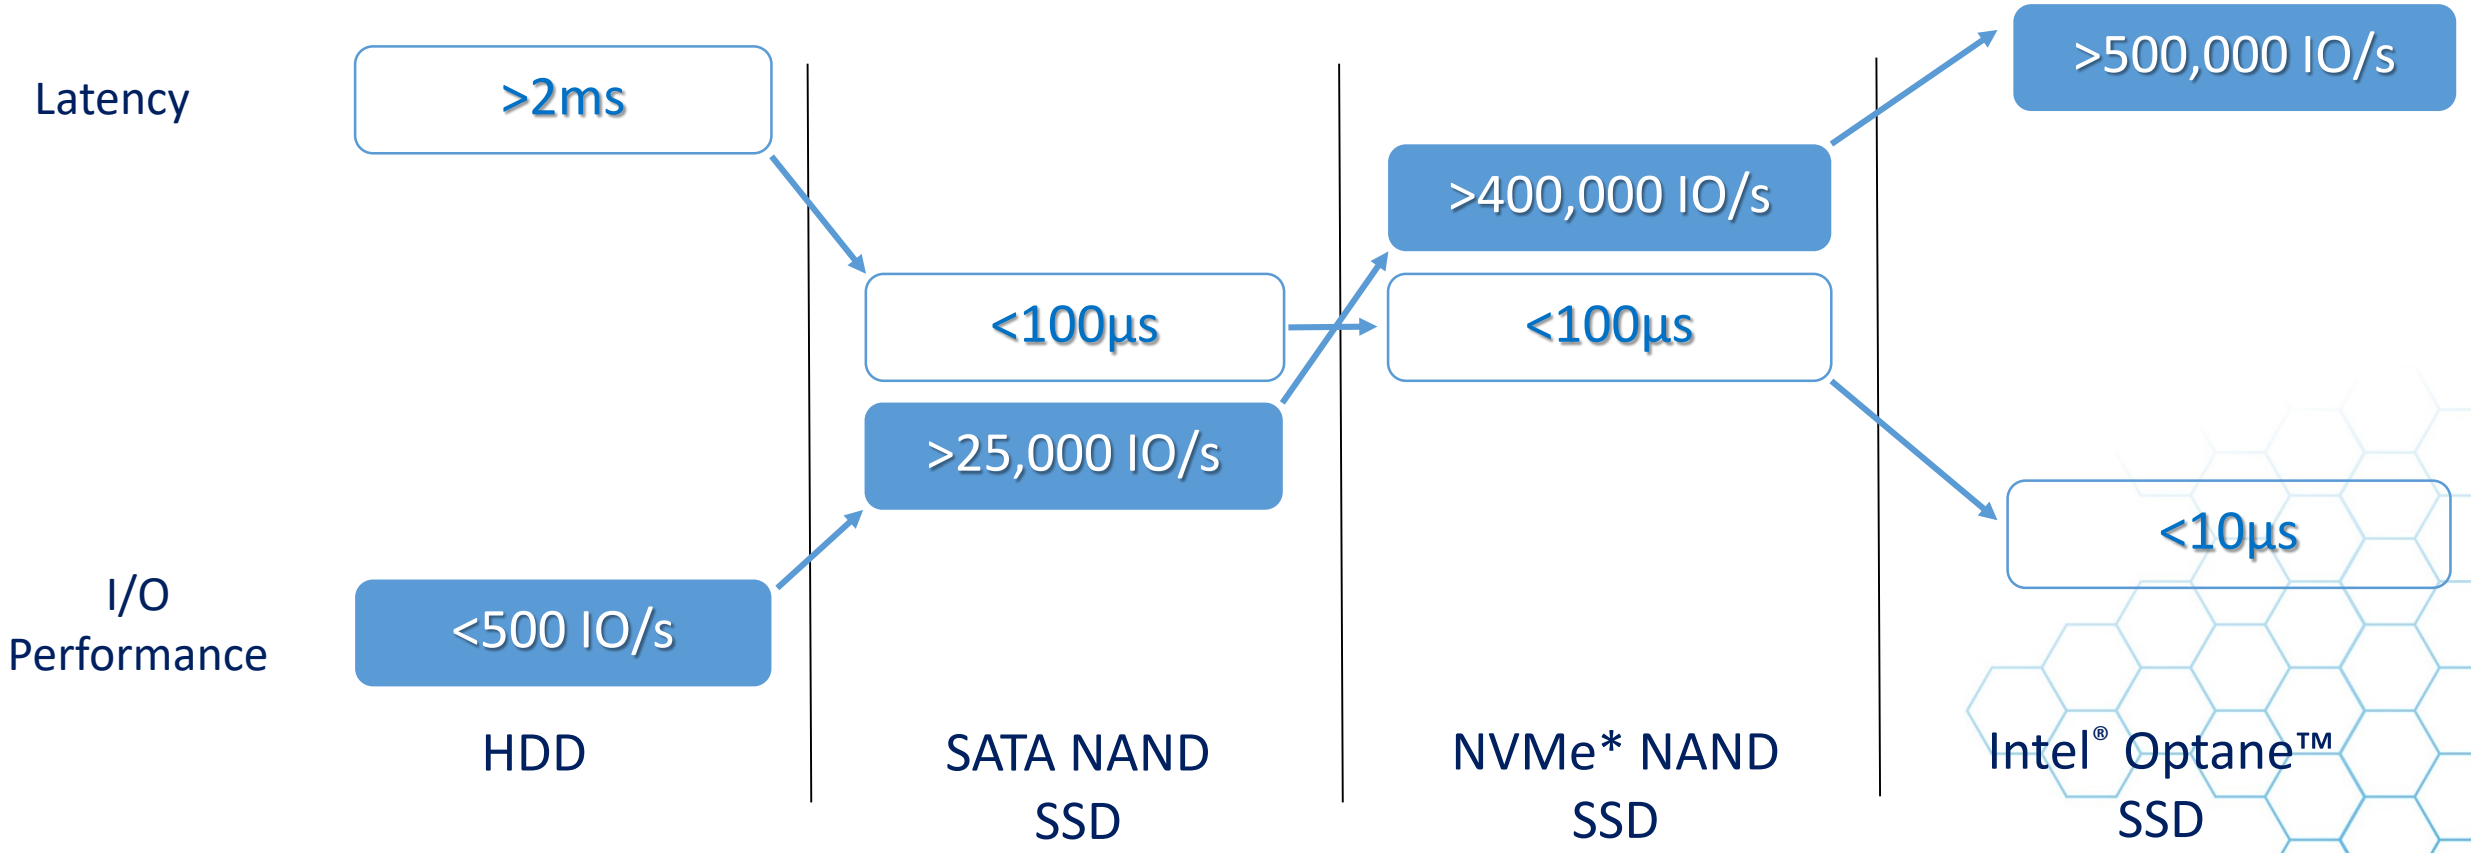

The Opportunity: Use Intel software ingredients to unlock the potential of new media

SPDK Architecture

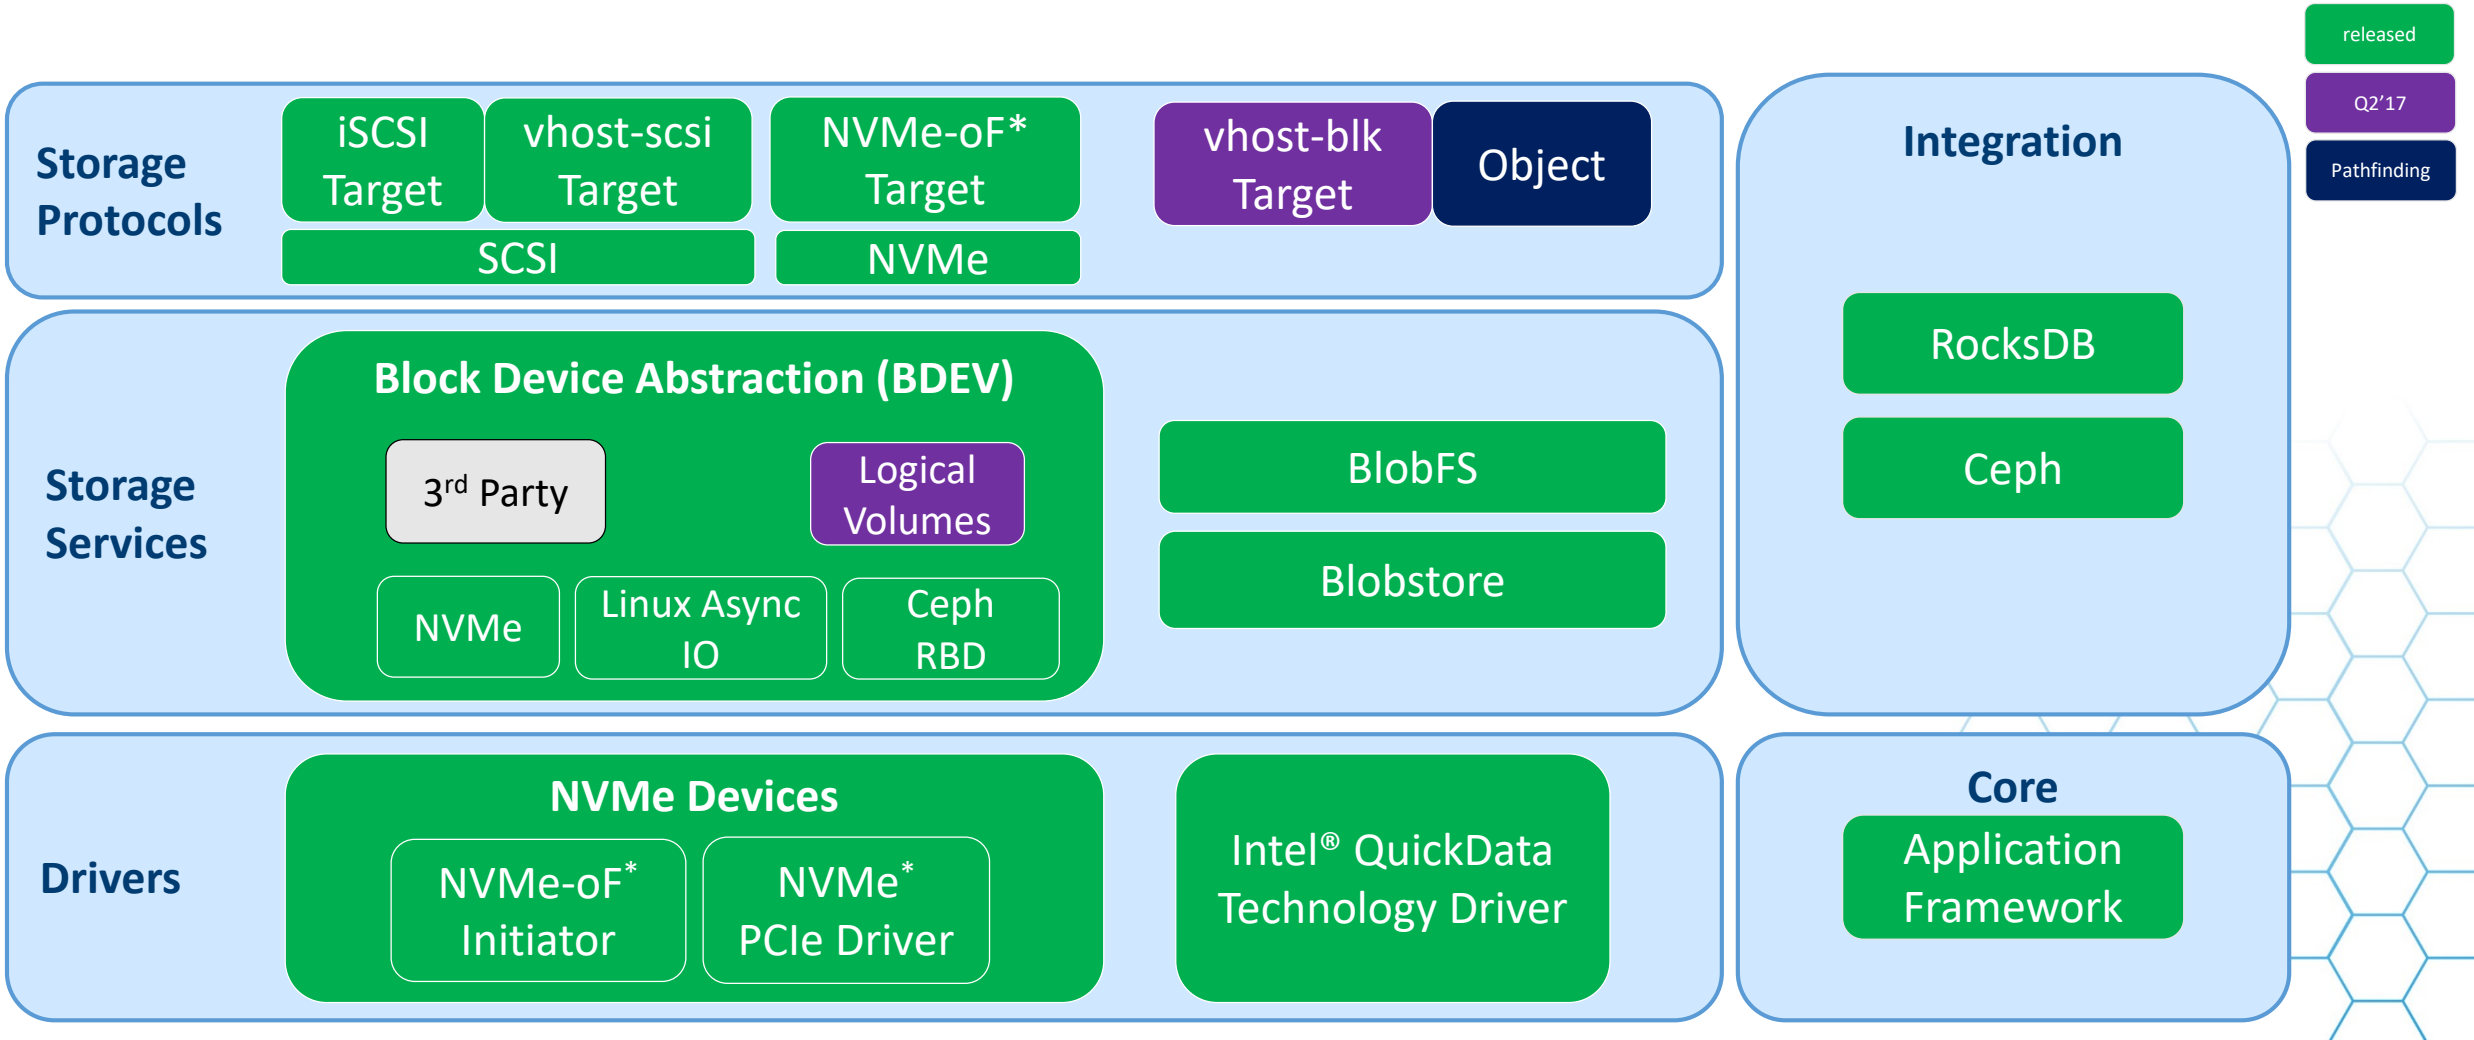

SPDK VHOST Architecture

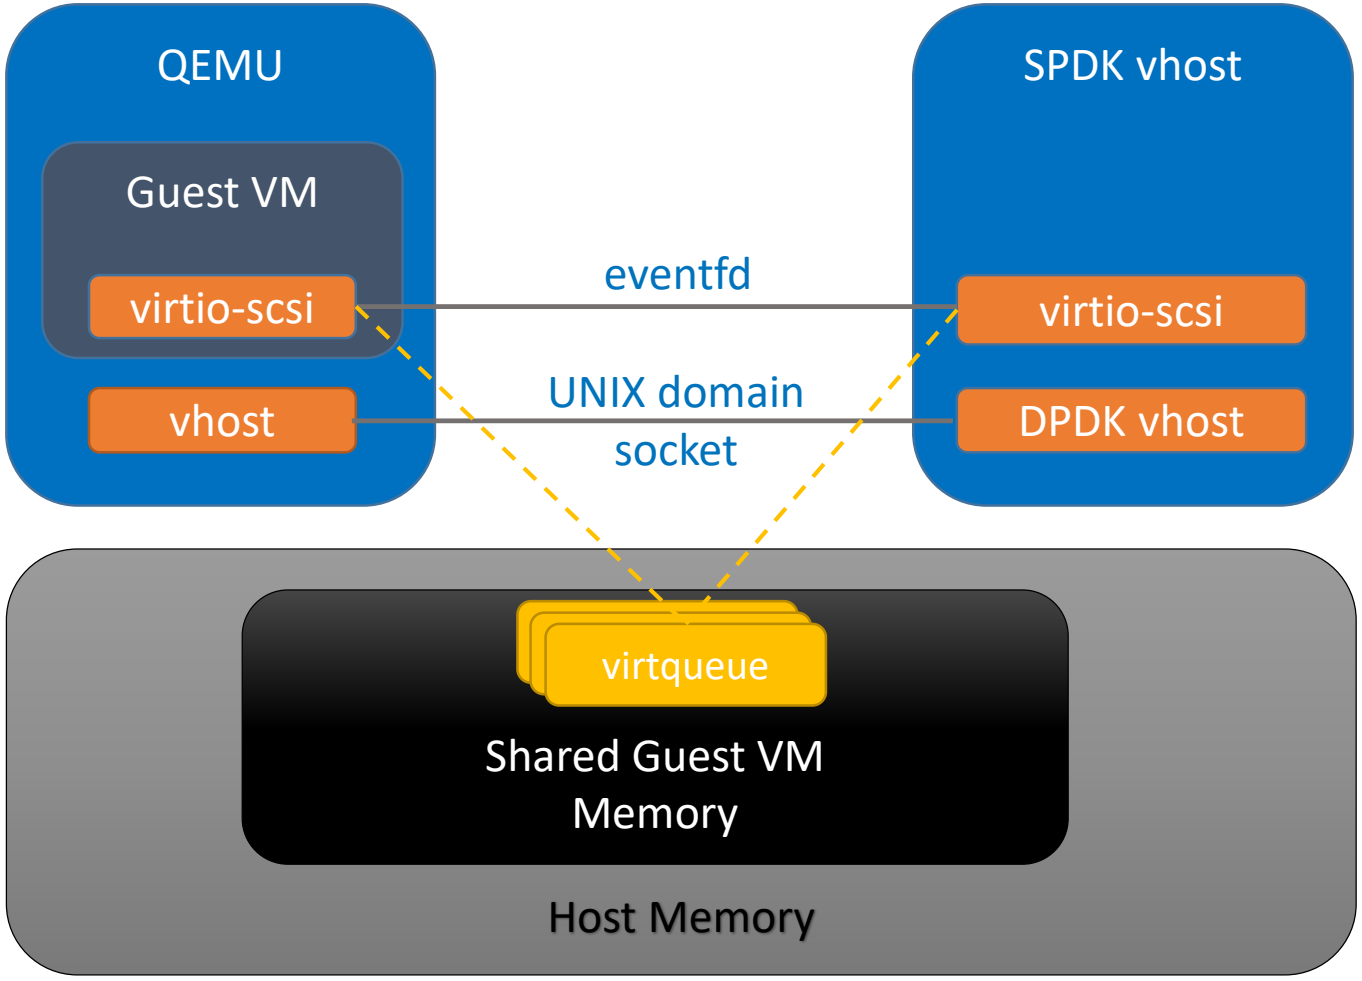

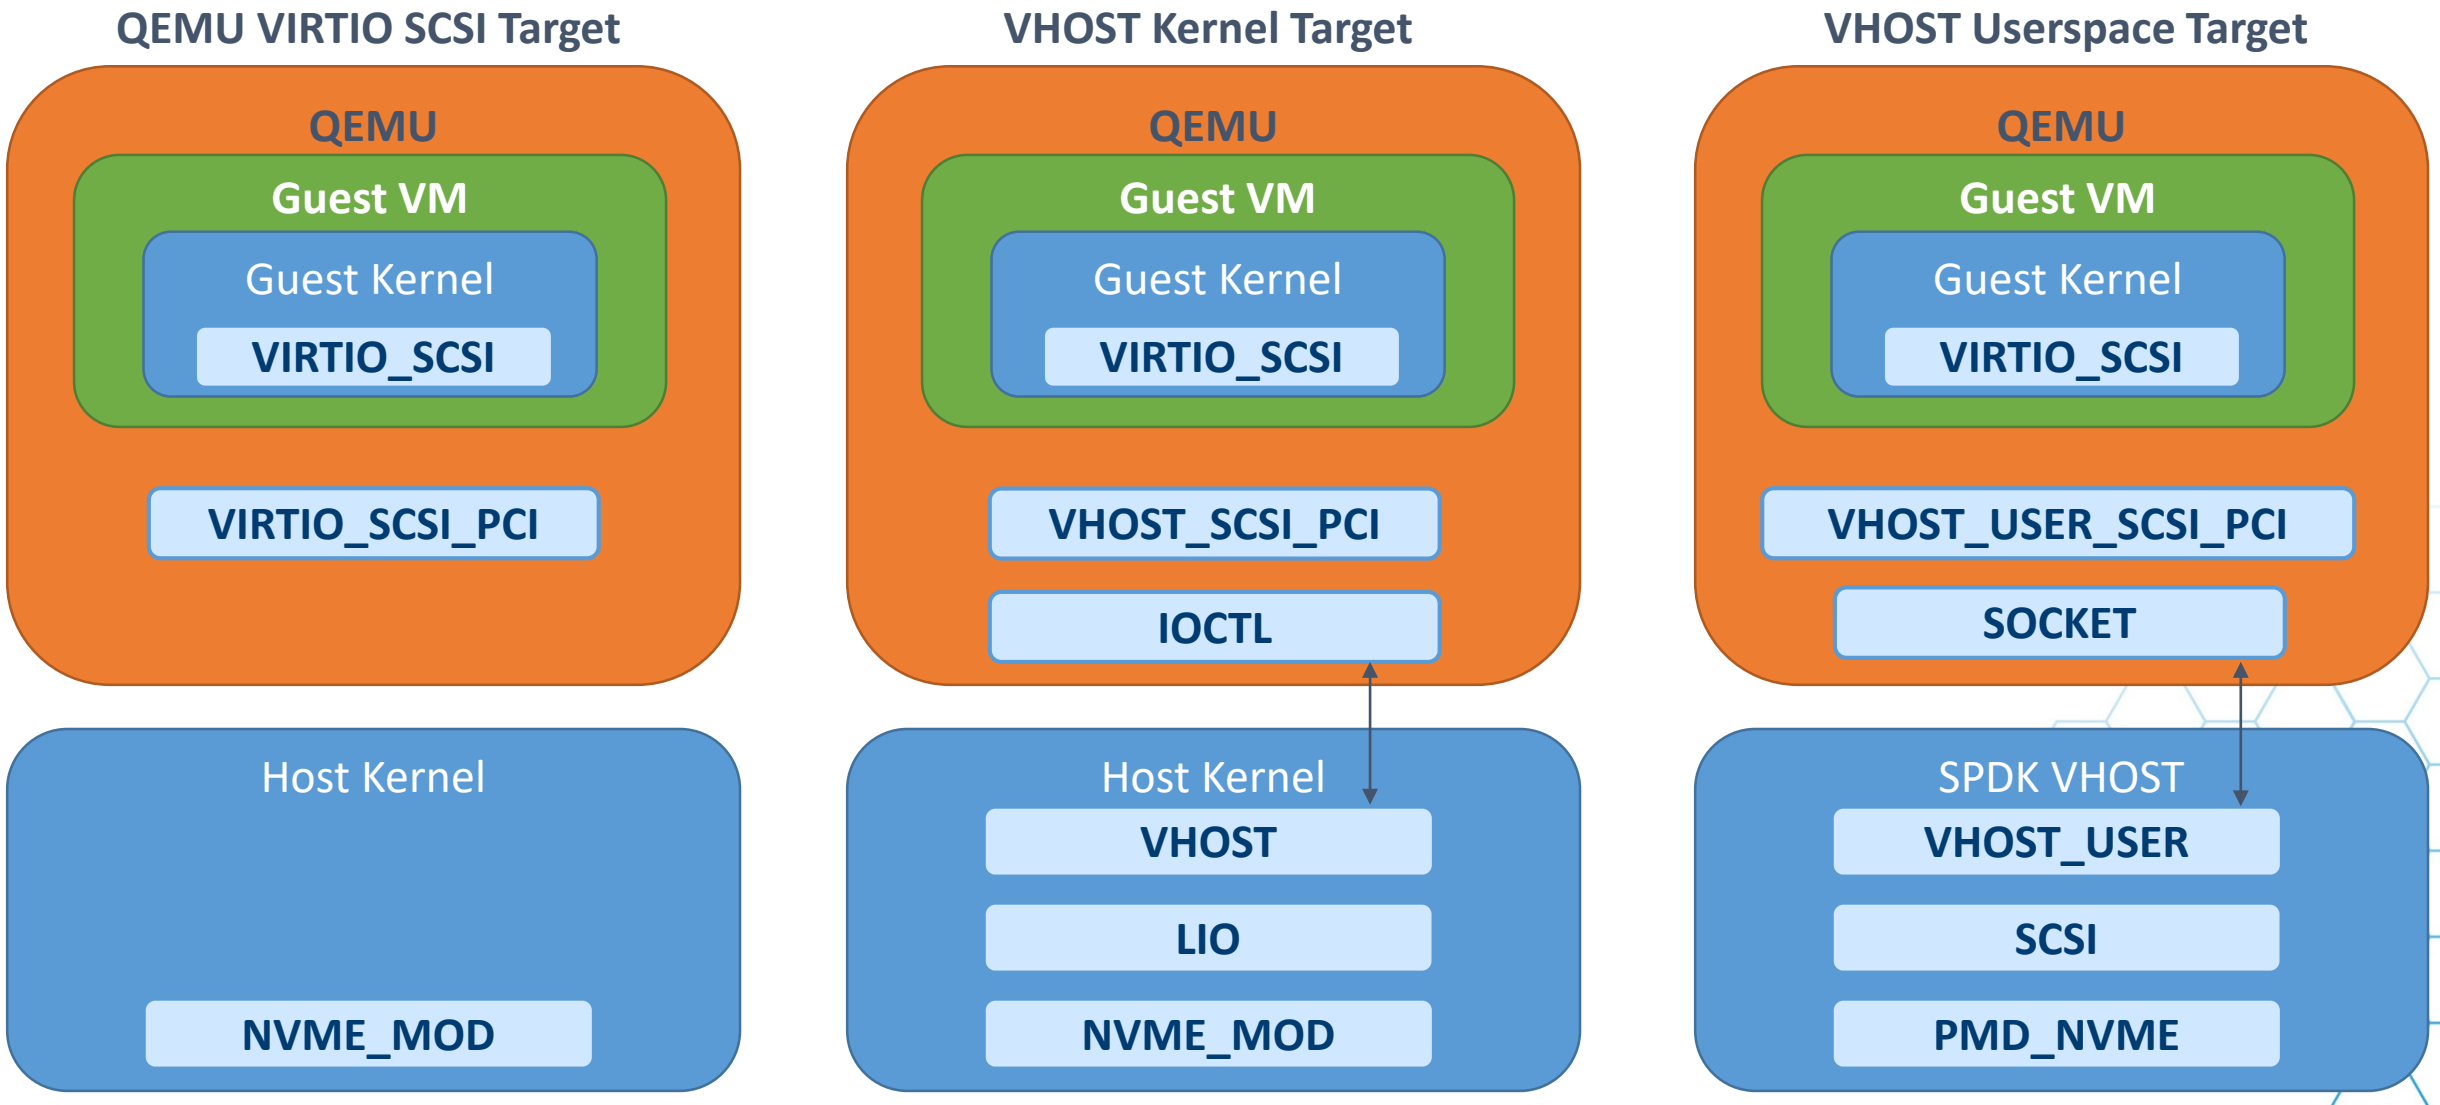
QEMU VIRTIO SCSI Target
QEMU
Guest VM
Guest Kernel
VIRTIO_SCSI
VIRTIO_SCSI_PCI
Host Kernel
NVME_MOD
VHOST Kernel Target
QEMU
Guest VM
Guest Kernel
VIRTIO_SCSI
VHOST_SCSI_PCI
IOCTL
Host Kernel
VHOST
LIO
NVME_MOD
VHOST Userspace Target
QEMU
Guest VM
Guest Kernel
VIRTIO_SCSI
VHOST_USER_SCSI_PCI
SOCKET
SPDK VHOST
VHOST_USER
SCSI
PMD_NVME

VM Ephemeral Storage

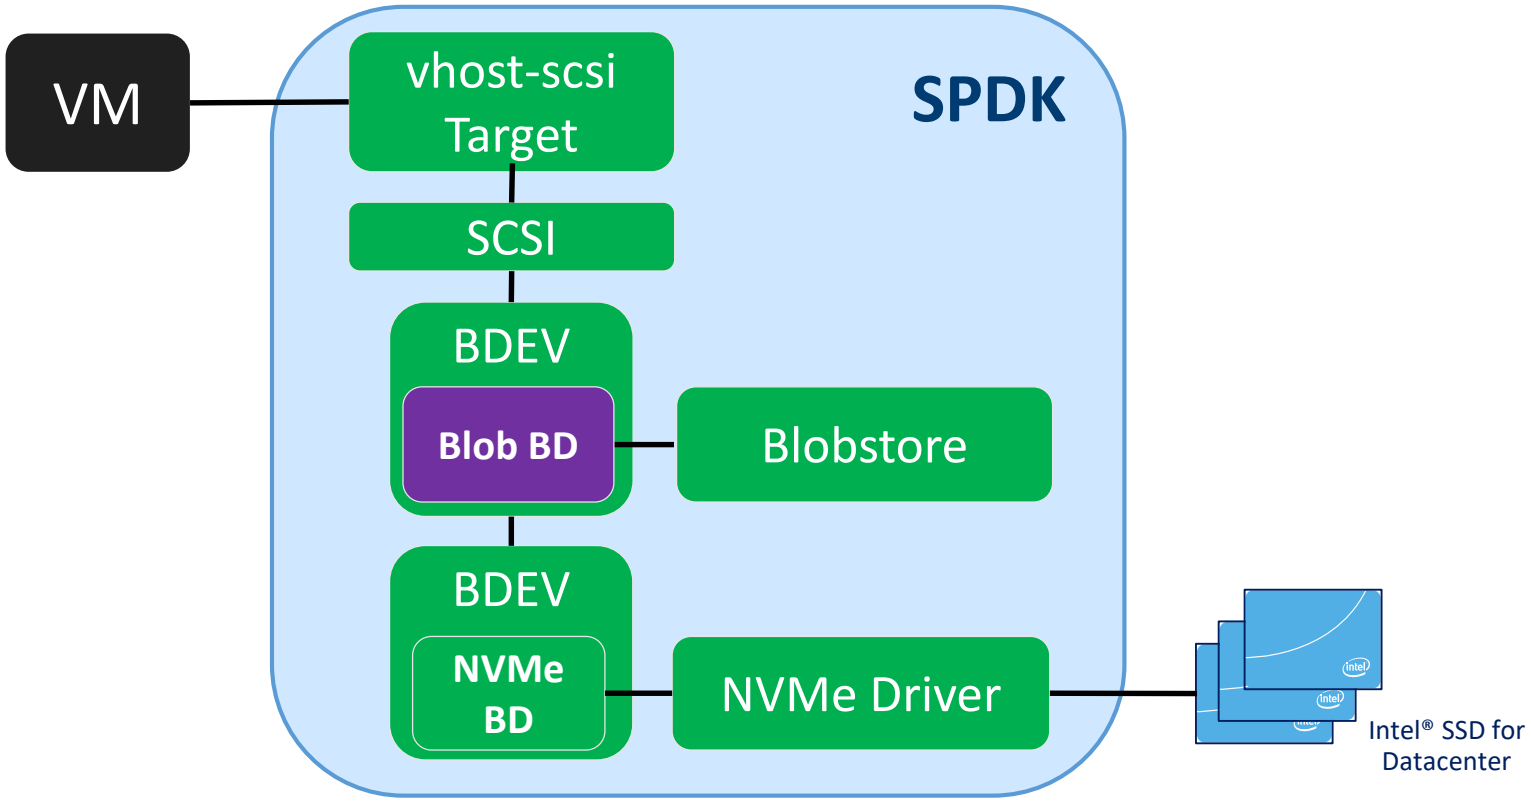

Released
Q2'17

- Improves Storage Virtualization
- Works with KVM/QEMU
- 6x efficiency vs. kernel vhost
- 10x efficiency vs. QEMU virtio
- Increased VM density

VM Remote Storage

Released

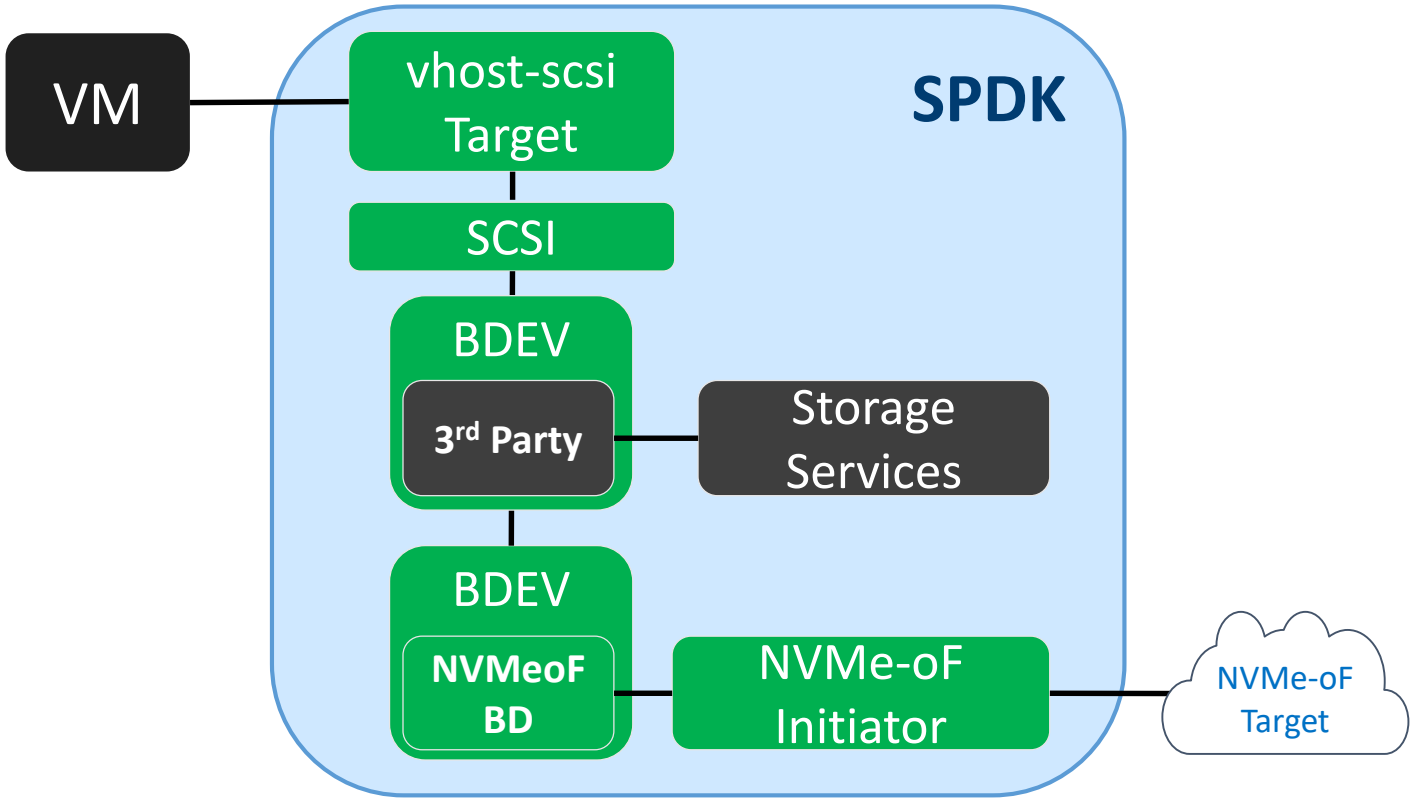

- Enable disaggregation and migration of VMs using remote storage
- Improves Storage Virtualization & Flexibility
- Works with KVM/QEMU

Benchmarks

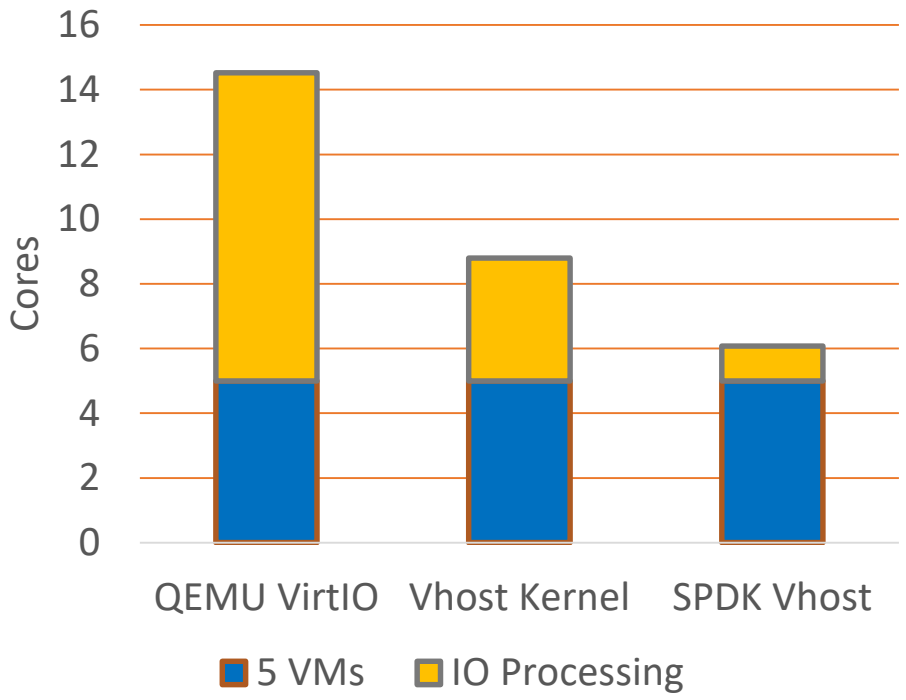

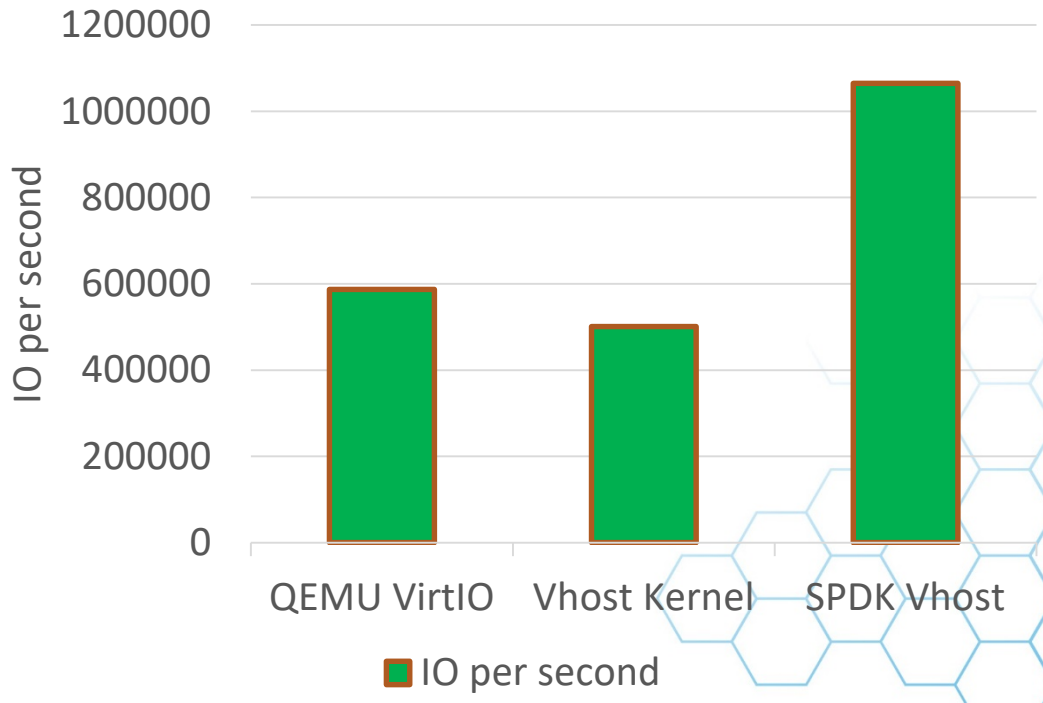

System configuration: 44x Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz (HT off); Cores per socket: 22; 8x Samsung 8GB DDR4 @2400 12x Intel SSD DC P3700 Series 1,5T @ FW 8DV101H0 DPDK: 17.02; Host Dist/Kernel: Fedora 25/Kernel 4.8.15-300; Guest Dist/Kernel: Ubuntu 16.04/Kernel 4.4.0-59-generic, mq enabled; Fio ver: fio-2.2.10; Fio workload: blocksize=4k, iodepth=512, iodepth_batch=128, iodepth_low=256, ioengine=libaio, size=10G, ramp_time=10, group_reporting, thread, numjobs=1, direct=1, rw=randread

Plans

- VFIO Support
- Support for vhost-blk protocol
- Live migration
- Performance tuning, including
  - -multiqueue
  - -completion event coalescing

Accelerate Crypto Service by DPDK vhost

Xin Zeng, Intel

主办方：

参与方：

协办方：

视频支持方：

Agenda

- Virtio Crypto Device Introduction
- Boost SSL/TLS Service by virtio-crypto
- DPDK vhost-user for virtio-crypto
- Plans
- Summary

Virtio Crypto Device

- A virtual cryptography device under virtio device framework
- Provides an set of operation interfaces for different cryptography services
- Mainly contributed by Huawei & Intel in community

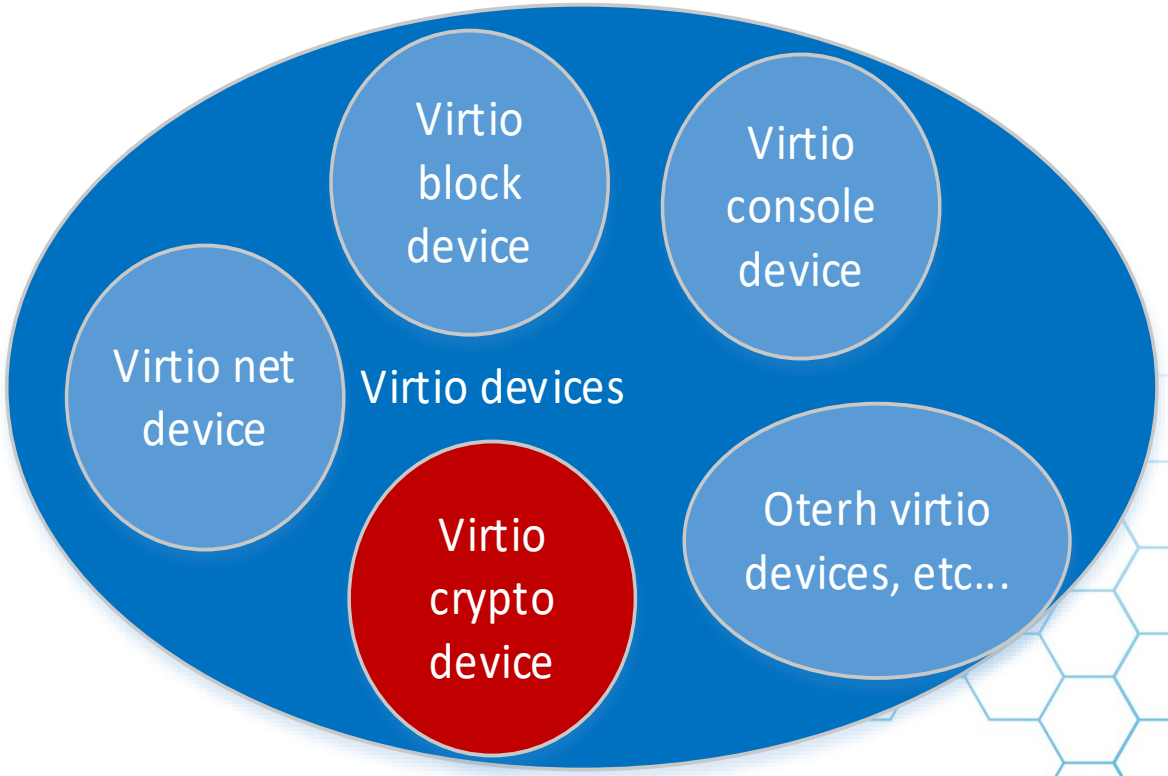

Boost SSL/TLS Service by virtio-crypto

- Motivation
  - Unified Driver in the Guest
  - Accelerator as a service for better performance
  - Friendly Cloud Characteristic
- PoC Workload
  - Nginx HTTPS Web Server
  - RSA2K session establishment
- Ingredients
  - virtio-crypto PMD
  - vhost-user for Crypto
  - Intel® QAT DH895XCC device driver in Linux
- Performance
  - ~4.5x throughput (TLS connection per second) compared to software solution

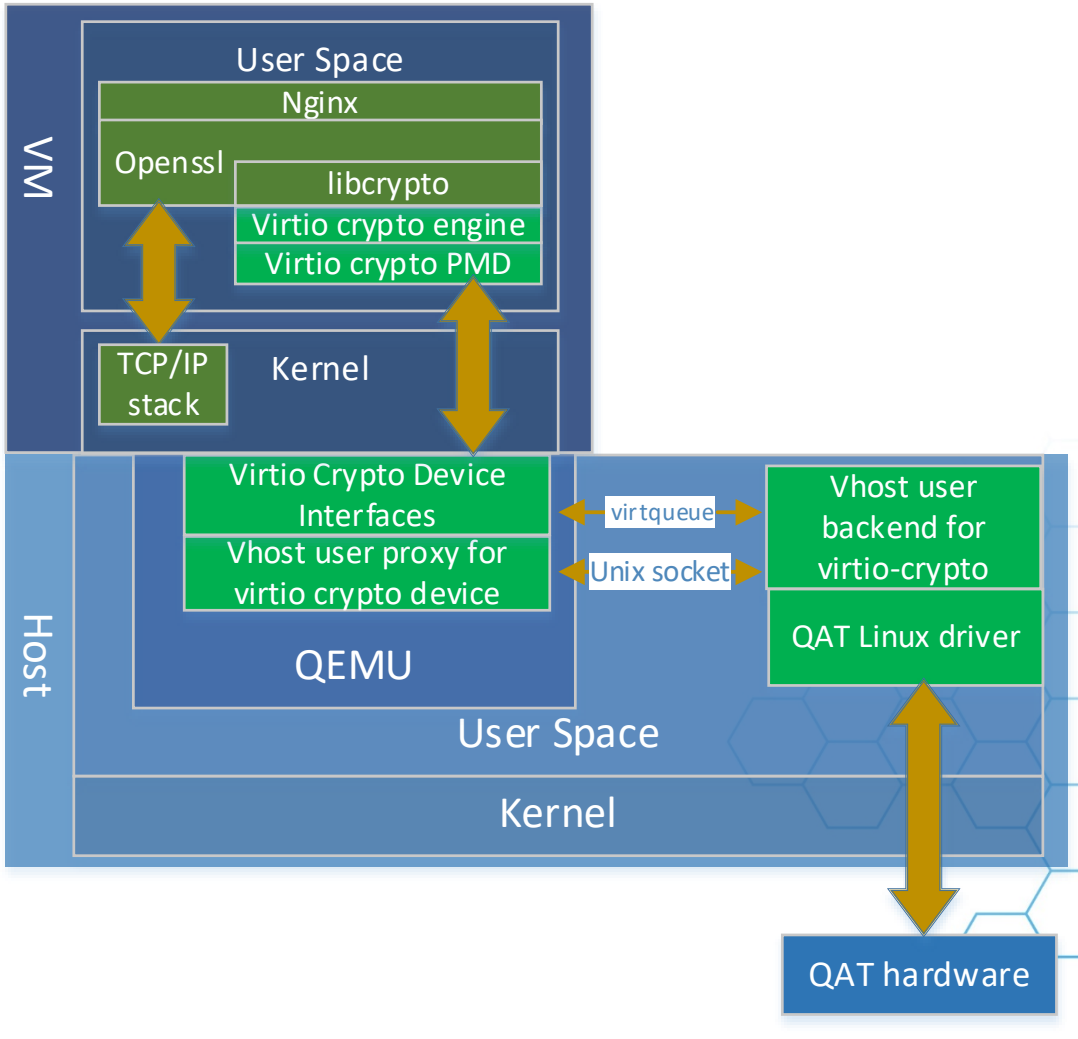

DPDK vhost-user for virtio-crypto

- viritio-crypto in VM
  - Crypto appliance
  - Under LKCF framework
  - virtio Crypto PMD
- New vhost proxy in QEMU
- virtio-crypto backend in Host
  - Build vhost user crypto target on top of DPDK generic vhost lib
  - Connect with DPDK crypto device

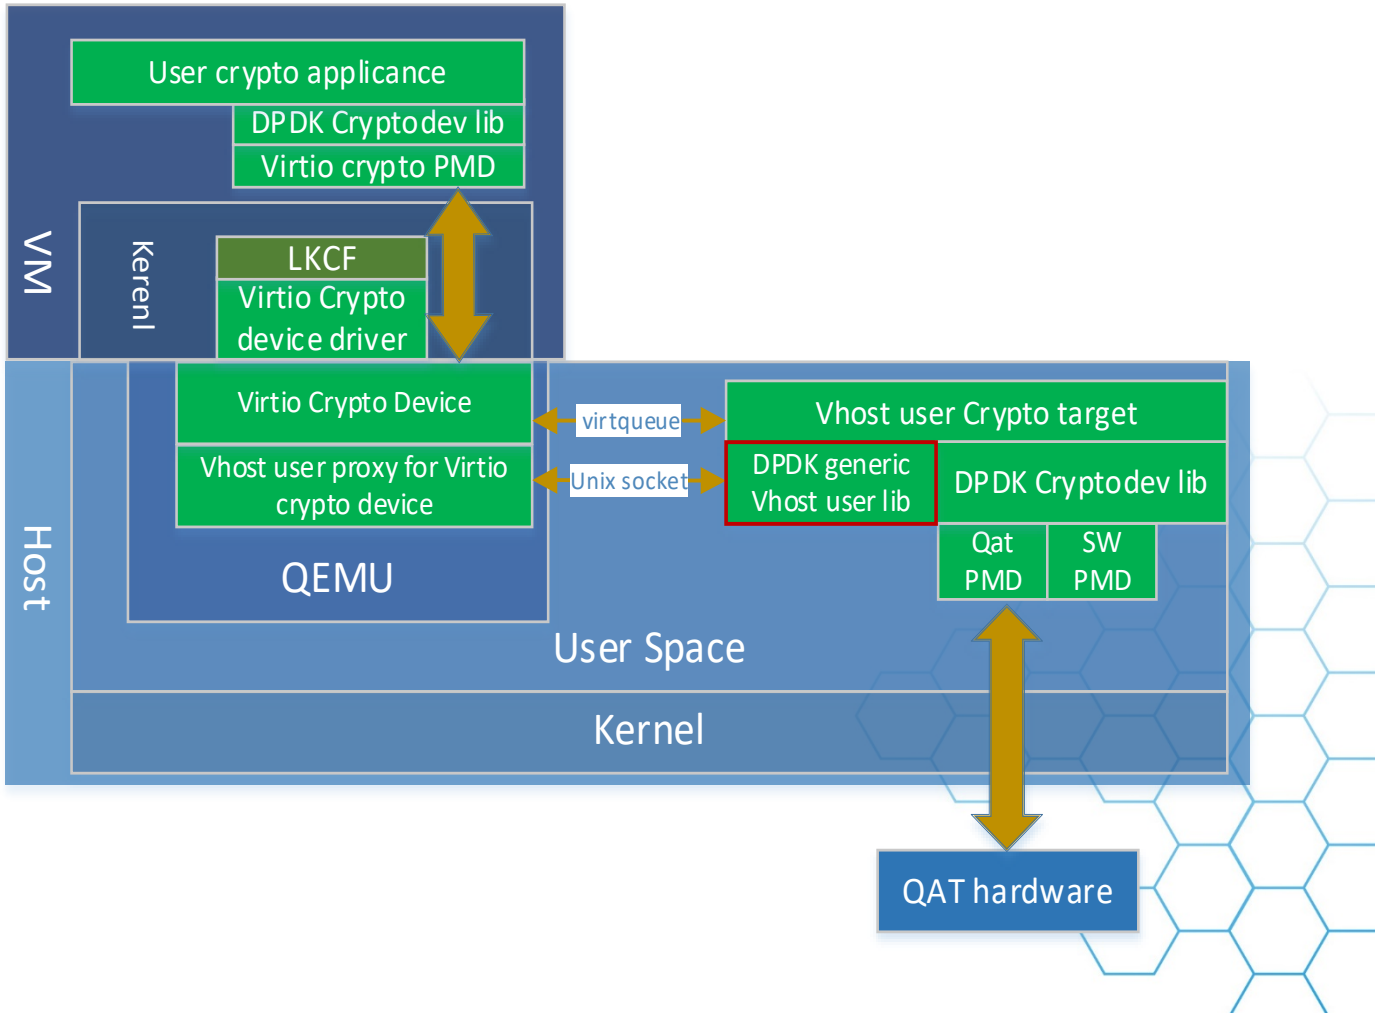

Intel® QAT Overview

- A hardware-based acceleration technology
- Accelerate compute-intensive security and compression operations
- For more details of Intel® QAT, visit here

WIP and Plans

- New device type (virtio-crypto) proposal in virtio spec. v1.1
- Upstream vhost user for virtio-crypto in DPDK community
- Live migration support
- Multi-queue support
- Performance optimization

Acknowledgement

- arei.gonglei@huawei.com
- Liang Ma
- John Griffin
- Brian Keating
- Jacqueline Jardim
- Cunming Liang

Summary

- DPDK generic vhost user library is ready (available in DPDK 17.05)
- vhost user for SCSI and Crypto devices are ongoing.
- Benefits from DPDK vhost library
  - Why Reinvent Wheel?
  - General APIs to build vhost user application
  - Leverage fast I/O capacity by DPDK PMD
  - High Performance
- Welcome contributions!

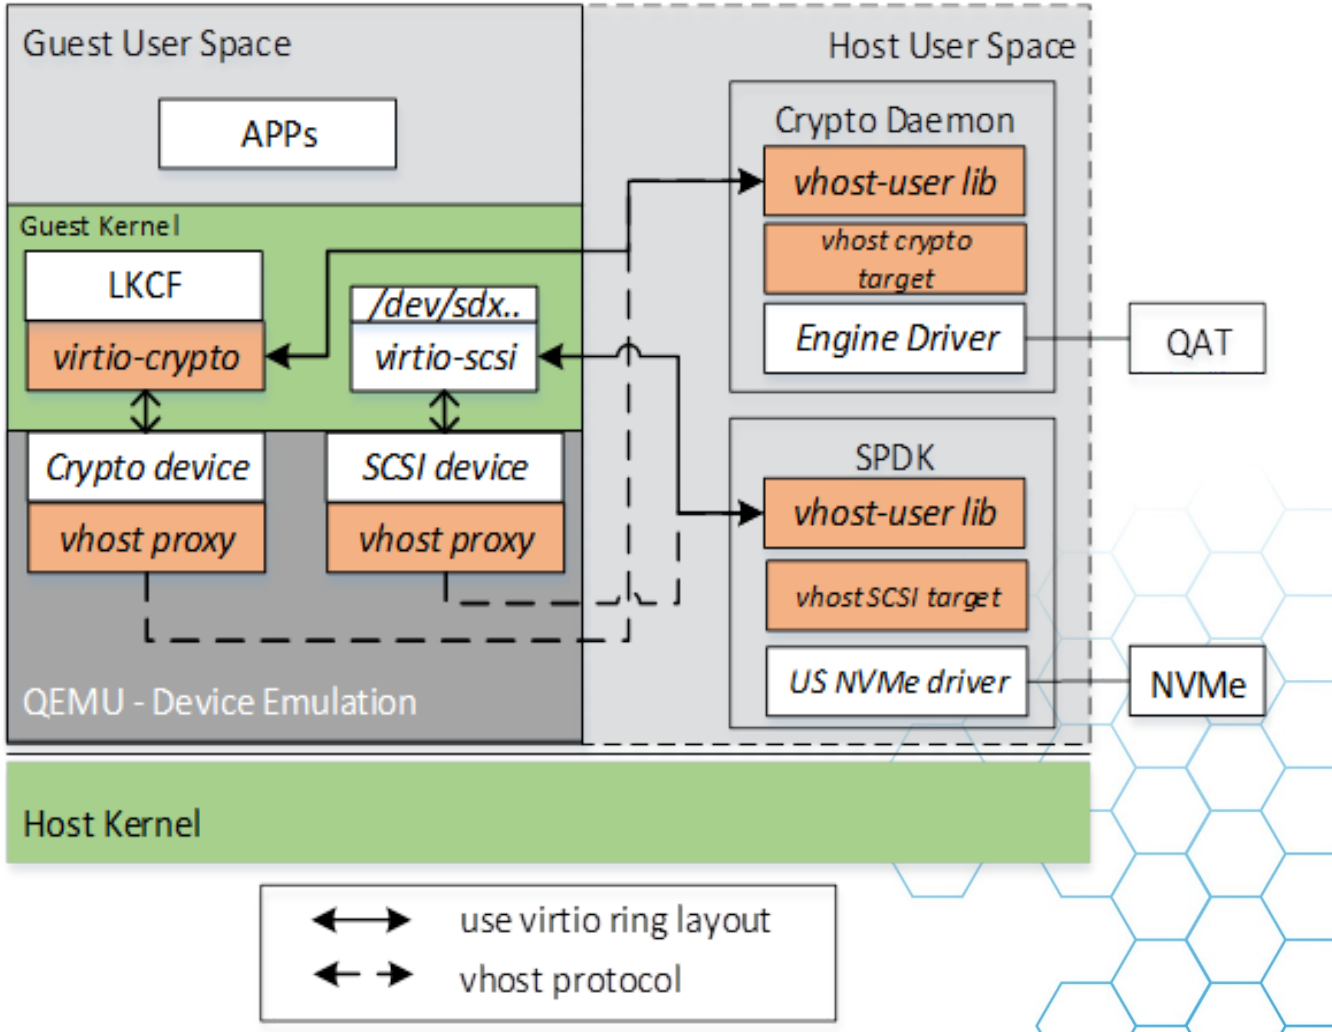

Thanks!!

Backup

- http://spdk.io
- Code available at https://github.com/spdk/spdk
- Submit your patch via https://review.gerrithub.io/spdk/spdk