

LoongArch I/O Virtualization Table Specification

Loongson Technology Corporation Limited

Version v0.1 , 2024-10-08

Revision History

Revision	Date	Type	Description
0.1	2024.10.08	A	Initial draft

*Types are divided into **A** - ADDED **M** - MODIFIED **D** - DEL

Contents

Revision History	II
Terms and Abbreviations.....	2
1. Introduction.....	3
2. LoongArch I/O Virtualization Table (IOVT).....	3
2.1. I/O Virtualization Table (IOVT).....	3
2.2. IOMMU Structure.....	3
2.3. Device Entry	5

Terms and Abbreviations

This specification uses the following terms and abbreviations:

Term	Meaning
ACPI	Advanced Configuration and Power Interface
IOMMU	Input Output Memory Management Unit

1. Introduction

This specification describes the I/O Virtualization Table (IOVT) on LoongArch platform, which provides information about LoongArch IOMMU and the I/O topology relevant to each IOMMU. The IOVT describes the configuration and capabilities of the IOMMUs present on the platform, along with information regarding the devices which each IOMMU manages.

2. LoongArch I/O Virtualization Table (IOVT)

2.1. I/O Virtualization Table (IOVT)

Table 1 shows the structure of IOVT. Bytes 0–32 are ACPI basic header information. The subsequent bytes are designed based on LoongArch IOMMU.

Table 1. LoongArch I/O Virtualization Table

Field	Byte Offset	Byte Length	Description
Signature	0	4	“IOVT” , I/O Virtualization Table.
Length	4	4	Length in bytes of the entire IOVT, including IOMMU structures.
Revision	8	1	1
Checksum	9	1	Checksum of entire structure must equal zero.
OEM ID	10	6	OEM ID.
OEM Table ID	16	8	For the IOVT, the table ID is the manufacture model ID.
OEM Revision	24	4	OEM revision of the IOVT for the supplied OEM Table ID.
Creator ID	28	4	The vendor ID of the utility that created the table.
Creator Revision	32	4	The revision of the utility that created the table.
IOMMU Count	36	2	Number of structures in the IOMMU structures.
IOMMU Offset	38	2	The offset from start of this table to the first structure in IOVT IOMMU structures.
Reserved	40	8	Must be zero.
IOMMU Structures[]	48	--	A list of structures. The list will contain one or more IOMMU structures.

2.2. IOMMU Structure

Each IOMMU structure describes an IOMMU device and reports its capabilities and configuration. LoongArch IOMMU can be implemented as a PCI device or a platform device. Table 2 shows the format of the IOMMU structure.

Table 2. IOMMU Structure

Field	Byte Offset	Byte Length	Description
Type	0	2	0: LoongArch IOMMUv1.
Length	2	2	The length of this structure in bytes, including device entries.
Flags	4	4	<p>IOMMU flags.</p> <ul style="list-style-type: none"> • Bit 0: PCI device flag. <ul style="list-style-type: none"> 1: The IOMMU is a <code>_PCI</code> device. 0: The IOMMU is a <code>_platform</code> device. • Bit 1: Proximity domain valid. • Bit 2: Whether the devices managed by the IOMMU can be classified by PCI segment. <ul style="list-style-type: none"> 1: All devices under the root PCI bridge can be managed. 0: Manageable devices are placed in the device entry list. • Bit3: Hardware capability support. • Bit4: MSI interrupt address bypass supported. • Bit5–31: Reserved, must be zero.
PCI Segment Number	8	2	The PCI segment group number.
Physical Address Width	10	2	The width of the physical address supported by the IOMMU.
Virtual Address Width	12	2	The width of the virtual address supported by the IOMMU.
Max Page Level	14	2	The max levels of page table supported by the IOMMU.
Page Size Supported	16	8	Indicates the page size that the TLB can support. When bit i is 1, it indicates that pages of 2^i bytes are supported.
IOMMU DeviceID	24	4	DeviceID of IOMMU. This field is valid only when the “PCI device flag” is set.
IOMMU Base Address	28	8	Base address of the IOMMU registers. This field is valid only when the “PCI device flag” is not set.
IOMMU Register Size	36	4	Size of IOMMU registers in bytes.
Interrupt Type	40	1	IOMMU page table exception interrupt type.
Reserved	41	3	Must be zero.
Global System Interrupt	44	4	Global System Interrupt (GSI) number of IOMMU page table exception interrupt. The exceptions include no read permission, no write

			permission, and an invalid page table. This field is valid only when the “PCI device flag” is not set.
Proximity Domain	48	4	If the “proximity domain valid” flag is set, this field indicates the proximity domain from which the IOMMU page table is allocated in the NUMA system.
Max Device Num	52	4	The maximum number of devices that the IOMMU structure can manage.
Number of Device Entries	56	4	Number of device entries.
Offset of Device Entries	60	4	The offset from the start of this IOMMU structure to the first device entry.
Device Entries[]	64	--	List of device entries that can be managed by the IOMMU structure.

2.3. Device Entry

Each device entry describes the devices that IOMMU can manage. Currently, only PCI devices are supported. Each device entry may be used to indicate a single PCI device or a range of devices. If wanting to indicate a range of devices, at least two device entries are needed to describe the start and end of that range. Table 3 shows the format of device entries.

Table 3. Device Entry

Field	Byte Offset	Byte Length	Description
Type	0	1	0: A single PCI device. 1: Start of range. 2: End of range. 3–255: Reserved.
Length	1	1	The length of this device entry in bytes.
Flags	2	1	Reserved.
Reserved	3	3	Must be zero.
DevID	6	2	DeviceID.