Scalable-I/O Virtualization (SIOV)

Contribution & Workstream Formation

Tom Stachura, Intel Corporation

Renee L'Heureux, Microsoft Corporation

OCP: Server Project Meeting, February 23, 2022

What is Scalable I/O Virtualization (SIOV)?

- SIOV is **hardware-assisted I/O virtualization** designed for the hyperscale era, with the potential to support **thousands of virtualized workloads** per server.
  - SIOV moves the non-performance-critical virtualization and management logic **off the device and into the virtualization stack**. It uses a **new scalable identifier** on the device to address the workloads' memory.

- SIOV delivers key benefits vs. prior art:
  - **Reduces the per-VM virtualization cost** on the devices
  - More efficiently **supports large numbers of VMs and containers**
  - Provides **more flexibility to the virtualization stack** for provisioning & composability.

- Impact: Virtualized I/O devices become much **more configurable and scalable** while delivering **near-native performance** to each VM/container/microservice.

Why Scalable I/O Virtualization (SIOV)?

- SIOV enables virtualization **modernization of I/O devices** (PCIe, CXL, etc.) for today's **massive-scale clouds**
  - SIOV is much **more scalable and flexible** than previous technologies, providing efficient I/O access to potentially **thousands of containers** or virtual machines
  - SIOV is simpler and **more efficient to implement** and deploy for I/O device manufacturers, software providers and cloud operators

- Driving as a standard **enables an open ecosystem** that fuels adoption, growth, and innovation
  - Host Devices
  - Target Devices
  - Software & Operating Systems

SIOV Contribution & Workstream Plan

- **Intel & Microsoft are contributing SIOV** to OCP to encourage an open ecosystem hyperscale I/O virtualization
  - Scalable I/O Virtualization (SIOV) Revision 1 contribution in process
    - Specification being contributed is posted for Server Project audience
  - **Call to Action**: Adopters/Suppliers are encouraged to register adoption

- **Intel, Microsoft & others are committed to evolve SIOV** to drive a broader ecosystem & further enhance data center I/O virtualization
  - Announcing intent to start a **SIOV Revision 2 workstream**. Scope focus:
    - Enhance definition where different implementations could cause incompatibilities
    - Extend definition to meet new, emerging, and complementary usages. e.g., Security
    - Enable a consistent multi-vendor experience. e.g., QoS, RAS, device mix
  - **Call to Action**: Get involved!
    - Parties interested in contributing resources and/or IP are encouraged to reach out

BACKUP

SR-IOV (Single-Root I/O V) → Scalable I/O Virtualization (Scalable-IOV)

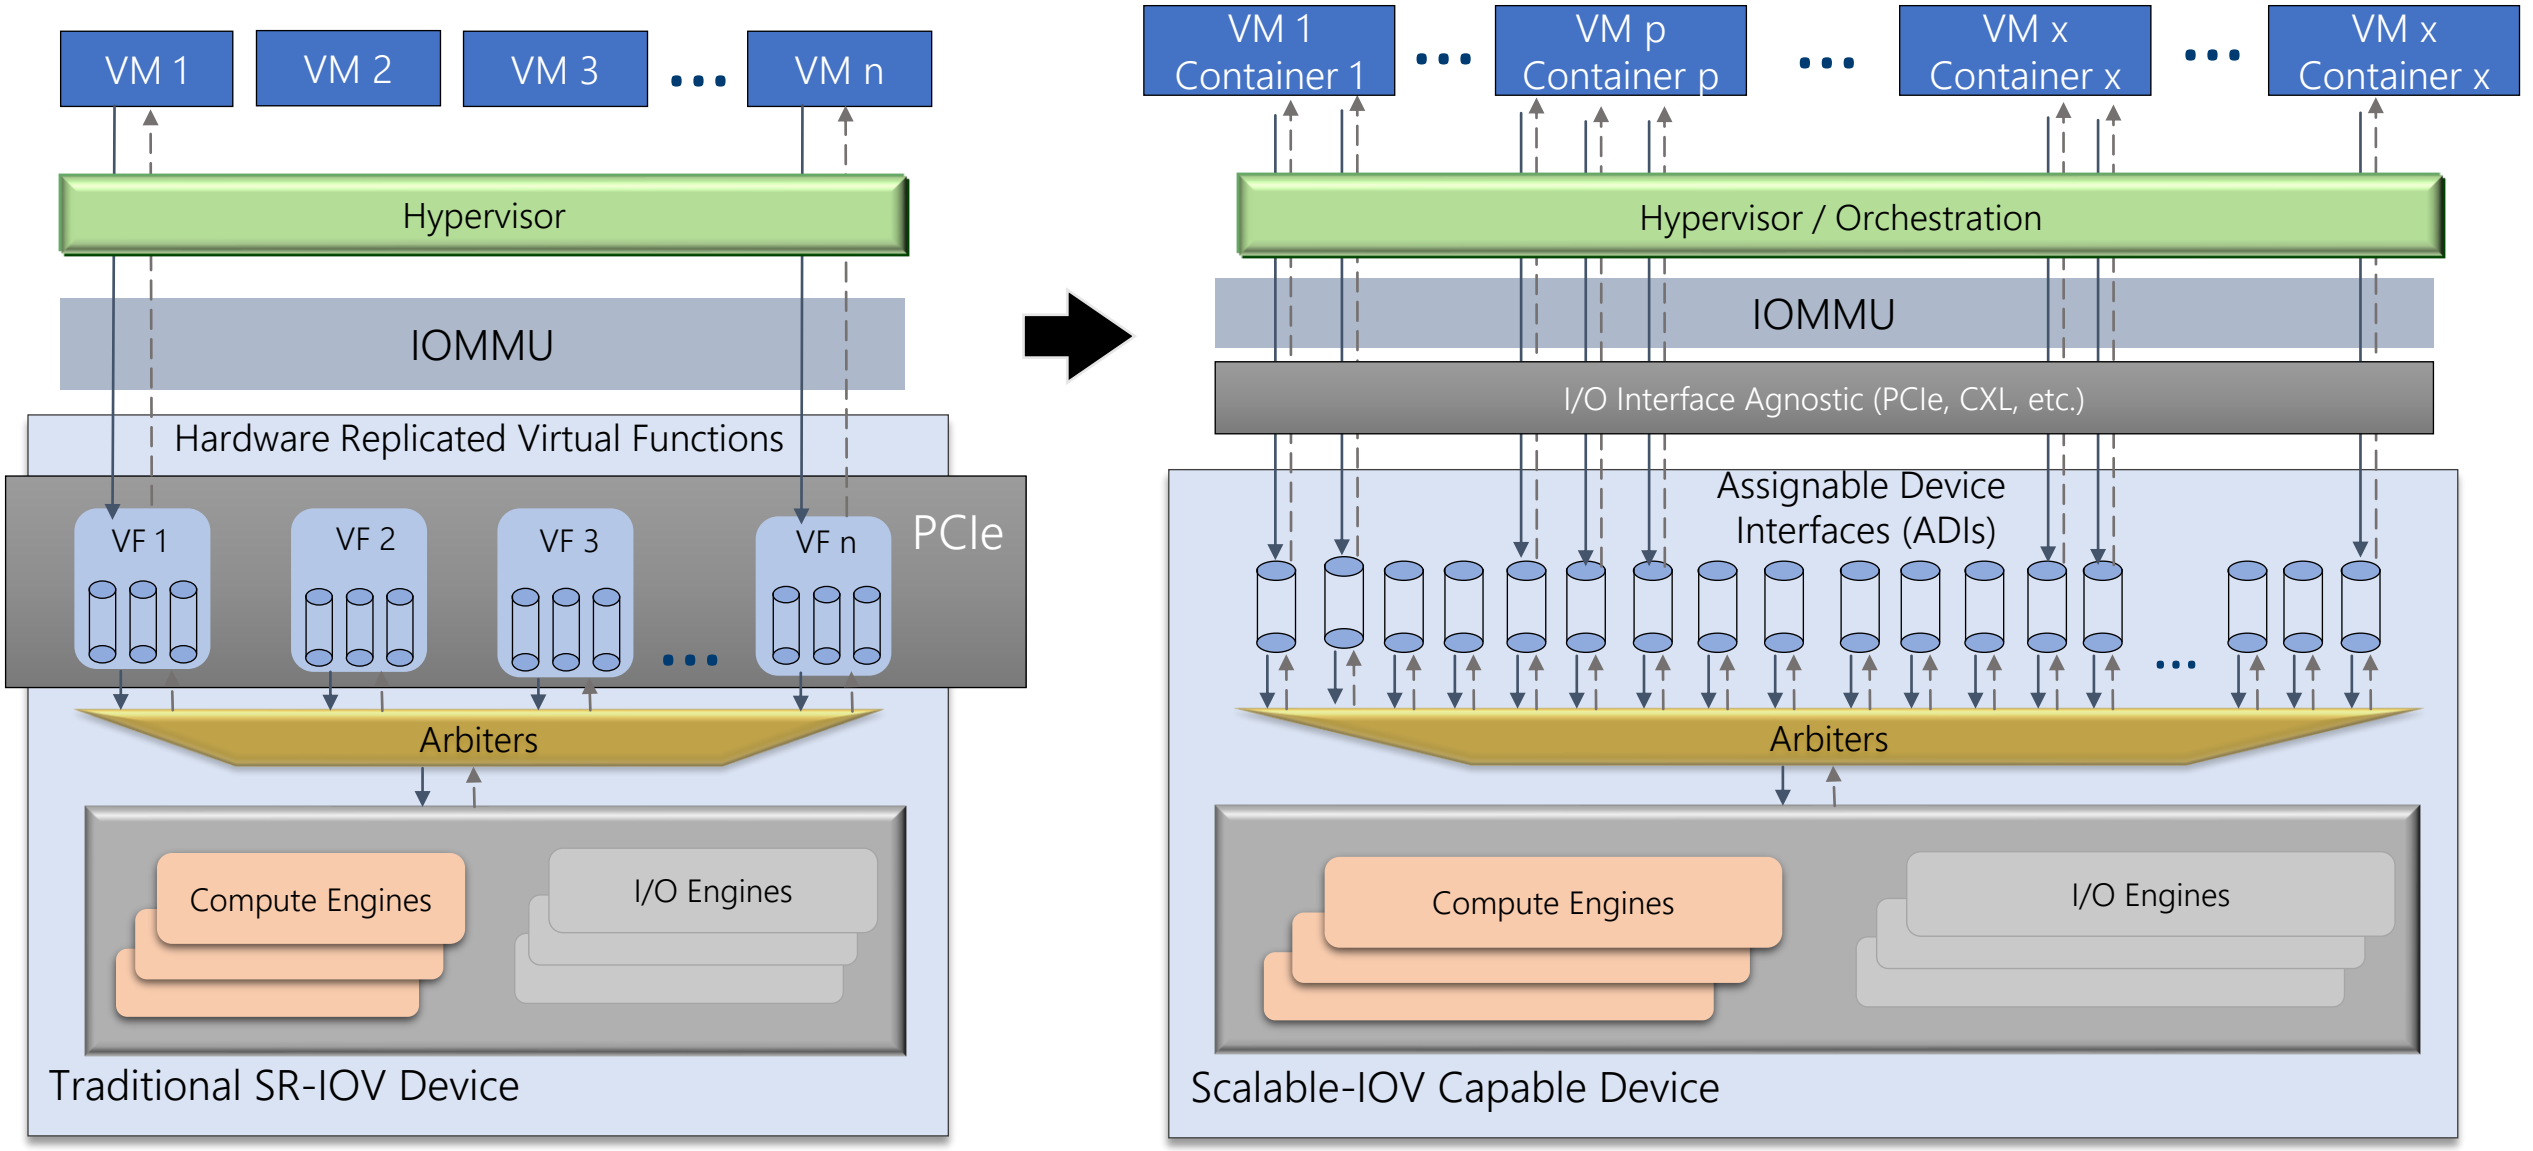

- While SR-IOV maps virtual devices as Virtual Functions (VFs) in HW, S-IOV composes virtual devices through light-weight and scalable Assignable Device Interfaces (ADIs) for fast-path operations and software mediated slow-path operations

Scalable IOV Benefits

Scalability

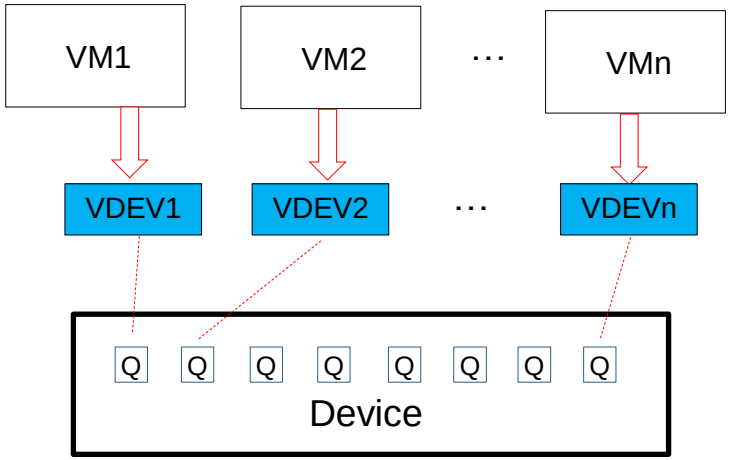

Flexibility

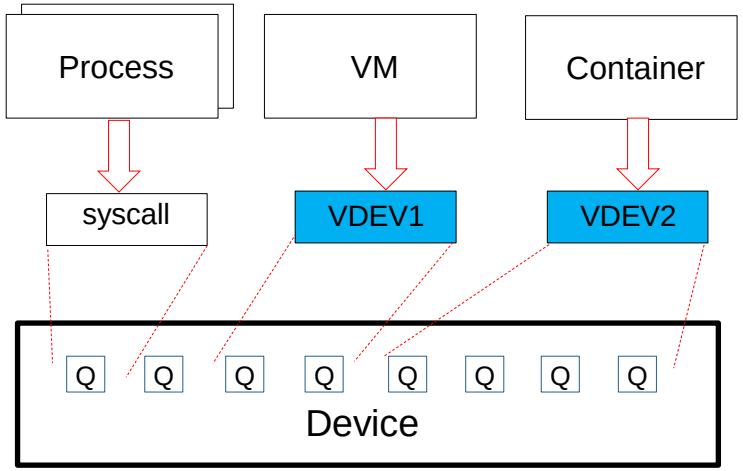

Over-provisioning

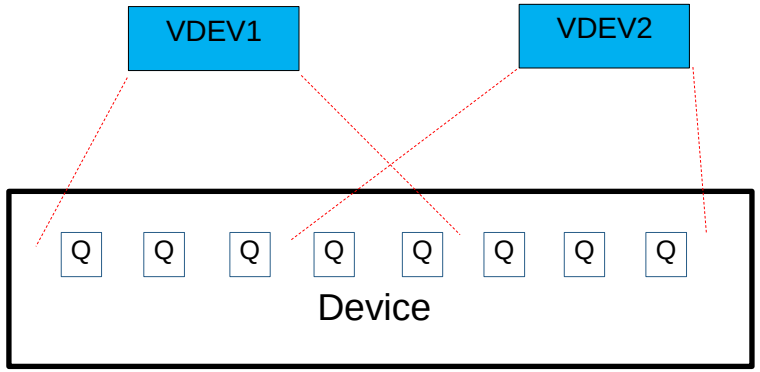

Composability

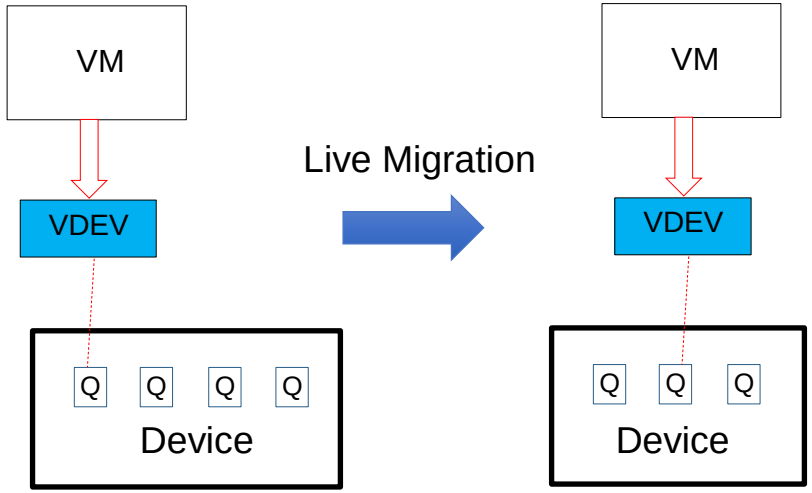

Scalable IOV: High Level Software Architecture

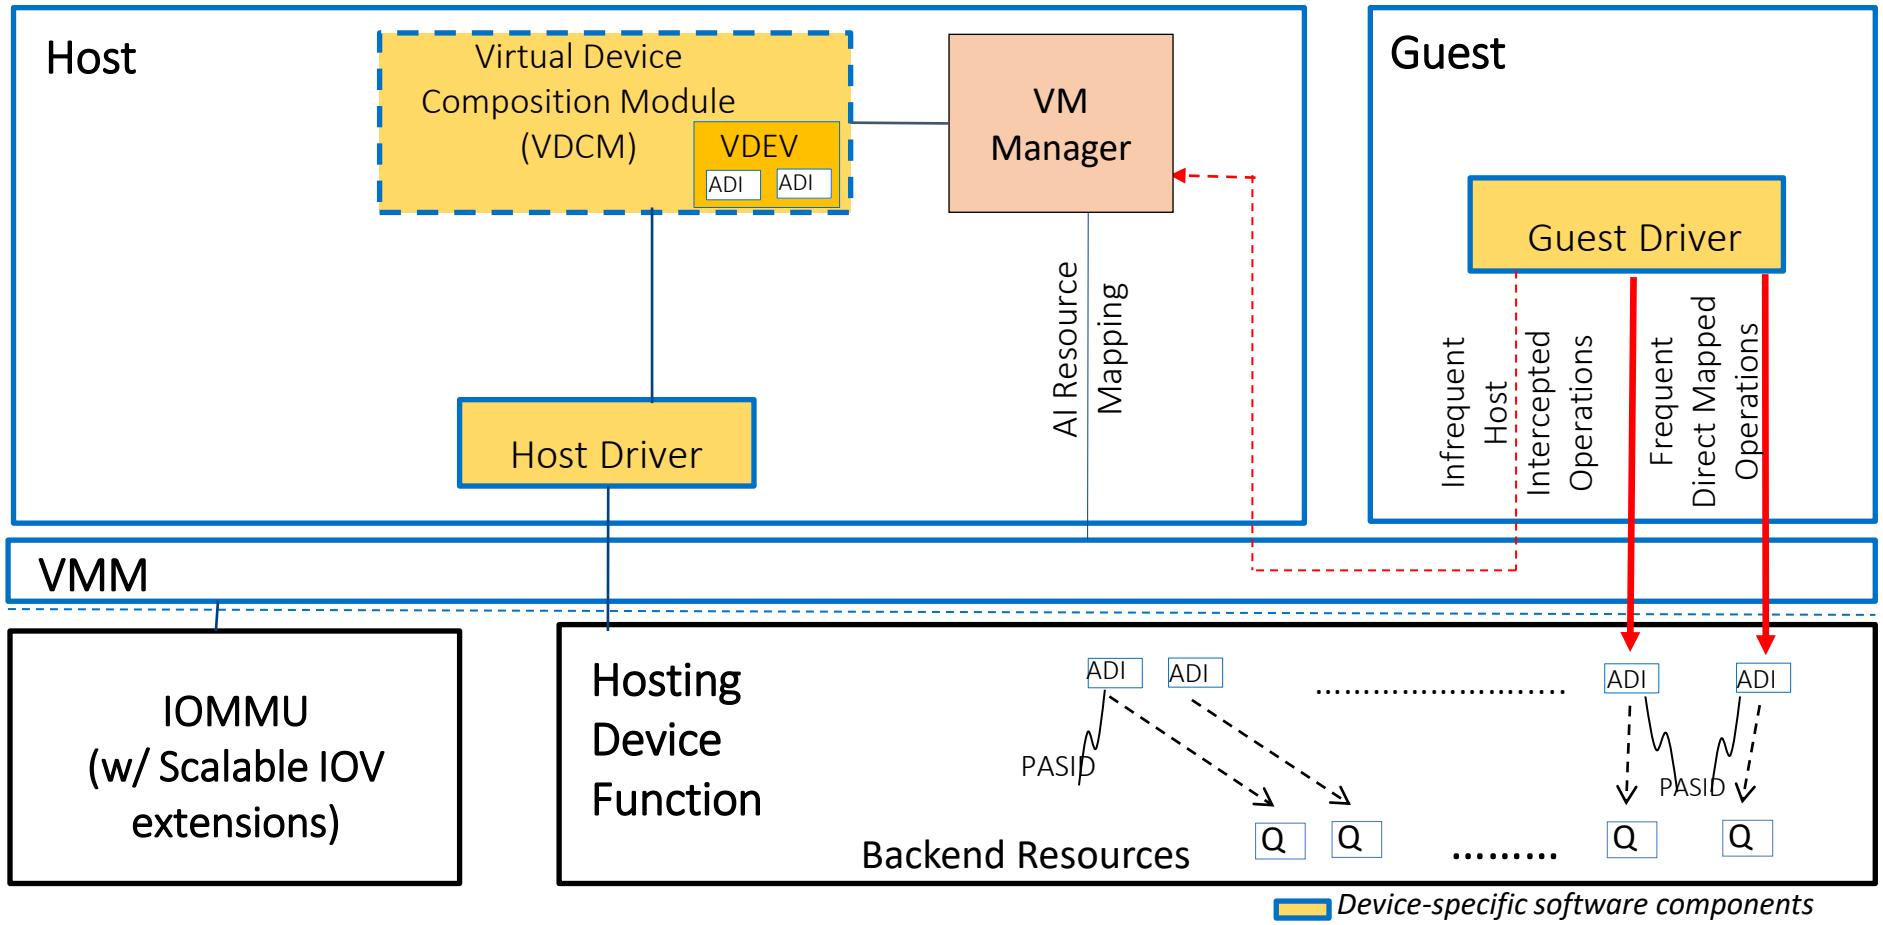

Device specific VDCM to decide what to intercept vs direct map

Emulating infrequent operations in VDCM simplifies device without compromising performance