Virtual Open Systems

AGL Virtio-loopback
Code review planned on (9/11/2022)

A work carried on by Virtual Open Systems, on behalf of Linux Foundation,
to enhance Automotive Grade Linux (AGL)

contact@virtualopensystems.com

www.virtualopensystems.com

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- Questions

Discussion Index

- **Brief design description and current status**
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- Questions

Virtio-loopback approach

Virtio Loopback describes a new Hardware Abstraction Layer (HAL) for non-Hypervisor environments based on virtio.

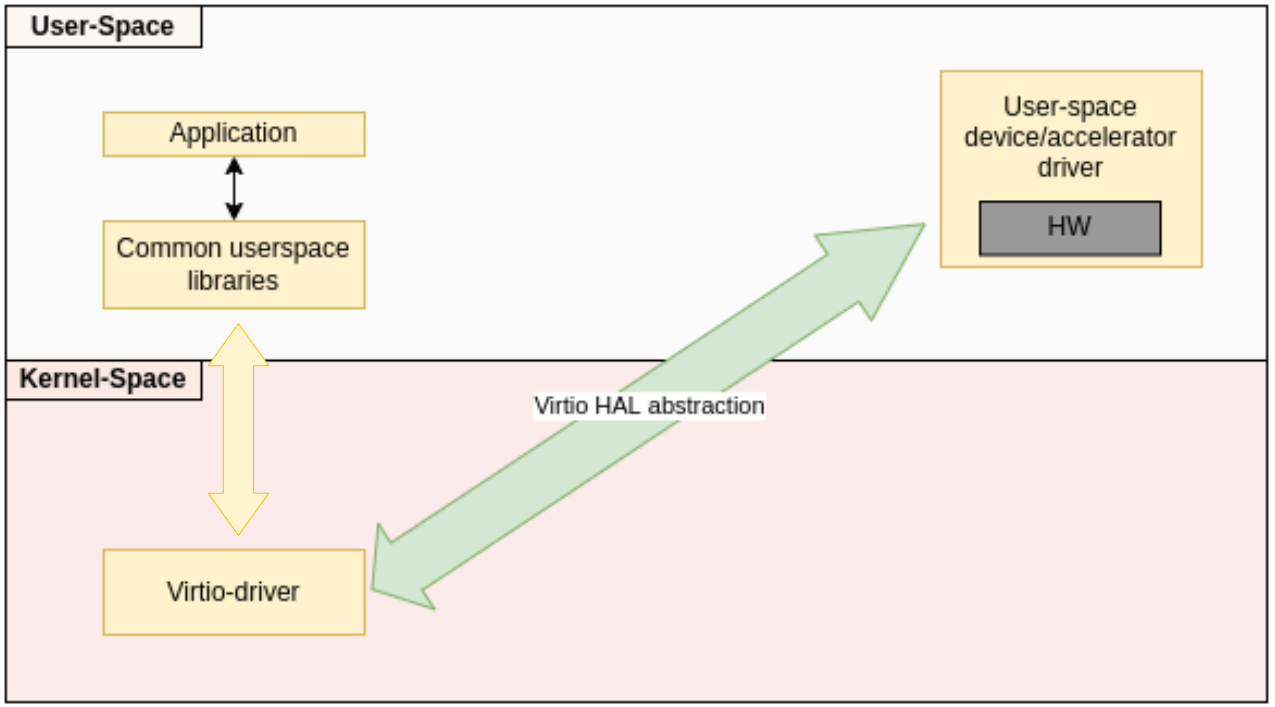

Virtio-loopback gives the ability to host user-space applications to take advantage of user-space drivers

Virtio-loopback components

- The kernel space component is a **new virtio transport** that forwards driver calls in user space where the device is implemented.
- The second component is an **application in user space** (virtio-loopback-adapter) that is particularly important for the set-up of the system configuration, but that does not impact the data plane path to avoid overhead.

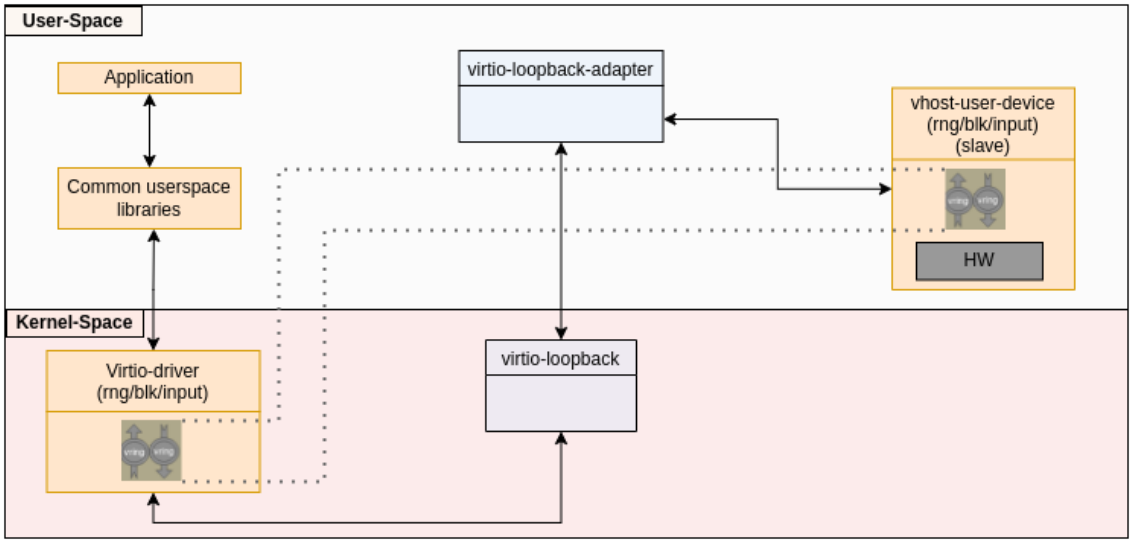

Current status of the activity

- **Alpha release**: Publicly released with docs and demo
- **Beta release**: Done for this review
- Next steps:
  - Merge the beta into the master
  - Benchmarks
  - Prepare AGL patches

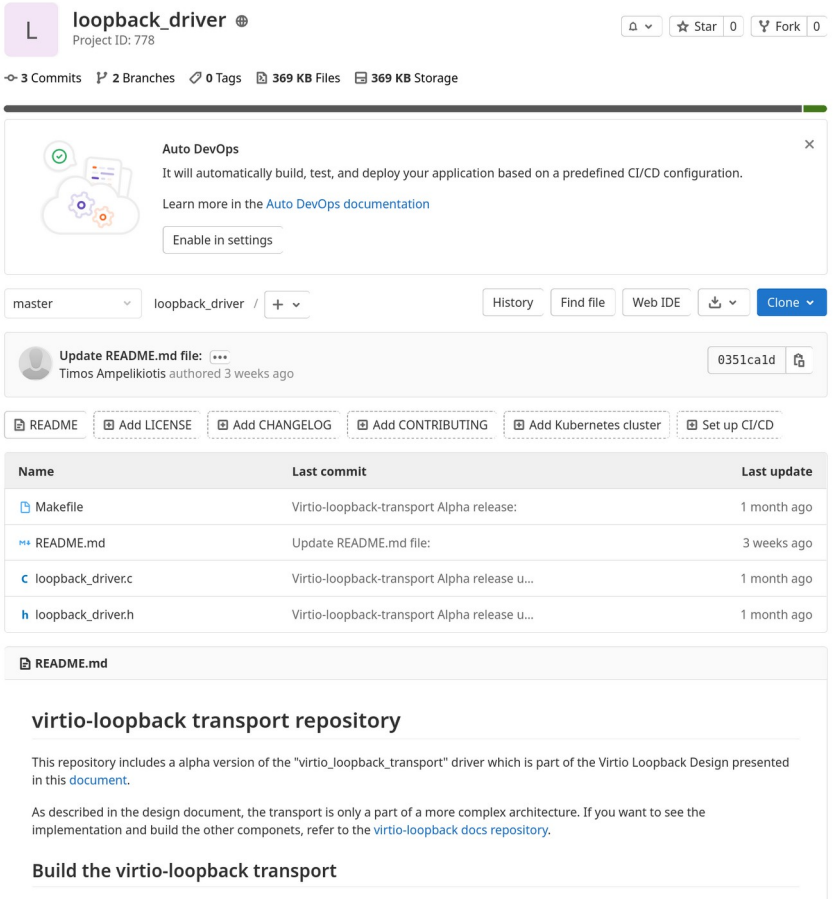

Discussion Index

- Brief design description and current status
- **Code review**
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- Questions

Discussion Index

- Brief design description and current status
- Code review
  - **Control plane**
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- Questions

Control Plane (1)

The following three steps are describing briefly the control plane which takes place before the whole system is ready to exchange any data.

The control plane consists of stages of communication:

- [Stage 1] Adapter ↔ Vhost-user-device
- [Stage 2] Adapter ↔ Transport driver
- [Stage 3] Adapter ↔ Vhost-user-device

Control Plane (2)

- [Stage 1] Adapter ↔ Vhost-user-device
  - The vhost-user-device sends via the Unix socket to the adapter:
    - Virtio features, vhost-user protocol features, virtio device configuration

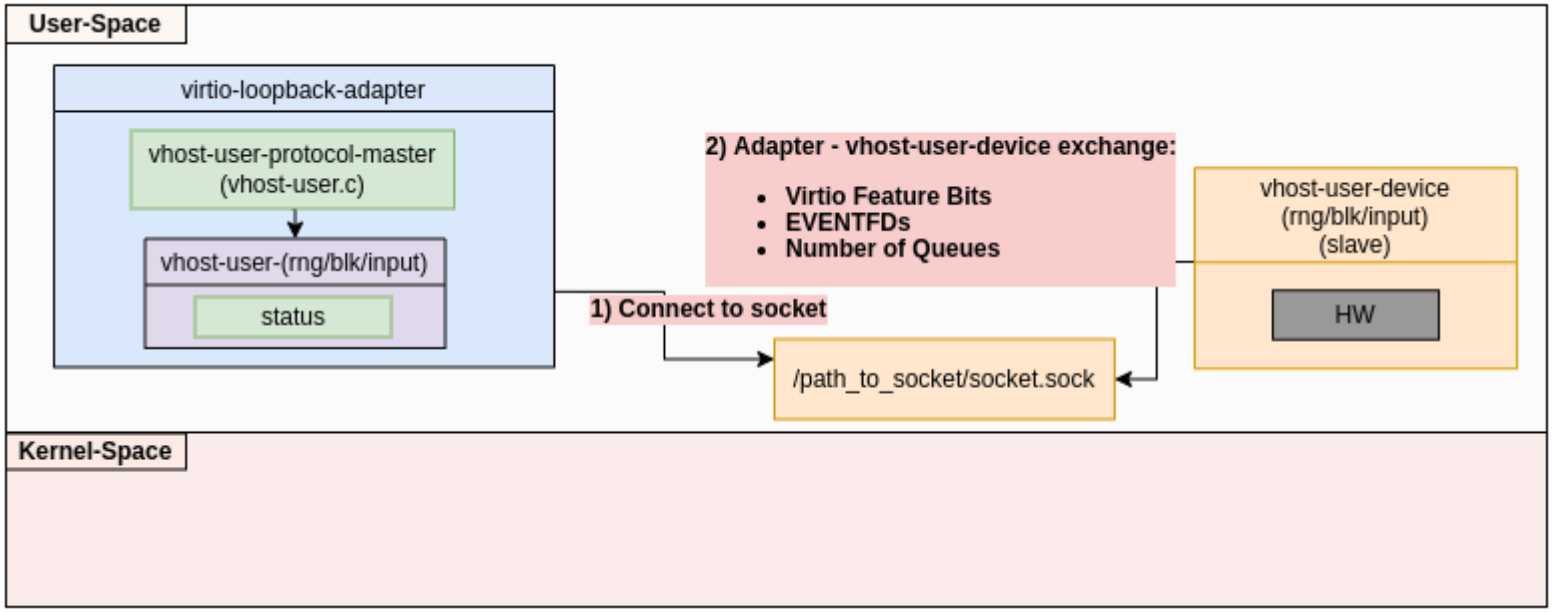

Control Plane (3)

- [Stage 2] Adapter ↔ Transport driver
  - The adapter sends to the loopback driver:
    - Virtio specific information: Device id, Vendor, magic number
    - Virtio device features
  - The virtio-loopback-transport starts and register the corresponding virtio device (blk, input, rng)
  - Acknowledges the features and writes back to the adapter

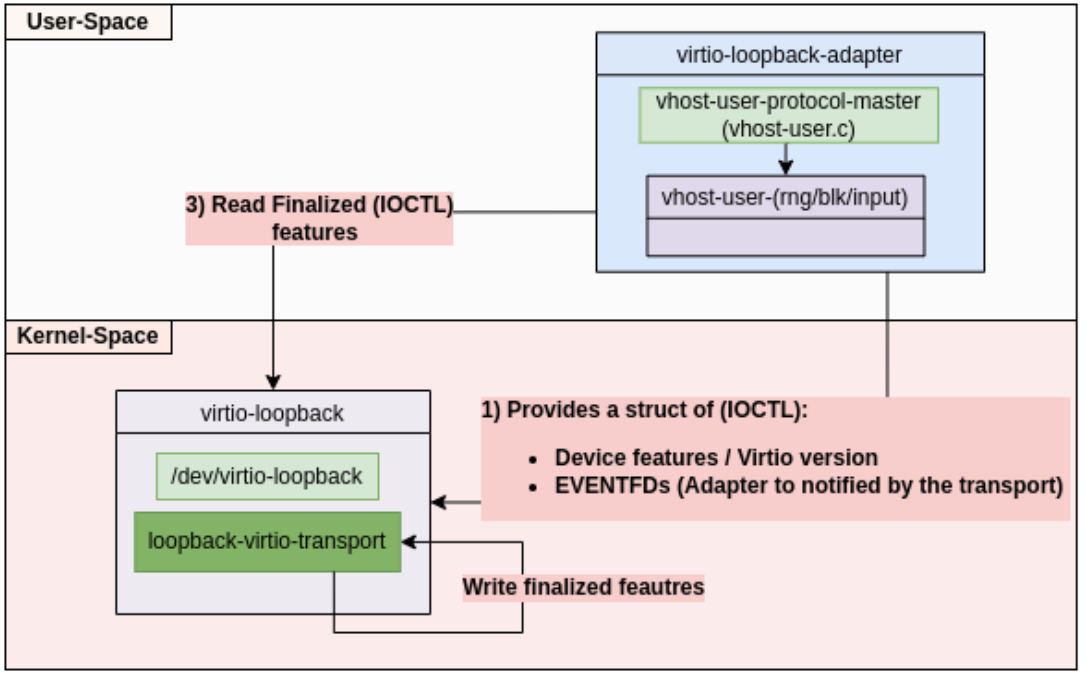

Control Plane (4)

- [Stage 3] Adapter ↔ Vhost-user-device
  - The adapter sends to the vhost-user-device:
    - The acknowledged virtio features

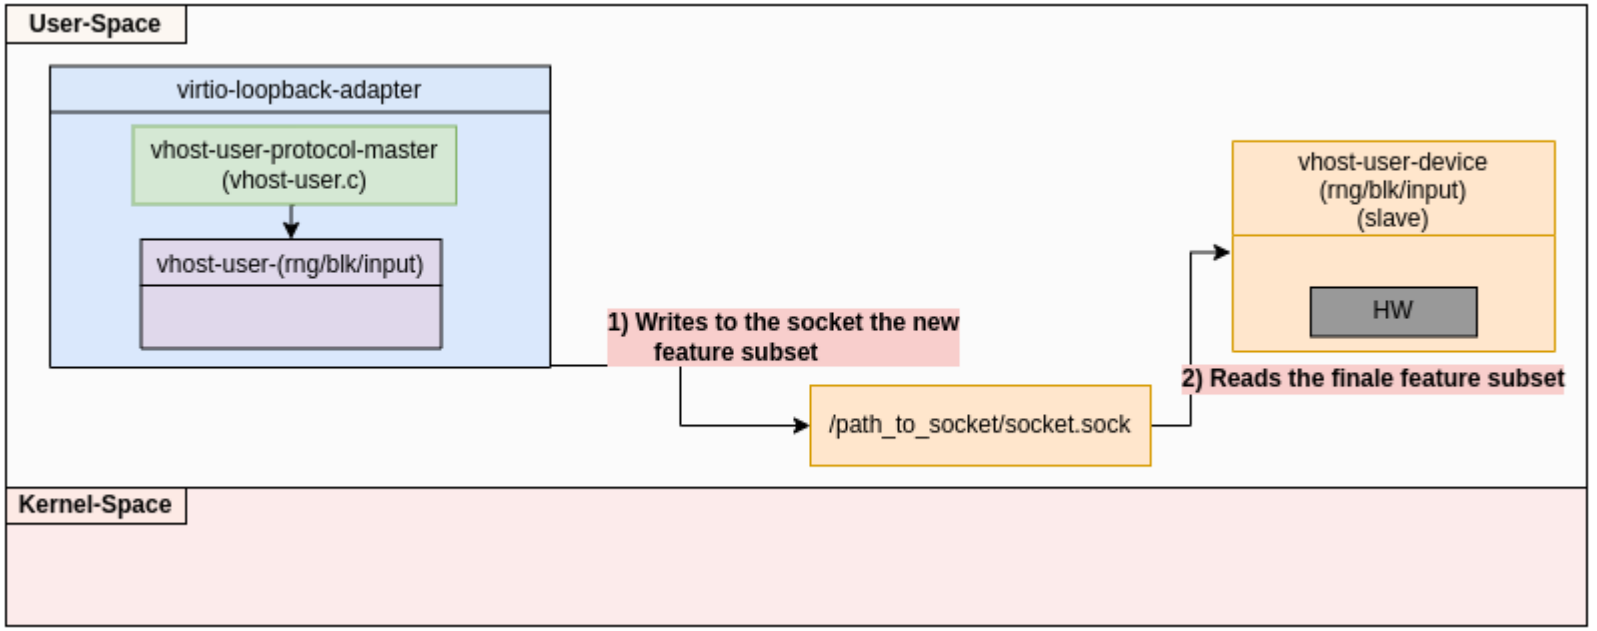

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - **Communication mechanisms**
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- Questions

Communication mechanisms used between the components

- **Adapter ↔ driver : Eventfds, SYS_CALLS**

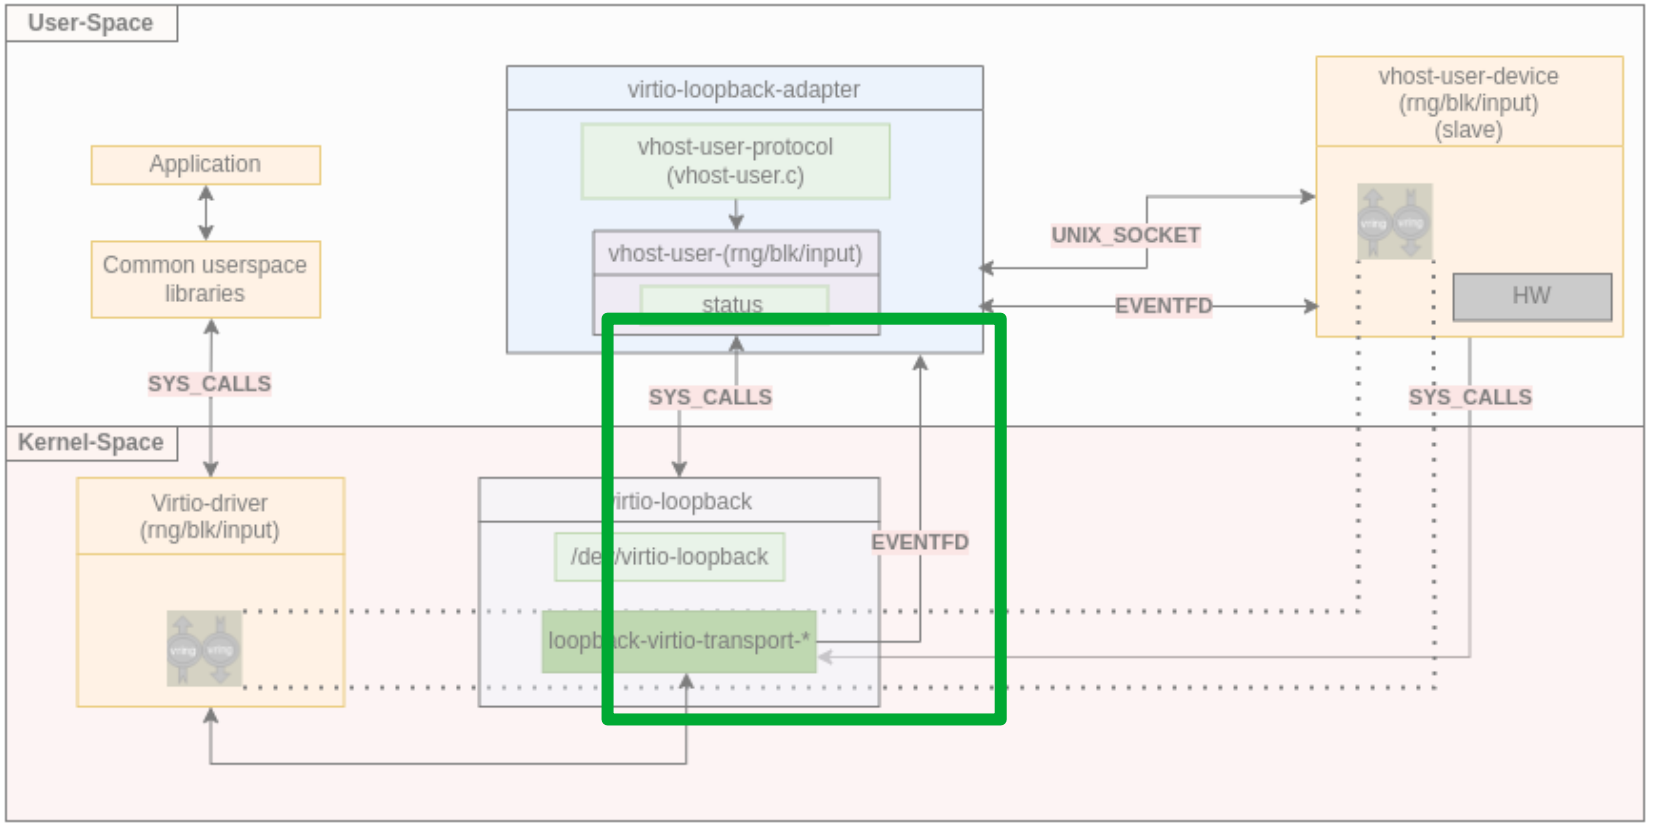

Communication mechanisms used between the components

- Adapter ↔ driver : Eventfds, SYS_CALLS
- **Vhost-user-device ↔ adapter** : Eventfds, Unix Socket

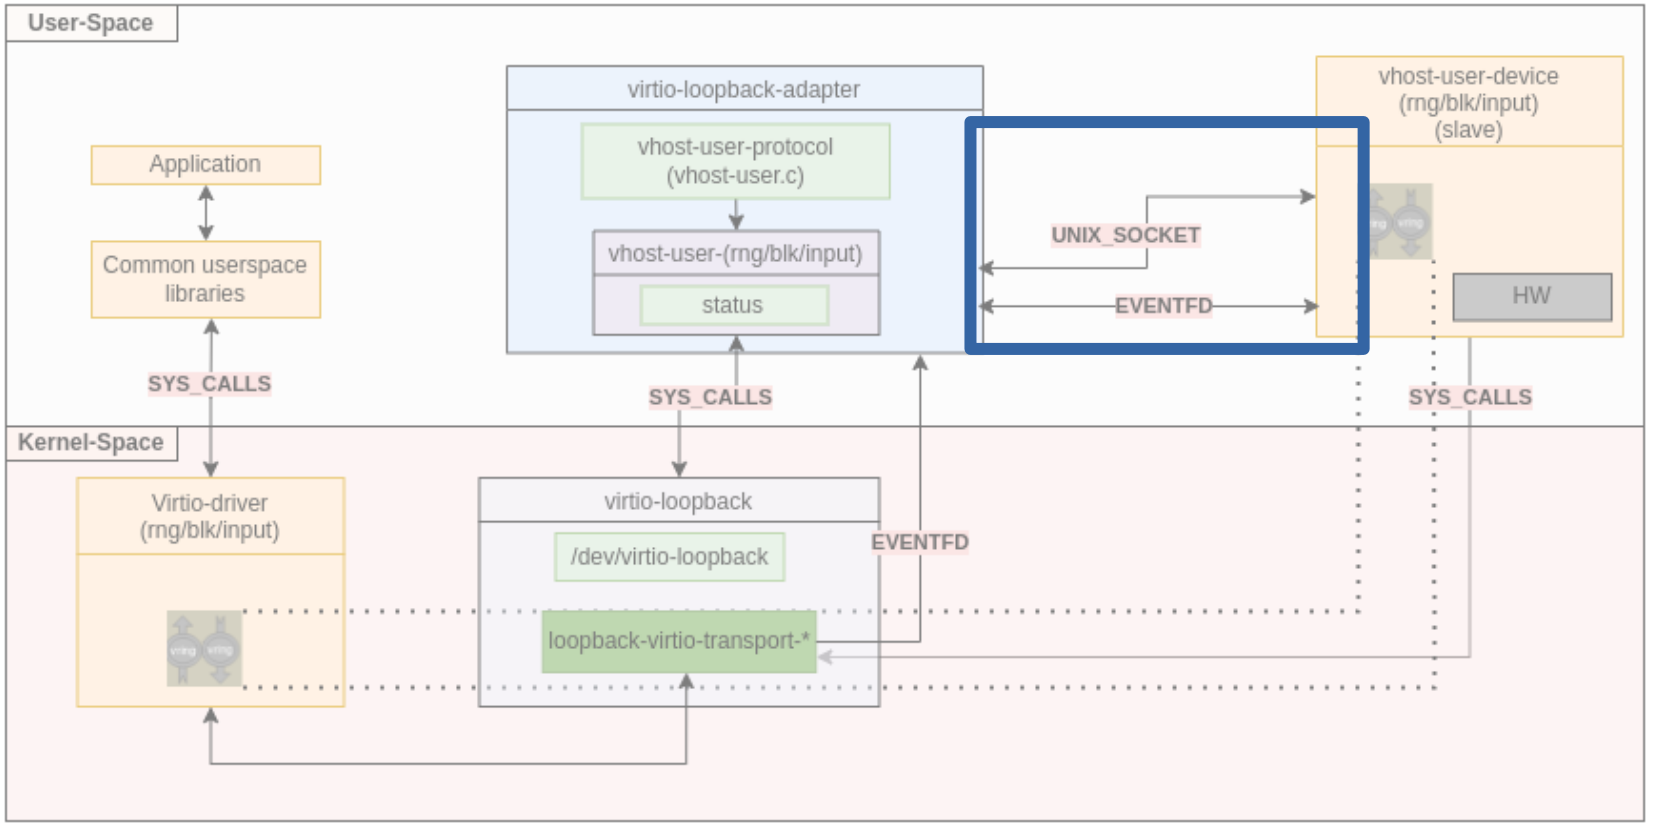

Communication mechanisms used between the components

- Adapter ↔ driver : Eventfds, SYS_CALLS
- Vhost-user-device ↔ adapter : Eventfds, Unix Socket
- **Vhost-user-device ↔ driver**: SYS_CALLS

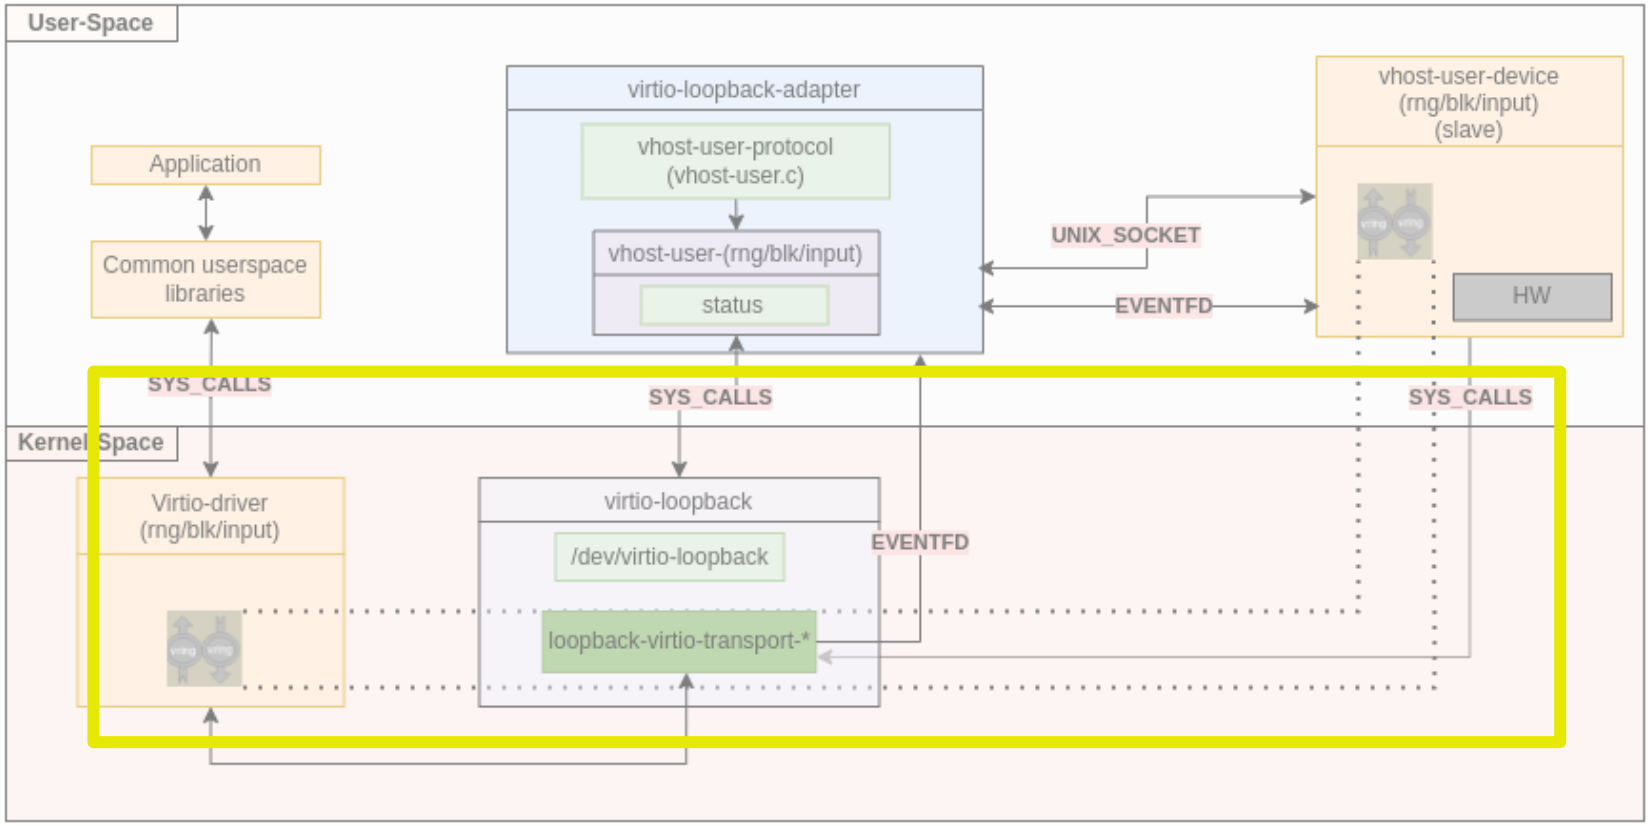

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - **Memory mapping (data plane)**
- Live demo
- Upstream
- Next steps
- Questions

Memory mapping mechanism (data plane)
Vhost-user mmap the Vrings (part 1)

The vhost-user-device access the data exchanged with the virtio device in two parts. First, access the vring data structure (in kernel space) and then access the SG list entries pointed by the vrings.

- **Vring data structure** The device uses **mmap** in order to obtain access

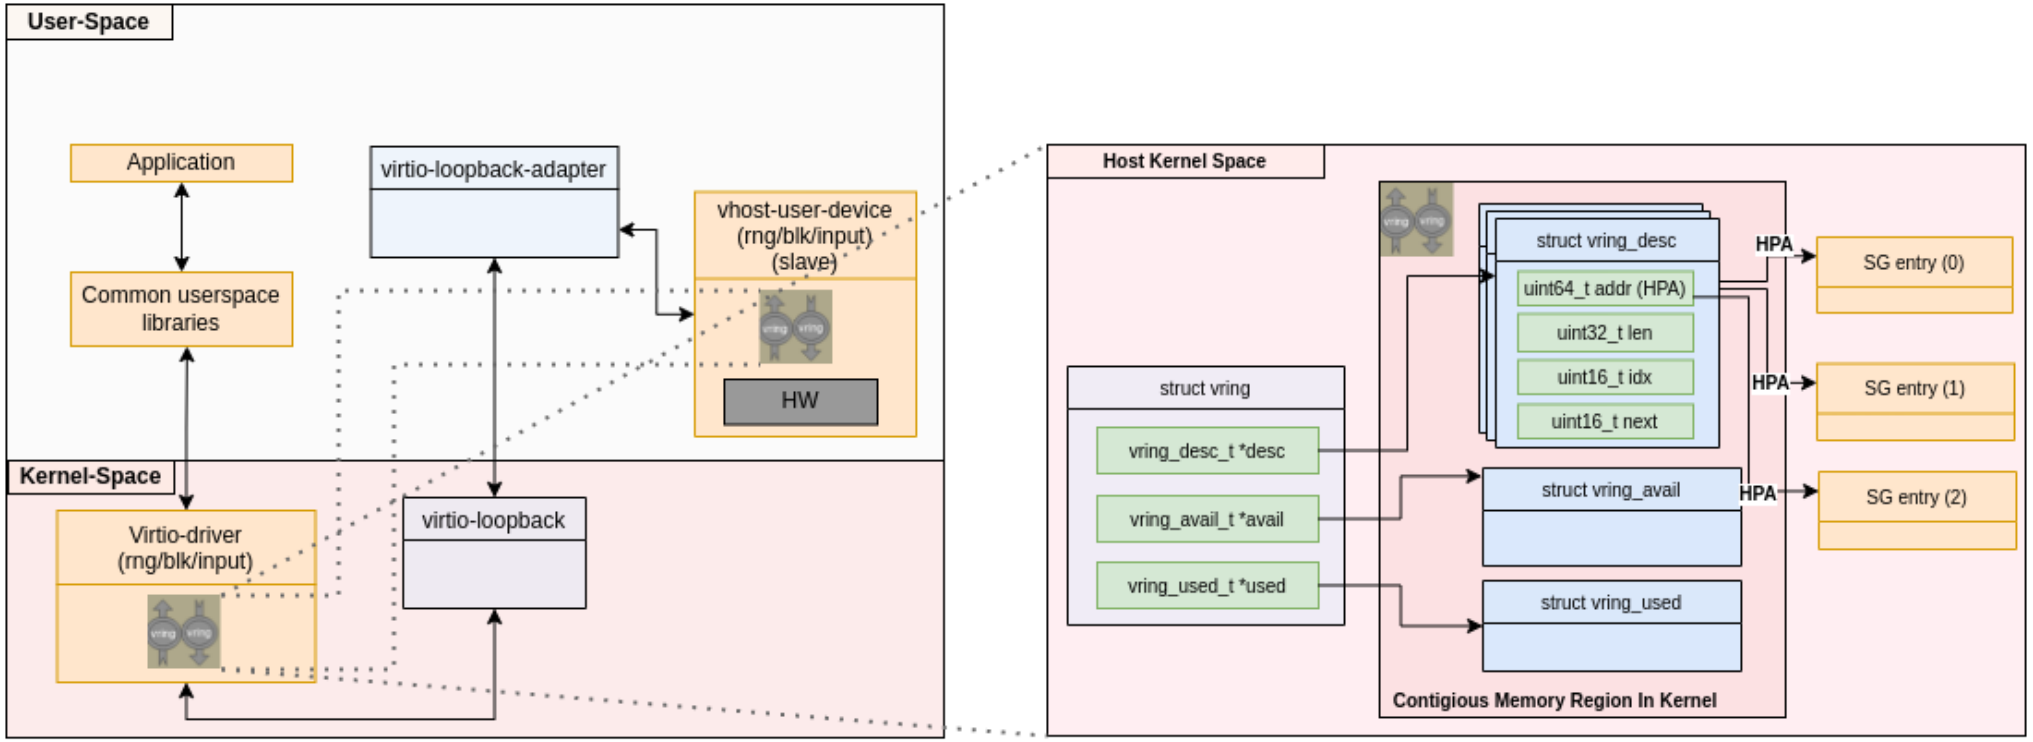

Memory mapping mechanism (data plane)
Vhost-user-device mmap the SG list (part 2)

SG list entries The vhost-user-device uses the ioctl in order to ask the transport driver for access to the buffer (Host Physical Address) pointed the descriptors' table

1) If the driver has already shared this memory page with the vhost-user-device, returns an VMA (Virtual Memory address) in the 'mmap_addr' argument.
2) Otherwise the driver returns 0 and vhost-user-device goes on with calling 'mmap' function
3) The vhost-user-device obtains a pointer to this SG list element

Memory mapping mechanism (data plane)
Vhost-user-device mmap the SG list (part 2)

1) **If the driver has already shared this memory page with the vhost-user-device, returns an VMA (Virtual Memory address) in the ‘mmap_addr’ argument.**
2) Otherwise the driver returns 0 and vhost-user-device goes on with calling ‘mmap’ function
3) The vhost-user-device obtains a pointer to this SG list element

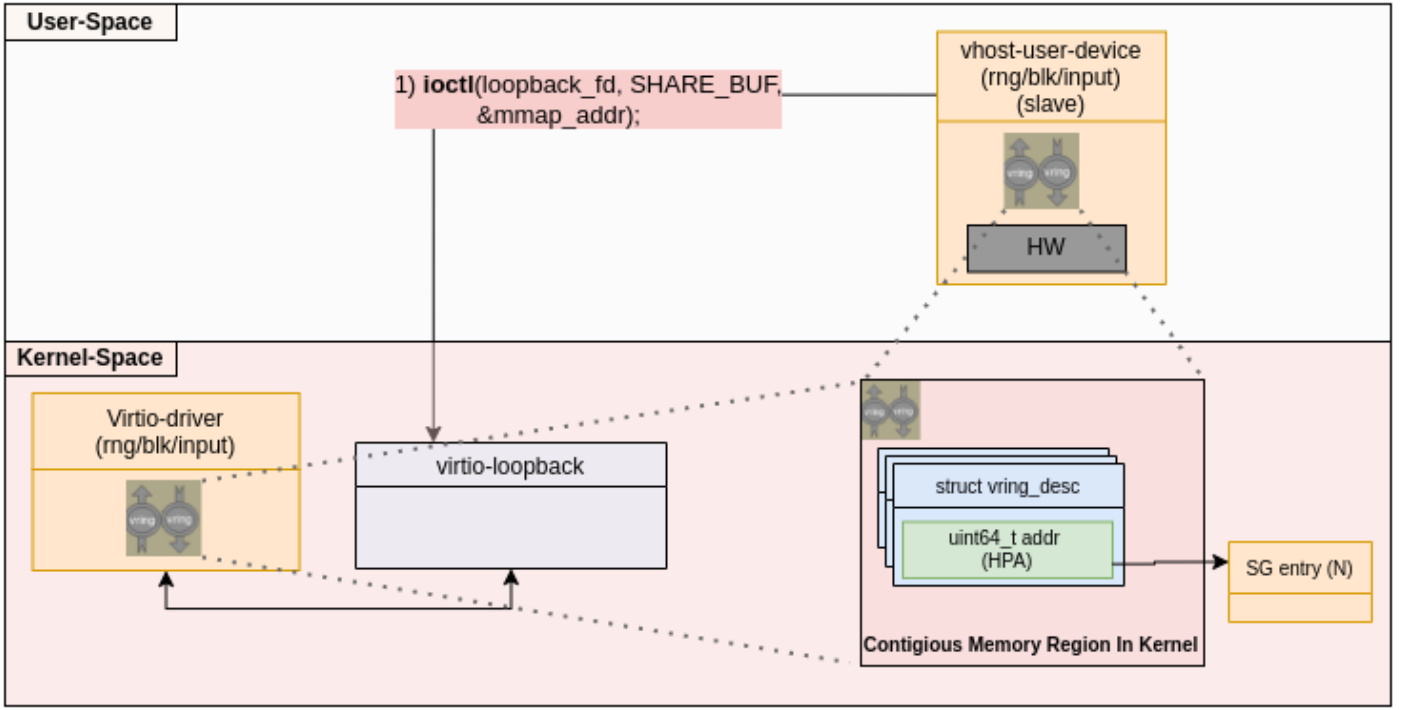

Memory mapping mechanism (data plane)
Vhost-user-device mmap the SG list (part 2)

1) If the driver has already shared this memory page with the vhost-user-device, returns an VMA (Virtual Memory address) in the ‘mmap_addr’ argument.
2) **Otherwise the driver returns 0 and vhost-user-device goes on with calling ‘mmap’ function**
3) The vhost-user-device obtains a pointer to this SG list element

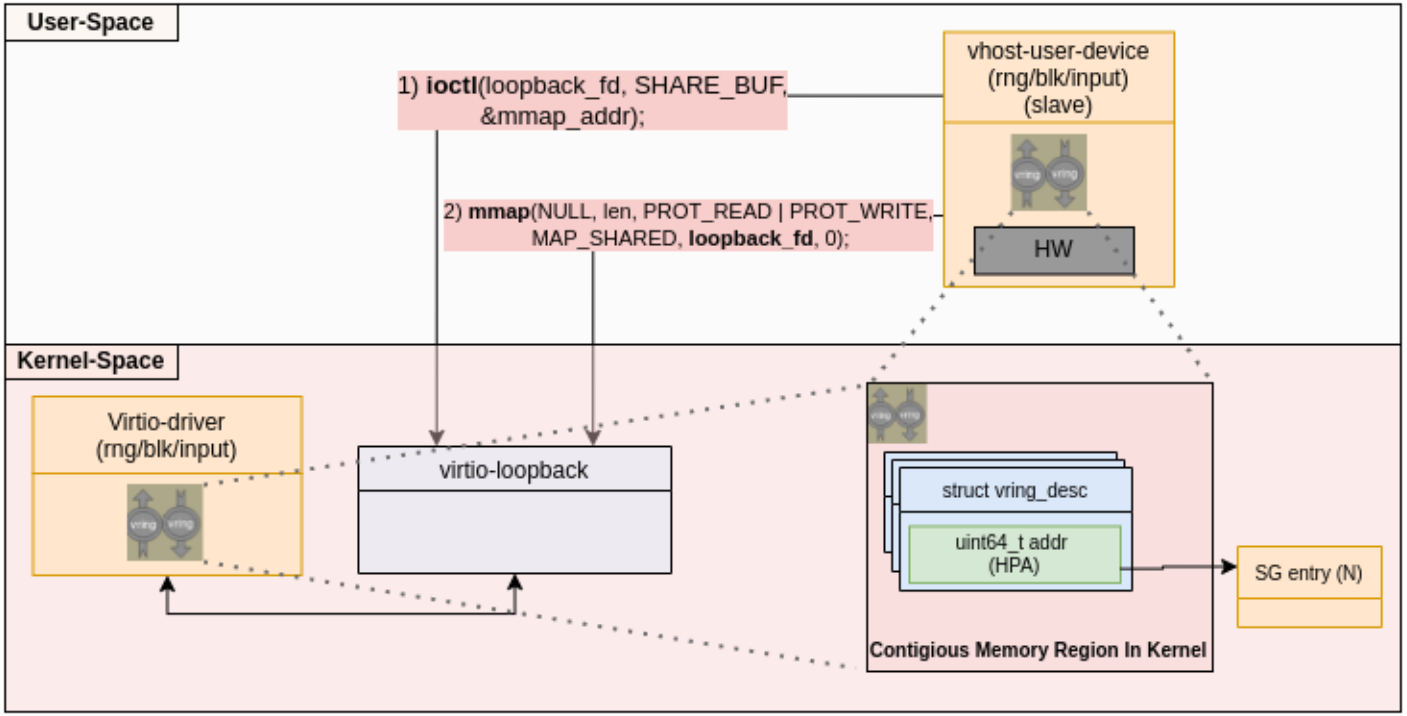

Memory mapping mechanism (data plane)
Vhost-user-device mmap the SG list (part 2)

1) If the driver has already shared this memory page with the vhost-user-device, returns an VMA (Virtual Memory address) in the ‘mmap_addr’ argument.
2) Otherwise the driver returns 0 and vhost-user-device goes on with calling ‘mmap’ function
3) **The vhost-user-device obtains a pointer to this SG list element**

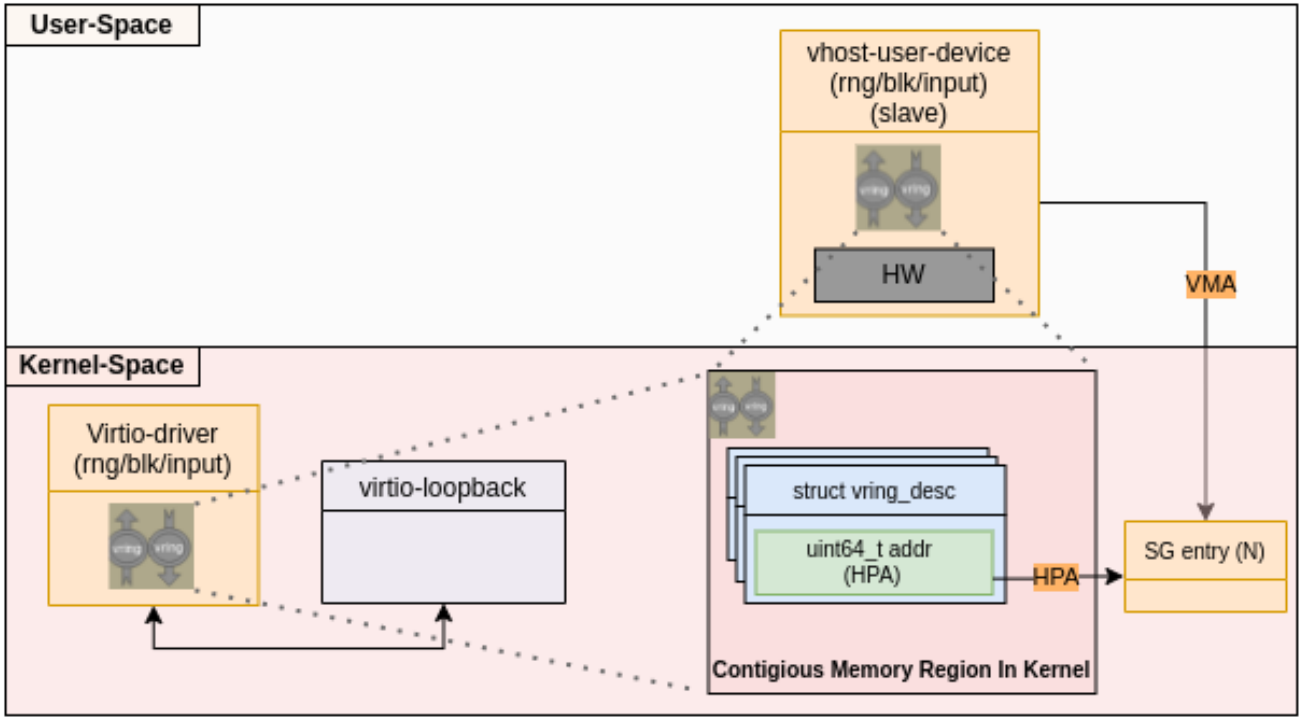

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- **Live demo**
- Upstream
- Next steps
- Questions

Live demo (input + blk)

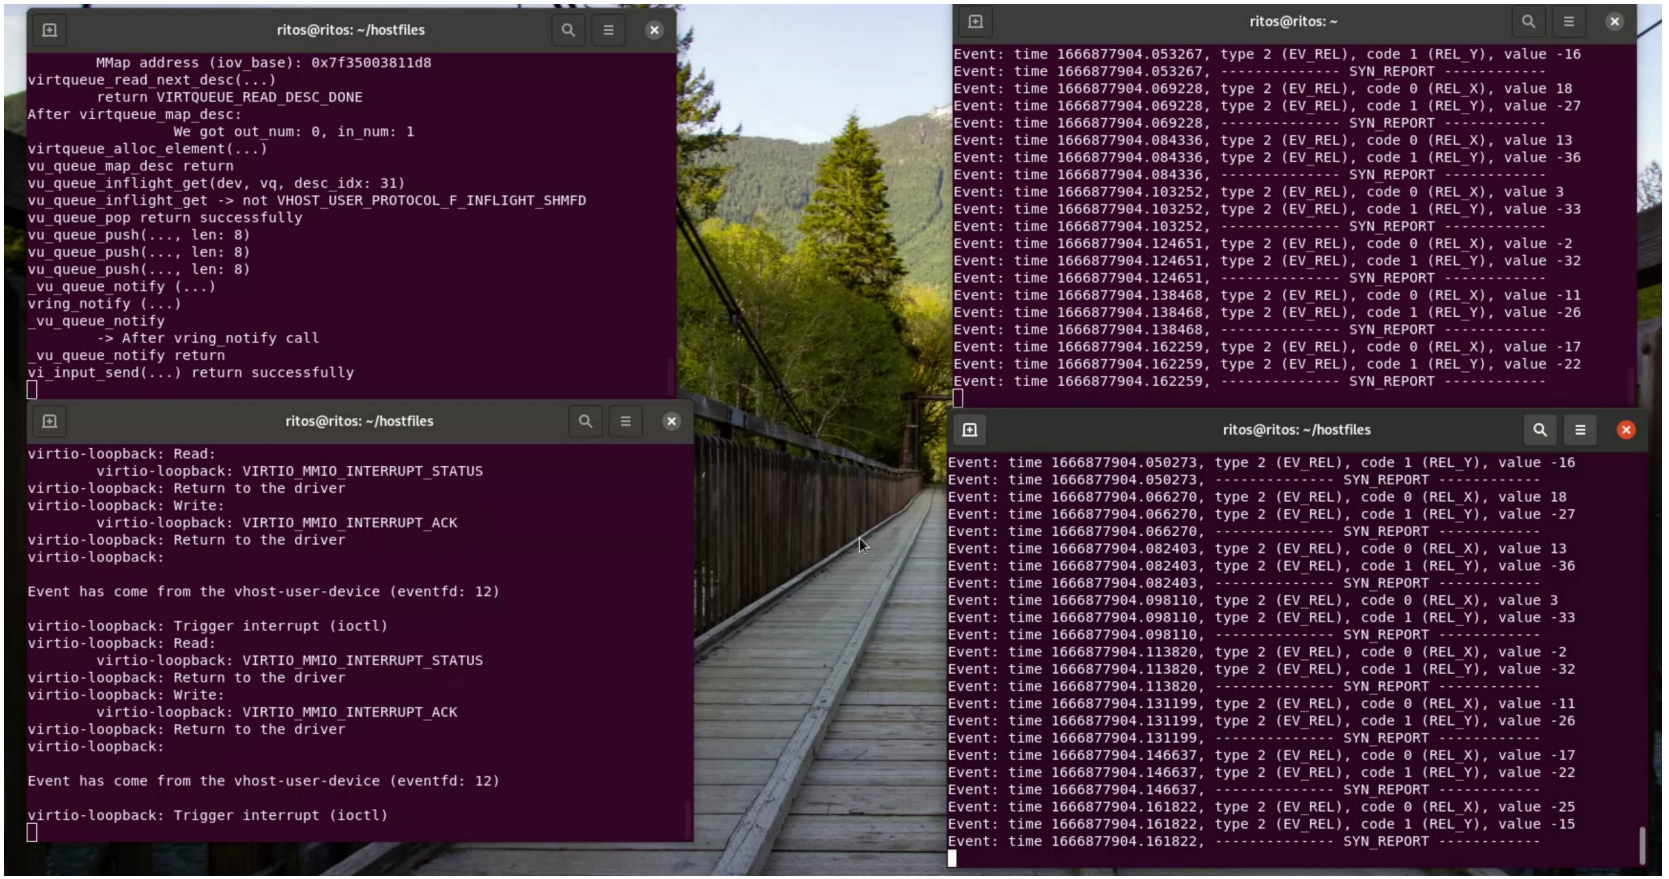

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- **Upstream**
- Next steps
- Questions

Upstream

The key upstream target for this work is AGL. Yocto layers and recipes will be created and integrated in meta-egvirt. However, each and every component could be integrated in the reference community

- Virtio-loopback-transport driver → Linux Kernel
  - (Mailing list: dev-mailing-lists-virtio)
- Virtio-loopback-adapter → Qemu project
- Vhost-user devices:
  - Vhost-user-library → Qemu project
  - Vhost-user-rng (RUST)
    - https://github.com/rust-vmm/vhost-device
    - https://github.com/rust-vmm/vm-virtio
    - https://github.com/rust-vmm/vhost-user-backend
    - https://github.com/rust-vmm/vhost

What's the best strategy to propose these components to the community(ies)?

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- **Next steps**
- Questions

Next steps

- Merge beta version into master branch
  - More testing
  - Polish the code / address comments
  - Update / add documentation
- Run benchmarks for vhost-user-blk & input
  - Using ‘dd’ for measuring blk throughput
  - Measure event latency with the help of ‘vinput’
- Create patches for the AGL
  - Meta-egvirt additions

Discussion Index

- Brief design description and current status
- Code review
  - Control plane
  - Communication mechanisms
  - Memory mapping (data plane)
- Live demo
- Upstream
- Next steps
- **Questions**

Questions

?

contact@virtualopensystems.com

Web: virtualopensystems.com

Products: http://www.virtualopensystems.com/en/products/

Demos: virtualopensystems.com/en/solutions/demos/

Guides: virtualopensystems.com/en/solutions/guides/

Research projects: virtualopensystems.com/en/research/innovation-projects/