

Per group default domain type

Lu Baolu
Intel Corporation

Agenda

- The problems we are suffering.
- How could we implement it.
- What's the user API .

Problems we are suffering

- A global `iommu_def_domain_type` determined by kernel configuration or boot parameters.
- Problem 1 - default domain type is DMA while the users might want IDENTITY for some fast-speed device for performance gains.
- Problem 2 - default domain type is IDENTITY while users might want DMA for some legacy 32-bit devices.
- More and more users are looking forward to a fine grained default domain type. Per group default domain type is proposed to address this.

How we do it? (discussion point)

- Current boot sequence

For each device on bus

Allocate group

Allocate default domain (if no)

Attach domain to device

IOMMU drivers request
different domain type if attach
failure

- Per-group domain type

Parse and save user specified
mappings between device and
domain type.

For each device on bus

Allocate group and determine
the domain type (driver > user
> global)

For each device on bus

Allocate default domain

Attach domain to device

User API (discussion point)

- `iommu.dma =`
Usage: Specify devices that want IOMMU translation
Format: [`<domain>:`]`<bus>:<slot>.<func>`[; ...]
- `iommu.identity =`
Usage: Specify devices that bypass IOMMU translation
Format: [`<domain>:`]`<bus>:<slot>.<func>`[; ...]

INTEL OPEN SOURCE TECHNOLOGY CENTER | 01.org

 @twitter handle