

Using IOMMUs for Virtualization in Linux and Xen

Muli Ben-Yahuda, [Jon Mason](#), Orran Krieger, Jimi Xenidis, Leendert Van Dorn, Asit Mallick, Jun Nakajima, Elsia Wahlig

jdmason@us.ibm.com

Where in the world is
Muli Ben-Yehuda?

Overview

- What is an IOMMU?
- Current Implementations
- Linux and the DMA-API
- Xen
- Road map
- Status
- Questions

What is an IOMMU?

An I/O Memory Management Unit (IOMMU) is a hardware component that provides two main functions: I/O Translation and Device Isolation.

- The IOMMU translates memory addresses presented by devices from “I/O space” to “machine space” to allow a particular device to access physical memory potentially out of its range. It does this by providing an “in range” address to the device and translating the “in range” address to the physical memory address on the fly.
- The same translation function, when coupled with access permissions (“who can access this memory?”) can limit the ability of devices to access specific regions of memory.

Why do we need an IOMMU?

- Pros
 - 32bit DMA capable, non-DAC devices cannot access physical memory addresses higher than 4GB.
 - IOMMUs can be programmed so that the memory region appears to be contiguous to the device on the bus (scatter/gather coalescing).
 - Device Isolation and other RAS features.
- Cons
 - TANSTAAFL, “there ain't no such thing as a free lunch”.
Remapping adds a performance hit to the transfer, albeit one which can be mitigated by an IOTLB.

Isolation – The Main Advantage

- For device isolation, an I/O translation needs to be available to a given device, but not some other device.
- We also need to restrict which domain can program which device – this is done by the hypervisor by controlling which domain can program which adapter.

Unfortunately, only one IOMMU available today on x86_64 based servers can do isolation, Calgary.

SWIOTLB

- Linux includes swiotlb which is a software implementation of the translation function of an IOMMU. Also known as “bounce buffers”.
- Linux always uses swiotlb on IA64 systems which have no hardware IOMMU, and can use it on x86-64 when told to do so or when a system has too much memory and not enough IOMMU.
- As of release 3.0.0, Xen always uses a modified version of swiotlb in dom0, since swiotlb provides machine contiguous chunks of memory (required for DMA) unlike the rest of the kernel memory allocation APIs when running under Xen.
- Using swiotlb (or any other IOMMU) is completely transparent to the drivers – everything is implemented in the architecture's DMA mapping API implementation.

AMD GART

- AMD Graphical Aperture Remapping Table (GART) provides a basic, translation only, IOMMU.
- Implemented in the on-chip memory controller
- Physical memory window and list of pages to be translated.
- Addresses outside the window are not translated.
- Fully supported in Linux; Xen support posted but not in main tree.

AMD IOV Technology and Intel VT-d

- Provides translation and isolation
- Devices are assigned into a protection domain with a set of I/O page tables defining the allowed memory addresses.
- Before a DMA transfer begins, the IOMMU intercepts the access and checks its cache (IOTLB) and (if necessary) the I/O page tables for that device, based on the devices Bus/Dev/Func.
- Can be arranged in a topology of IOMMUs
- I/O page tables maintained in system memory by host software; with AMD's implementation, the page table format is compatible with the MMU's page table format.

Calgary TCEs

- Calgary's Translation Control Entries (TCEs) provide functionality to translate and isolate.
- Calgary provides a unique I/O address space up to 4GB in size to all devices behind each PCI Host Bridge (PHB).
- Calgary uses the DMA address as an index into its IOTLB. If a translation is not found in the IOTLB, the address is used as an index into a system controlled translation table in memory.
- Calgary is found in some of IBM's System P and System X servers.

Linux Status

- Linux has a standard API for dealing with DMA memory which all well written drivers are already using, the DMA API.
- First we cleaned up the x86-64 DMA-API implementation to support more than nommu, swiotlb, and gart cleanly – the dma-ops patch.
- Then we did Calgary bring up on bare metal Linux
- And implemented the DMA-API for Calgary on the server formerly known as “xSeries x366”.
- Despite the hardware having never been validated, it actually works.
- We had to work around a few oddities creatively

Linux, continued

- Calgary merged in 2.6.18-rc1
- This provides an isolation capable IOMMU on Intel based servers – get your DMA handling wrong and the DMA will be stopped by the IOMMU with an informative message rather than corrupting memory!
- Some open issues: direct userspace access (i.e., X), graceful handling and recovery of driver errors, better integration with swiotlb, NUMA enhancements, etc. etc...

Direct Hardware Access

One of the main selling points of virtualization is machine consolidation. So let's assume for a second that you put your database virtual machine and your web server virtual machine on the same physical machine. Your database needs fast disk access; your web server, fast network access.

Xen supports the ability to allocate different physical devices to different domains (multiple “driver domains”). However, due to architectural limitations of most PC hardware, this cannot be done securely. In effect, any domain that has direct hardware access has to be considered “trusted”.

The Problem with Direct Access

The reason why is that all I/O is done in physical addresses.

Consider the following case:

- Domain A is mapped in 0-2GB of physical memory
- Domain B is mapped in 2-4GB of physical
- Domain A has direct access to a PCI NIC
- Domain A programs the NIC to DMA in the 2-4GB physical memory range, overwriting domain B's memory.
Oops!

The solution – an IOMMU!

Xen

- Main goal: using Calgary to provide direct access to devices from multiple driver domains.
 - Almost there – dom0 running with Calgary enabled, with Xen programming the IOMMU.
- Working on getting another driver domain running with Calgary enabled as well.
- Dom0 detect Calgary in the machine – notifies the hypervisor which initializes Calgary support.
- New privileged hypercalls: iommu detected, create and destroy I/O space. I/O spaces are identified by BDF or parts of BDF.
- New hypercalls: map and unmap translation entry in I/O space.

Road map

- Integration with Xen's grant tables and PCI frontend and backend drivers
- Support for Intel and AMD's upcoming IOMMUs
- Migration?
- Fully virtualized OS's with direct device access?
- Performance...
- How do we build better IOMMUs?

Status and Availability

- Linux code merged into mainline: 2.6.18-rc1 should have it all.
- Xen trees available on xenbits:
 - <http://xenbits.xensource.com/ext/xen-iommu.hg>
 - <http://xenbits.xensource.com/ext/linux-iommu.hg>

Questions?