Xen Source Code Structure and Disk in Xen

Zhiqiang Ma

The Department of Computer Science and Engineering
The Hong Kong University of Science and Technology

Xen and Linux kernel packages

- Xen 3.4.3
  - http://bits.xensource.com/oss-xen/release/3.4.3/xen-3.4.3.tar.gz
- Linux kernel 2.6.32.13
  - http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.32.13.tar.bz2
  - Xen patch for Dom0 (DomU)
    - http://gentoo-xen-kernel.googlecode.com/files/xen-patches-2.6.32-2.tar.bz2

Xen source code

- Xen hypervisor
- Xen tools
- Drivers in Dom0
- Drivers in DomU

Xen interfaces

- xen/include/
  - xen/include/public/
  - xen/include/xen/
  - Interfaces with comments
- xen/include/public/xen.h
  - Guest OS interface to Xen; Hypercall, VIRQ, Shared VCPU info, etc.
- xen/include/public/event_channel.h
  - Event channels between domains

Xen implementation

- xen/arch/x86
  - Architecture independent implementation
- xen/common
  - Common components' implementation
- xen/arch/x86/x86_64/entry.S
  - Hypercall and fault low-level handling routines
- xen/common/event_channel.c
  - Event channel implementation

Xen tools

- tools/
  - Tools for management and driver programs in user space
- tools/python/xen/xend/
  - Xend Daemon in Dom0 for management
- tools/xenstore/
  - XenStore: Hierarchical namespace shared between domains
  - Xenbus: In-kernel API for I/O driver to interact with XenStore
- tools/hotplug/Linux
  - Hotplug scripts
- tools/pygrub
  - Grub-like bootloader for xen to boot DomU images
- tools/blktap
  - User space part for blktap driver

Kernel

- Backend driver
  - drivers/xen/
  - blkback, netback, etc.
- Frontend driver
  - drivers/xen/
  - blkfront, netfront, etc.

Xen split drivers

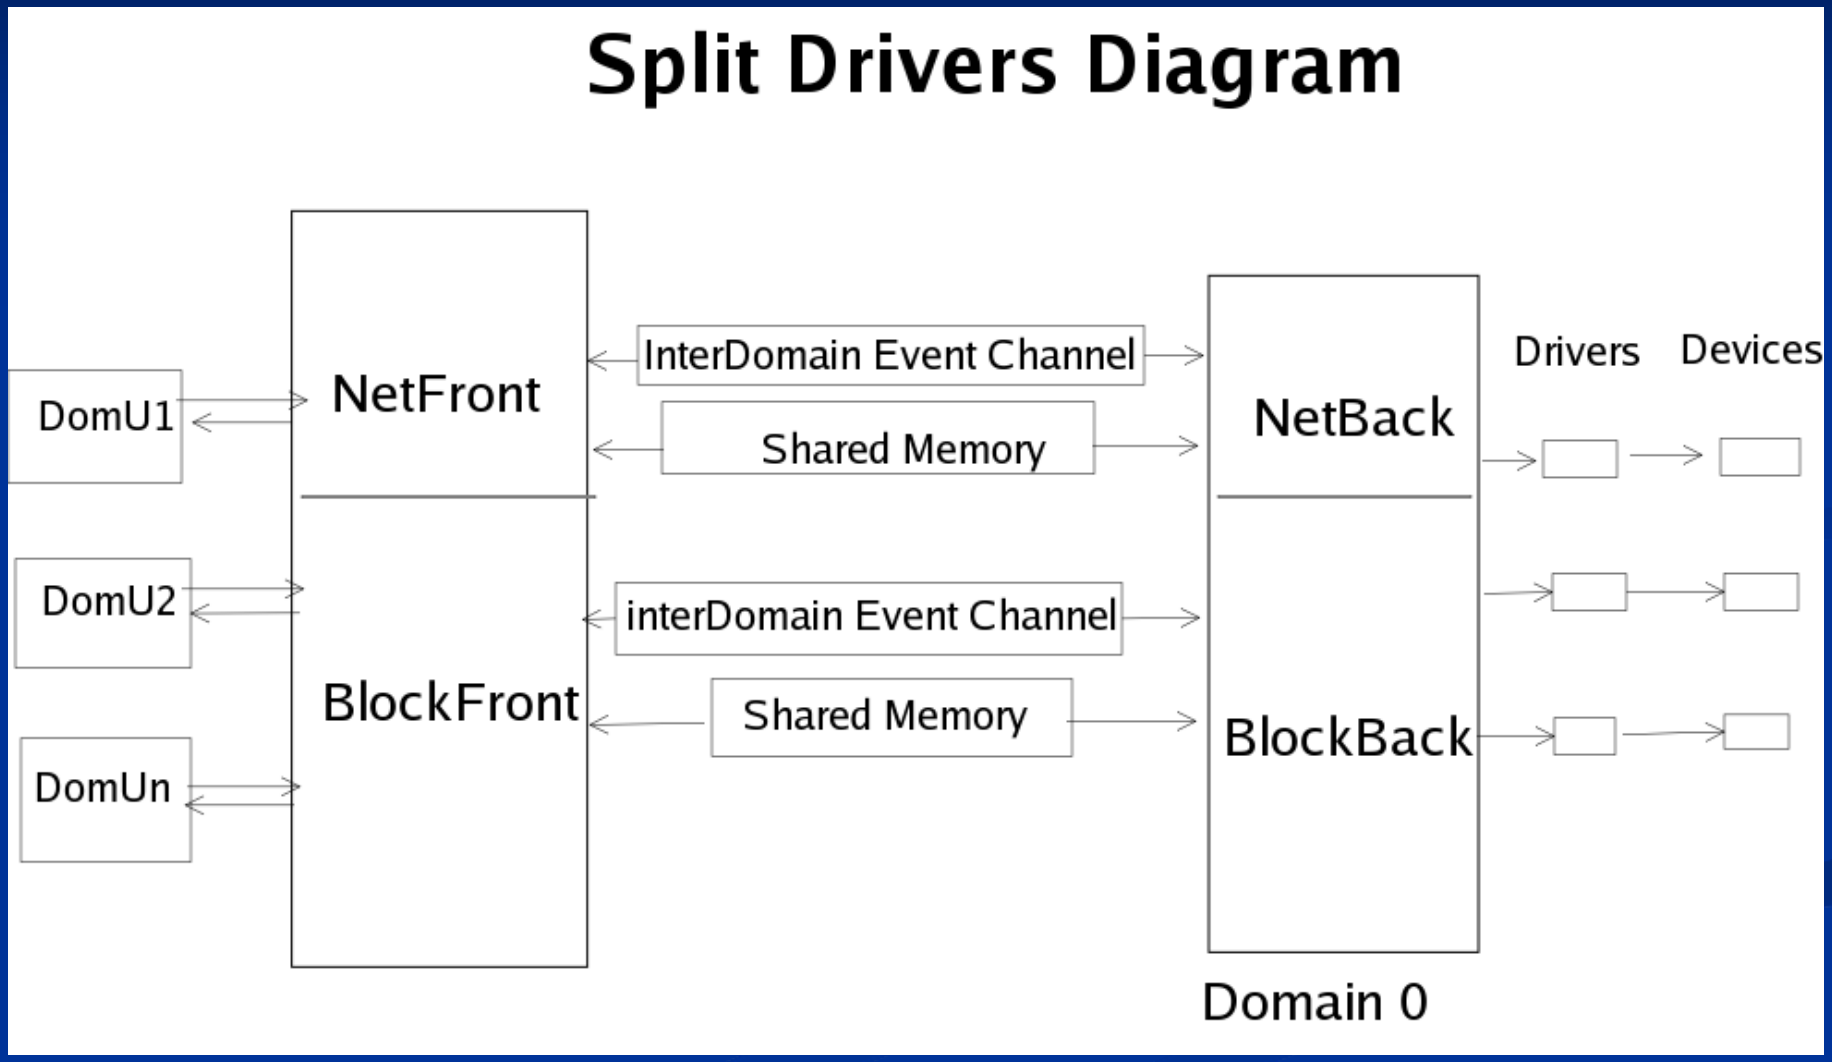

Graph from: http://wiki.xensource.com/xenwiki/XenIntro

Xen split drivers

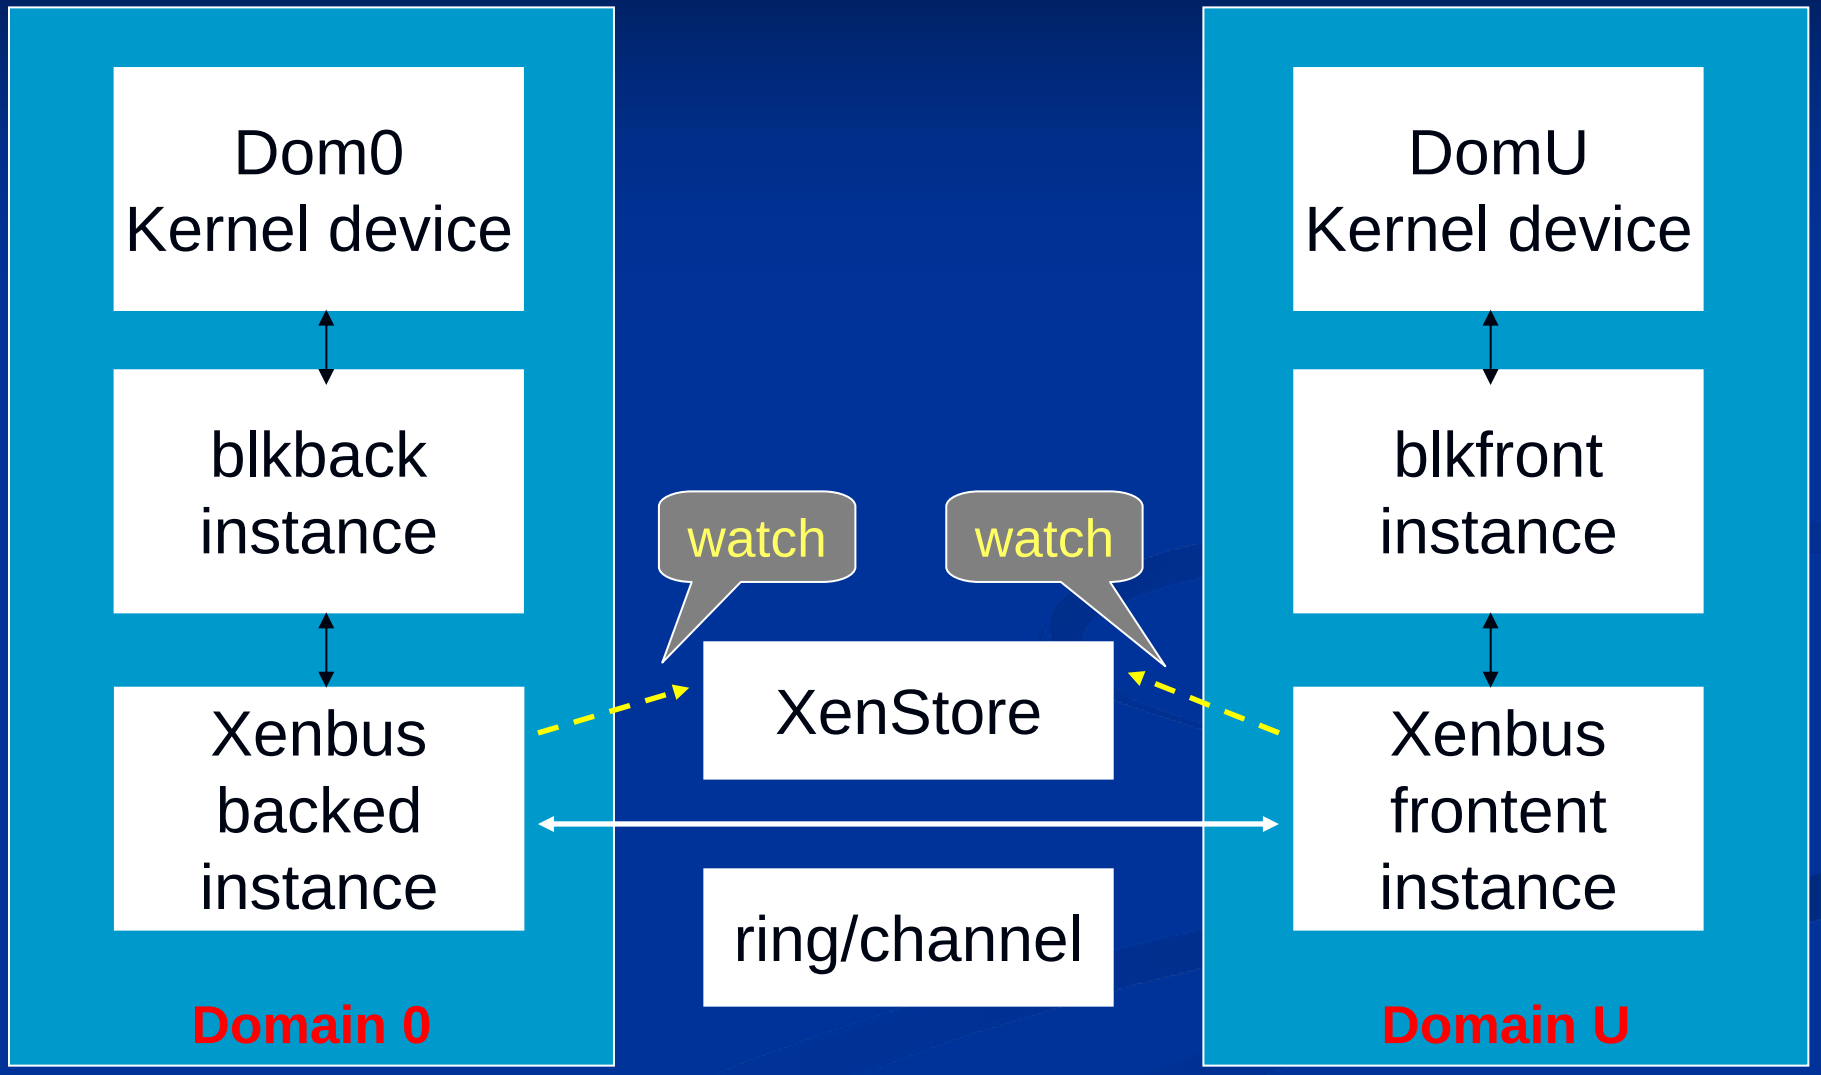

Device initialization

1. Xend (or another tool) writes frontend and backend details to the store
   - `/local/domain/0/backend/vbd/U/<deviceID>/...`
   - `frontend /local/domain/U/device/vbd/<deviceID>`
   - `frontend-id U`
   - `state XenbusStateInitialising`
   - `... <device-specific details>`
   - `/local/domain/U/device/vbd/<deviceID>/...`
   - `backend /local/domain/0/backend/vbd/U/<deviceID>`
   - `backend-id 0`
   - `state XenbusStateInitialising`
   - `... <device-specific details>`
2. Watches fire, and the Xenbus instances begin negotiation.

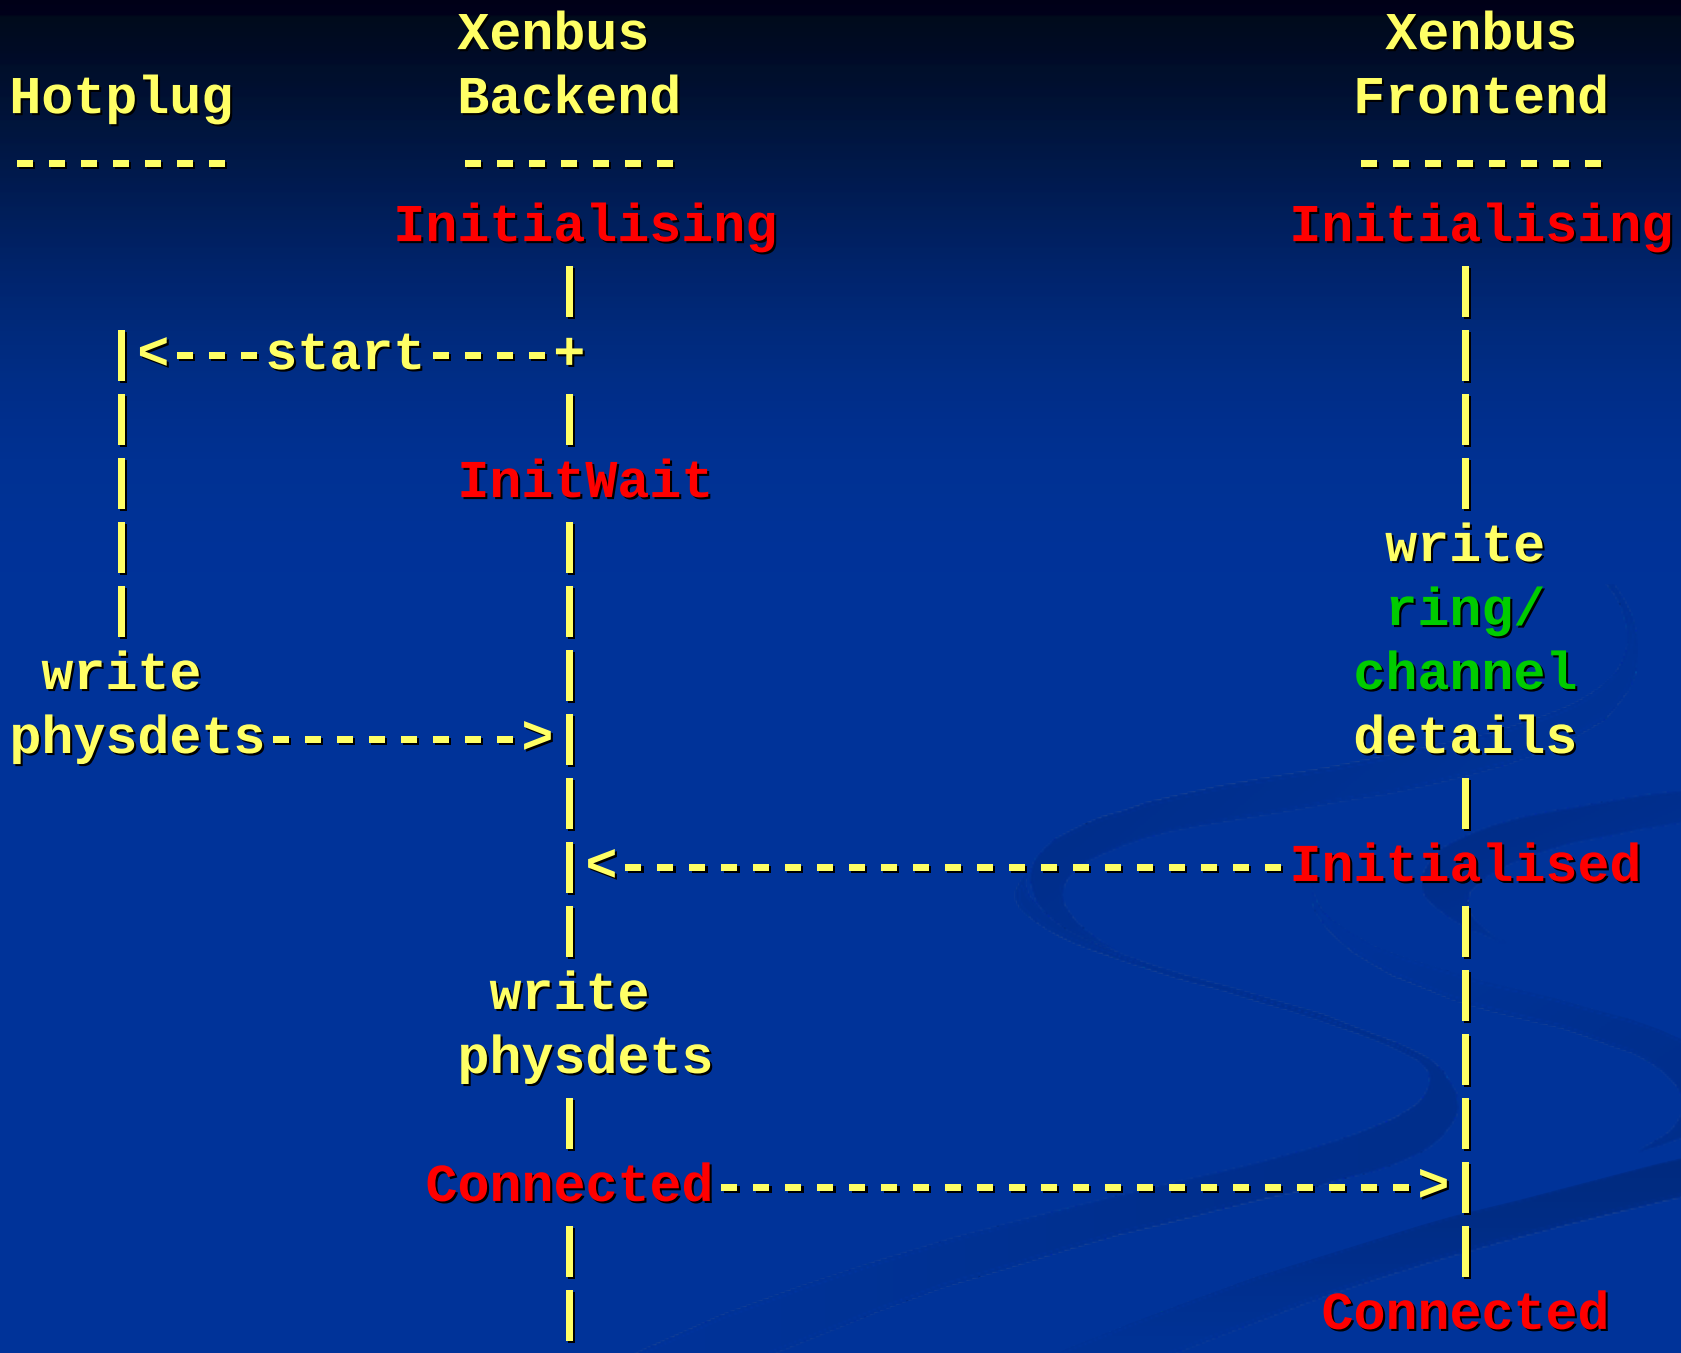

Graph from: http://wiki.xensource.com/xenwiki/XenSplitDrivers

blktap

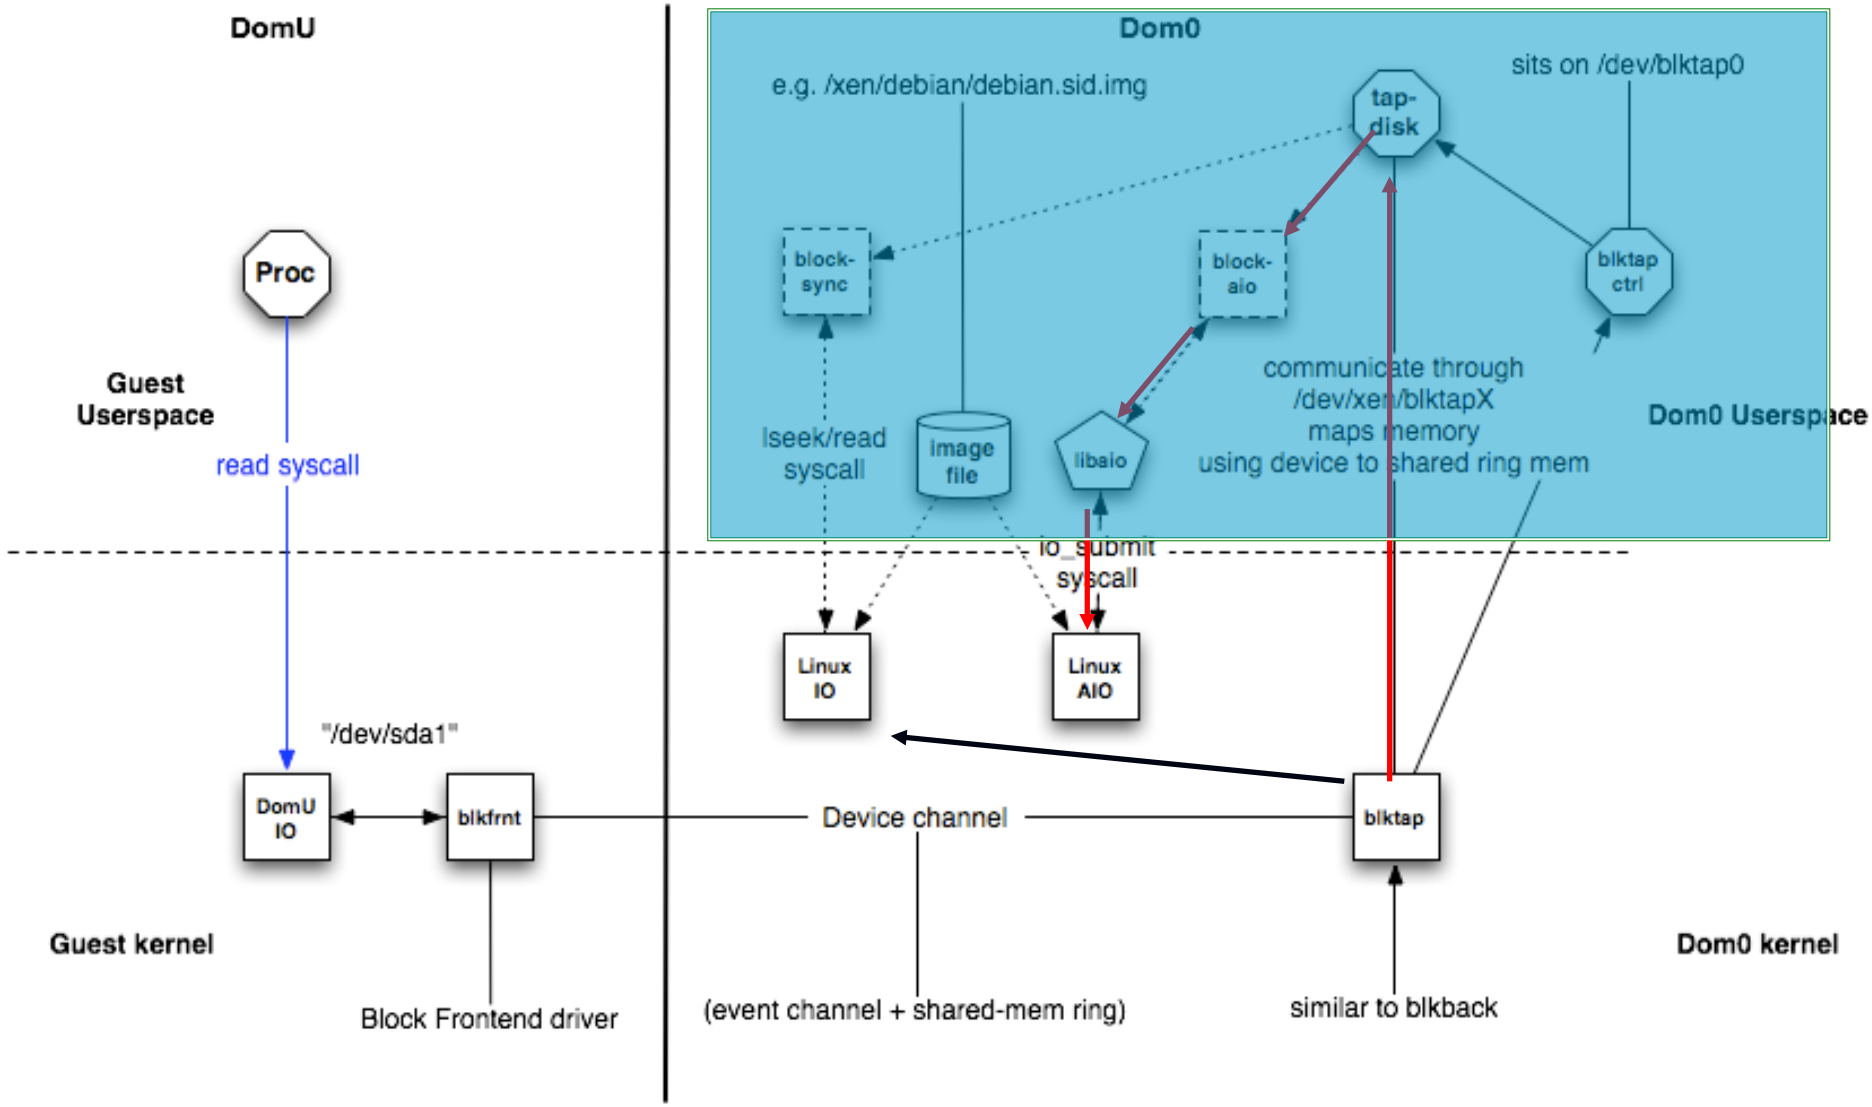

Graph from: http://wiki.xensource.com/xenwiki/blktap

Thank you!