

# DMA-IOMMU for VT-d

# The purpose of this talk?

We want to replace the Intel IOMMU driver's IOVA (IO Virtual address) handling code with generic code already used by other drivers

To do this we need to go over the issues preventing us merging this replacement.

# Some context

Previously each IOMMU driver had their own implementation for IO virtual addresses handling (allocating, freeing and finding free IO virtual address ranges)

Now Arm and AMD drivers now use the same generic path. We want intel to use this path too.

We call this generic path the dma-iommu path. It's implemented in "drivers/iommu/dma-iommu.c"

# Why replace the existing code?

The logical for the driver specific code and the genetic code is pretty much identical but just different in ways which make it difficult for developers

It allows us to remove a lot of code from the intel driver.

```
drivers/iommu/Kconfig      | 1 +  
drivers/iommu/intel-iommu.c | 742 +++-----  
include/linux/intel-iommu.h | 1 -  
3 files changed, 55 insertions(+), 689 deletions(-)
```

Problems that  
need solving

# Bounce buffer code

We need to port the Intel IOMMU implementation of bounce buffers for untrusted devices to the generic dma-iommu path

[<https://lkml.org/lkml/2019/8/30/178>]

As well as the trace events for IOMMU map/unmap

[<https://lkml.org/lkml/2019/8/30/175>]

# i915 gpu driver ignores map\_sg return value

The i915 driver ignores the number of elements returned by map\_sg() and instead relies on the list ending in NULL.

For example:

```
/* iterate over the pages of the given sg_table */  
#define for_each_sgt_page(__pp, __iter, __sgt)
```


```
/* iterate over the device addresses of the given sg_table */  
#define __for_each_sgt_daddr(__dp, __iter, __sgt, __step)
```

[\[https://github.com/torvalds/linux/blob/d012a7190fc1fd72ed48911e77ca97ba4521bccd/drivers/gpu/drm/i915/i915\\_scatterlist.h#L76\]](https://github.com/torvalds/linux/blob/d012a7190fc1fd72ed48911e77ca97ba4521bccd/drivers/gpu/drm/i915/i915_scatterlist.h#L76)

# I915 gpu driver ignores map\_sg return value

By ignoring the map\_sg return value we get memory corruption which causes the driver to display artifacts when we switch to the generic dma-iommu path:

<https://patchwork.kernel.org/cover/11306999/>





Let's chat