Cuda Tensor Core

[illegible]

[illegible] GPU[illegible]

- WGMMA[illegible]
- [illegible]
- [illegible]
- [illegible]

[illegible] GPU[illegible] wmma.mma.sync[illegible] mma.sync[illegible] SM[illegible] 32[illegible] warp [illegible]

wgmma.mma_async[illegible] 128[illegible] SM [illegible] warp [illegible]

[illegible] L2 [illegible]
[illegible] wgmma.mma_async[illegible]

[illegible] wgmma[illegible] bank conflicts[illegible] TMA[illegible] L2[illegible] flash attention [illegible] bank conflicts[illegible]

TMA [illegible]
[illegible] pipelines[illegible]

	Volta	Turing	Ampere	Hopper
Base Size*	64[illegible] F16(4x4x4)	64[illegible] F16(4x4x4)	256[illegible] F16(8x4x8)	512[illegible] F16(8x4x16)
tensorCore[illegible]	672(8/SM)	576(8/SM)	512(4/SM)	576(4/SM)
SM[illegible] *	512[illegible] F16	512[illegible] F16	1024[illegible] F16	2048[illegible] F16
[illegible] FOPS(F16)	125T	130.5T	312T	1000T
boost clock(MHz)	1530	1455	1410	≈A100x1.3
FP64	-	-	√	√
TF32	-	-	√	√
BF16	-	-	√	√
FP16	√	√	√	√

	Volta	Turing	Ampere	Hopper
FP8	-	-	-	√
INT8/UINT8	-	√	√	√
INT4/UINT4	-	√	√	-
INT1	-	√	√	-

	Volta	Turing	Ampere	Hopper
SM/TPC	2	2	2	2
process blocks/SM	4	4	4	4
FP64/PB	8	-	8	16
FP32/PB	16	16	16	32
INT32/PB	16	16	16	16
tensorCore/PB	2	2	1	1
LSU/PB	8	4	8	8
register file	64KB(16384*32bit)	64KB(16384*32bit)	64KB(16384*32bit)	64KB(16384*32bit)
L0 ICache/PB	1	1	1	1
L1/SHM/SM	128KB	96KB	192KB	256KB
warp scheduler/PB	1(32thread/clk)	1(32thread/clk)	1(32thread/clk)	1(32thread/clk)
dispatch uint/PB	1(32thread/clk)	1(32thread/clk)	1(32thread/clk)	1(32thread/clk)

[![image.png](Cuda Tensor Core/Ubsimage-png.png)](Cuda Tensor Core/Ubsimage-png.png)

tensorCore□□□□□□ SM□□ thread□ warp□□□□□□□□ thread
□□□□□□□□□□□□□□□□□□□□□□□□□□□□ SM□□□ thread□ tensorCore□□□□□□□

Revision #2
Created 2025-01-11 09:46:27 UTC by Colin
Updated 2026-04-29 07:33:23 UTC by Colin