

Pascal MMU Format Changes:

Highlights:

- Expanded Virtual Addressing - Upto 49 Bits of VA.
- Expanded physical addressing for system memory – Up to 47 bits of system PA.
- Support for 2MB big pages.
- Dropped support for 128KB Big Pages.
- 5 Level Format to support the expanded virtual addressing.
- All page levels/tables are now 4KB in size.

Page Table Format:

PDE1 and up																																																43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Physical Address																																				VOL		A		V										A=00		0		Not Valid																																					
0																																				x		x		x										A=01		0		Vid Mem PDE																																					
See PTE Formats																																				x		x		x										A=1x		0		Sys Mem PDE																																					
See PTE Formats																																				1		1		1										1		1		PTE																																					
PDE0																																																43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Physical Address																																				VOL		A		V										AB=00		0		Not Valid																																					
0																																				x		x		x										AB=01		0		Vid Mem PDE																																					
See PTE Formats																																				x		x		x										AB=1x		0		Sys Mem PDE																																					
See PTE Formats																																				1		1		1										1		1		PTE																																					
Physical Address																																				VOL		A		V										AS=00		0		Not Valid																																					
0																																				x		x		x										AS=01		0		Vid Mem PDE																																					
See PTE Formats																																				x		x		x										AS=1x		0		Sys Mem PDE																																					
See PTE Formats																																				1		1		1										1		1		PTE																																					
Unused																																																																																											
PTE																																																43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Kind												CompTag Line												Physical Address																								RO		E		VOL		A		V																																			
x												x												x																								x		x		x		x		x																																			
K[7:0]												CTL[17:0]												PA[36:12]																								ad		ro		P		E		x		A=00		0		Not Valid																													
K[7:0]												CTL[17:0]												PA[36:12]																								ad		ro		P		E		x		A=01		1		Vid Mem PTE																													
K[7:0]												CTL[17:0]												PA[57:12]																								ad		ro		P		E		x		A=1x		1		Sys Mem PTE																													

PDE0 Fields AB = Aperture for big page page table VOLB = Volatile for big page page table in System memory PAB = Physical Address of big page page table AS = Aperture for small page page table VOLs = Volatile for small page page table in System memory PAS = Physical Address of small page page table	PTE Fields V = Valid A = Aperture for pages VOL = Volatile page PA = Physical Address of page PEER = Peer ID CTL = CompTag Line E = Encrypted R = Privilege bit RO = Read-only page AD = Atomic Disable K = Kind	PDE Aperture Decode: 00 Invalid 01 Local Video memory 10 System memory, cache coherent 11 System memory, not cache coherent PTE Aperture Decode: 00 Local Video memory 01 Peer Video Memory 10 System memory, cache coherent 11 System memory, not cache coherent
--	--	--

Page Level Sizes and Coverage:

Page Type		4KB	64KB	2MB
Page	Size	4KB	64KB	2MB
	Alignment	4KB	64KB	2MB
Page Table	Number of Entries	512	32	Uses PD0
	Size	4KB	256B Multiple page tables can be packed into one 4K page	
	Coverage	2MB	2MB	
Page Directory0	Number of Entries	256	256	256
	Number of bits	8	8	8
	Size	4KB	4KB	4KB
	Coverage	512MB	512MB	512MB
Page Directory1	Number of Entries	512	512	512
	Number of Bits	9	9	9
	Size	4KB	4KB	4KB
	Coverage	256GB	256GB	256GB
Page Directory2	Number of Entries	512	512	512
	Number of Bits	9	9	9
	Size	4KB	4KB	4KB
	Coverage	128TB	128TB	128TB
Page Directory3	Number of Entries	4	4	4
	Number of Bits	2	2	2
	Size	4KB	4KB	4KB
	Coverage	512TB	512TB	512TB

49 Bit VA breakdown:

PD3 [48:47]	PD2 [46:38]	PD1 [37:29]	PD0 [28:21]	2MB PAGE [0:20]	
				PT (64K) [20:16]	64k Page [15:0]
				PT (4K) [20:12]	4k Page [11:0]

Page Level structure:

Upper Page Directory Levels(PDE1 and above)

Field	Mnemonic	Width	Semantics
Valid	V	1	Indicates that the entry is a PTE. Should be 0 for all levels.

Aperture	A[1:0]	2	Aperture for the next level: 0 = invalid 1 = local video memory 2 = coherent system memory 3 = non-coherent system memory. Setting this field to 0 (invalid) indicates that the next level is not valid. Except for VOL, the remaining fields are undefined.
Physical Address	PA[39:12] PA[57:12]	28 46	Pointer to physical memory containing the next level. The five level page table format has 46 bits for system addresses. The format supports a 58 bit physical address space for system. In the PTE, the upper 18 bits of the physical address share the bits and CompTag Line. Although the page table support 58 bits, the hardware supports 47 bits of system addresses.
Volatile	VOL	1	Accesses to the next level are volatile. If Valid and Aperture are 0, and Volatile is 1, then the PDE is sparse.
Total		50	

Lower Page Directory (PDE0):

Field	Mnemonic	Width	Semantics
Valid	V	1	If valid is 0, the entry is a PDE, and the following fields apply. If valid is 1, the entry is a 2MB PTE, and the PTE fields apply.
Aperture (Big pages)	AB[1:0]	2	Aperture for the big page page table: 0 = invalid 1 = local video memory 2 = coherent system memory 3 = non-coherent system memory. Setting this field to 0 (invalid) indicates that the big page page table is not valid. Except for VOLB, the remaining big page fields are undefined.
Physical Address (Big pages)	PAB[39:12] PAB[57:12]	28 46	Pointer to physical memory containing the big page page table. The five level format has 46 bits for system addresses, but only 35 bits are used by the hardware.
Aperture (Small pages)	AS[1:0]	2	Aperture for the small page page table: 0 = invalid 1 = local video memory 2 = coherent system memory 3 = non-coherent system memory. Setting this field to 0 (invalid) indicates that the small page page table is not valid. Except for VOLB, the remaining small page fields are undefined.
Volatile (Small pages)	VOLS	1	Small page page table accesses are volatile.
Volatile (Big pages)	VOLB	1	Big page page table accesses are volatile. If Valid and both Aperture fields are 0, and Volatile is high, then the PDE is sparse.
Physical Address (Small pages)	PAS[39:12] PAS[57:12]	28 46	Pointer to physical memory containing the small page page table. The five level page table format has 46 bits for system addresses, but only 35 bits are used.
Total		64	64 bits for the two level format and 99 bits for the five level format

		99	
--	--	----	--

Page Table Format:

Field	Mnemonic	Width	Semantics
Valid	V	1	Indicates if the PTE is valid. If the PTE is not valid, all other fields, except Privileged and Volatile, are ignored.
Privileged	P	1	Setting this bit causes non-privileged requests to fault when accessing this page. Setting this bit in a invalid big page PTE indicates that no small page PTEs in this range are valid.
Read-Only	RO	1	Setting this bit causes writes to this page to fault.
Atomic-Disable	AD	1	Setting this bit disables atomic accesses.
Encryption	E	1	Do not use. Set to 0
Physical Address	PA[39:12] PA[57:12]	28 46	Pointer to the physical address of the base of the page. PA[39:37] is shared with PEER[2:0]. The five level page table format has 46 bits for system addresses, but only 35 bits are used. PA[57:40] is shared with CTL[17:0].
Peer ID	PEER[2:0]	3	Indicates the target peer when the aperture is peer. Shared with PA[39:37].
Aperture	A[1:0]	2	Aperture for pages: 0 = local video memory 1 = peer video memory 2 = coherent system memory 3 = non-coherent system memory
Volatile	VOL	1	Page accesses are volatile. If Valid is 0 and Volatile is 1, the PTE is sparse.
Kind	K[7:0]	8	Compression info
CompTag Line	CTL[17:0]	18	Indicates comp tag number. Shared with PA[57:40].
Total		62	

Details:

VOL:

The VOL bit indicates that the data for this page or page table is volatile and should not be cached in the L2 cache. This bit applies to system memory and peer apertures. On regular GPU's, L2 always caches local video memory pages, this bit is ignored when the aperture is local video memory. On Tegra chips, VOL applies to all apertures. Pages must be flushed from L2 when their cache policy changes.

VOL has meaning for invalid PDEs and PTEs. If the PDE or PTE is invalid and VOL is high, the MMU treats the range mapped by the PDE or PTE as sparse. The MMU redirects accesses to sparse ranges to one of the two SM debug dummy pages, unless the client is TEX or PROP. Redirection happens even if SM debug mode is disabled. The MMU returns an ACK to naïve clients. It returns a silent NACK to clients that are aware of sparse accesses. The lowest level PDE has two pointers, one to map small pages and one to map big pages. The big page VOL bit determines sparseness when both pointers are not valid. The lowest level PDE cannot be sparse if either or both pointers are valid.

Atomic Disable:

The Atomic Disable bit disables atomic requests to the page. An atomic request to a page with Atomic Disable high will fault. This bit is used to grant exclusive access to a page by revoking the mappings to this page as part of handling the atomic fault.