

# PowerEdge Server GPU Matrix

Brand	Model	GPU Memory	Memory ECC	Memory Bandwidth	Max Power Consumption	Host Interface (BW)	GPU-to-GPU Peer Interconnect (BW)	Slot Width	GPU Height/Length <sup>5</sup>	Auxiliary Cable	Workload <sup>1</sup>
<a href="#">AMD</a>	MI300X (x8) OAM	192 GB HBM3	Y	5.3 TB/sec	750W	AMD Infinity Fabric Links	896 GB/sec	N/A	N/A	N/A	AI / HPC
<a href="#">AMD</a>	MI210	64 GB HBM2e	Y	1638 GB/sec	300W	PCIe Gen4x16/Infinity Fabric Link bridge <sup>8</sup>	64 GB/sec (PCIe 4.0)	DW	FHFL	CPU 8 pin	HPC/Machine learning training
<a href="#">Intel</a>	Gaudi3 PCIe	128GB HBM3	Y	3.7 TB/s	600W	PCIe Gen5x16 (128 GB/sec)	3x200G RoCE	DW	FHFL		AI / HPC
<a href="#">Intel</a>	Gaudi3 OAM	128GB HBM3	Y	3.6TB/s	850W	PCIe Gen5x16 (128 GB/sec)	3x200G RoCE	N/A	N/A	N/A	AI / HPC
<a href="#">Intel</a>	Max1550 OAM	128 GB HBM2e	Y	3276.8 GB/sec	600W	PCIe Gen5x16 (128 GB/sec)	Intel XeLink	N/A	N/A	N/A	AI / HPC
<a href="#">Intel</a>	Max1100	48 GB HBM2e	Y	1228.8 GB/sec	300W	PCIe Gen5x16 (128 GB/sec)	XeLink bridge <sup>8</sup> (600GB/s)	DW	FHFL	PCIe 16 pin	AI / HPC
<a href="#">Intel</a>	Flex 140	12 GB GDDR6	Y	336 GB/Sec	75W	PCIe Gen4 x8 (32 GB/sec)	XeLink bridge <sup>8</sup> (336 GB/sec)	SW	HHHL/FHHL	N/A	Inferencing/Edge
<a href="#">Nvidia</a>	H200 NVL	141GB HBM3e	Y	4.8 TB/s	600W	PCIe Gen 5x16 (128 GB/sec)	NVLink (4th Gen - 900 GB/s)	DW	FHFL	PCIe 16 pin	AI / HPC
<a href="#">Nvidia</a>	H100 NVL	94 GB HBM3	Y	3.9 TB/s	350W- 400W	PCIe Gen 5x16	600GB/s (PCIe 5.0)	DW	FHFL	PCIe 16 pin	AI / HPC
<a href="#">Nvidia</a>	H200 SXM5 (x8) DLC	141GB HBM3e	Y	4.8 TB/s	700W	PCIe Gen 5x16 (128 GB/sec)	NVLink (4th Gen - 900 GB/s)	DW	N/A	N/A	AI / HPC
<a href="#">Nvidia</a>	H200 SXM5 (x8)	141GB HBM3e	Y	4.8 TB/s	700W	PCIe Gen 5x16 (128 GB/sec)	NVLink (4th Gen - 900 GB/s)	N/A	N/A	N/A	AI / HPC
<a href="#">Nvidia</a>	H100 SXM5 (x8)	80 GB HBM3	Y	3 TB/sec	700W	PCIe Gen 5x16 (128 GB/sec)	NVLink (4th Gen - 900 GB/s)	N/A	N/A	N/A	AI / HPC
<a href="#">Nvidia</a>	H100 SXM5 (x4)	80 GB HBM3	Y	3 TB/sec	700W	PCIe Gen 5x16 (128 GB/sec)	NVLink (4th Gen - 900 GB/s)	N/A	N/A	N/A	AI / HPC
<a href="#">Nvidia</a>	A100 SXM4 (x8)	80 GB HBM2	Y	2039 GB/sec	500W	PCIe Gen 5x16 (128 GB/sec)	NVLink (3rd Gen - 600 GB/s)	N/A	N/A	N/A	HPC/AI/ Database Analytics
<a href="#">Nvidia</a>	A100 SXM4 (x4)	80 GB HBM2	Y	2039 GB/sec	500W	PCIe Gen 5x16 (128 GB/sec)	NVLink (3rd Gen - 600 GB/s)	N/A	N/A	N/A	HPC/AI/ Database Analytics
<a href="#">Nvidia</a>	A100 SXM4 (x4)	40 GB HBM2	Y	1555 GB/sec	400W	PCIe Gen 5x16 (128 GB/sec)	NVLink (3rd Gen - 600 GB/s)	N/A	N/A	N/A	HPC/AI/ Database Analytics
<a href="#">Nvidia</a>	H100	80 GB HBM2e	Y	2000 GB/sec	300-350W	PCIe Gen 5x16 (128 GB/sec)	NVLink8 (4th Gen - 600 GB/s)	DW	FHFL	PCIe 16 pin	HPC/AI/ Database Analytics
<a href="#">Nvidia</a>	A100	80 GB HBM2e	Y	1935 GB/sec	300W	PCIe Gen4 x16 (64 GB/sec)	NVLink8 (3rd Gen - 600 GB/s)	DW	FHFL	CPU 8 pin	HPC/AI/ Database Analytics

Brand	Model	GPU Memory	Memory ECC	Memory Bandwidth	Max Power Consumption	Host Interface (BW)	GPU-to-GPU Peer Interconnect (BW)	Slot Width	GPU Height/Length <sup>5</sup>	Auxiliary Cable	Workload <sup>1</sup>
<a href="#">Nvidia</a>	L40S	48 GB GDDR6	Y	864 GB/sec	350W	PCIe Gen4 x16	64 GB/sec <sup>5</sup> (PCIe 4.0)	DW	FHFL	PCIe 16 pin	AI/ Performance graphics/VDI
<a href="#">Nvidia</a>	L40	48 GB GDDR6	Y	864 GB/sec	300W	PCIe Gen4 x16	64 GB/sec (PCIe 4.0)	DW	FHFL	PCIe 16 pin	Performance graphics/VDI
<a href="#">Nvidia</a>	A16	64 GB GDDR6	Y	800 GB/sec	250W	PCIe Gen4 x16	64 GB/sec (PCIe 4.0)	DW	FHFL	CPU 8 pin	VDI
<a href="#">Nvidia</a>	L4	24 GB GDDR6	Y	300 GB/s	72W	PCIe Gen4 x16	64 GB/sec (PCIe 4.0)	SW	HHHL	N/A	Inferencing/ Edge/VDI
<a href="#">Nvidia</a>	L4	24 GB GDDR6	Y	300 GB/s	72W	PCIe Gen4 x16	64 GB/sec (PCIe 4.0)	SW	FHHL	N/A	Inferencing/ Edge/VDI
<a href="#">Nvidia</a>	A2 (v2)	16 GB GDDR6	Y	200 GB/sec	60W	PCIe Gen4 x8	32 GB/sec (PCIe 4.0)	SW	HHHL	N/A	Inferencing/ Edge/VDI
<a href="#">Nvidia</a>	A2 (v2)	16 GB GDDR6	Y	200 GB/sec	60W	PCIe Gen4 x8	32 GB/sec (PCIe 4.0)	SW	FHHL	N/A	Inferencing/ Edge/VDI
<a href="#">Nvidia</a>	A10	24 GB GDDR6	Y	600 GB/sec	150W	PCIe Gen4 x16	64 GB/sec (PCIe 4.0)	SW	FHFL	PCIe 8 pin	mainstream graphics/VDI

Brand	Model	Part Numbers	Tied SKU <sup>2</sup>			Customer Kit SKU <sup>2</sup>	
<a href="#">AMD</a>	MI300X OAM	W9DCT	490-BJZD				
<a href="#">AMD</a>	MI210	0RJ6Y/YJFV0/880VN	490-BHUU	490-BJYD		490-BHUT	490-BJYF
<a href="#">Intel</a>	Gaudi3 PCIe						
<a href="#">Intel</a>	Gaudi3 OAM	3PNFV	489-BBGR				
<a href="#">Intel</a>	Max1550 OAM	PK79H	490-BJDX				
<a href="#">Intel</a>	Max1100	WG7J6	490-BJYL			490-BJYN	
<a href="#">Intel</a>	Flex 140	P7M0V/DHG15	490-BIIQ	490-BIIM		490-BIIK	490-BIIO
<a href="#">Nvidia</a>	H200 NVL	H80JD					
<a href="#">Nvidia</a>	H100 NVL	MM3MP	330-BCFX				
<a href="#">Nvidia</a>	H200 SXM5 (x8) DLC	RFR22					
<a href="#">Nvidia</a>	H200 SXM5 (x8)	1WPD4	490-BKHF				
<a href="#">Nvidia</a>	H100 SXM5 (x8)	W2YC4	489-BBGD				
<a href="#">Nvidia</a>	H100 SXM5 (x4)	V0W6P	489-BBGJ				
<a href="#">Nvidia</a>	A100 SXM4 (x8)	DPYN4	489-BBGC				
<a href="#">Nvidia</a>	A100 SXM4 (x4)	VT1NH/6CNP5	489-BBFS				
<a href="#">Nvidia</a>	A100 SXM4 (x4)	700T5	489-BBFT				
<a href="#">Nvidia</a>	H100	DGP4C/YR3WC	490-BIGQ	490-BJBZ		490-BIGZ	
<a href="#">Nvidia</a>	A100	WGWR4/575TK	490-BHBL	490-BHBM	490-BHTB, 490-BHUV, 490-BJKG	490-BHBK	
<a href="#">Nvidia</a>	L40S	7WK28	490-BJVT			490-BJVV	
<a href="#">Nvidia</a>	L40	VHNWT/20FCY	490-BJHP	490-BJHQ		490-BJHN	
<a href="#">Nvidia</a>	A16	8F04V/D1P1T	490-BHER	490-BHET		490-BHES	
<a href="#">Nvidia</a>	L4 (HHHL)	V9XT2	490-BJRN			490-BJRR	
<a href="#">Nvidia</a>	L4 (FHHL)	NG3PY	490-BJRP			490-BJRQ	
<a href="#">Nvidia</a>	A2 (v2)	H98K0	490-BHVO			490-BHVQ	
<a href="#">Nvidia</a>	A2 (v2)	MG5JX	490-BHVR			490-BHVP	
<a href="#">Nvidia</a>	A10	CGNVH/6KNWW	490-BGKF	490-BGLB	490-BGYD	490-BGKH	

15<sup>th</sup> Gen Server Support

16<sup>th</sup> Gen Server Support

Brand	Model	15 <sup>th</sup> Gen Server Support															16 <sup>th</sup> Gen Server Support														
		R750	R650	R6525	R7525	XR4520C	XR12	XR11	XE9680L	XE9680	XE9640	XE8640	R760xa	R760	R760xs	R760xd2	R660	R7625	R7615	R6625	R6615	R960	T560	C6620	XR8620t	XR7620	XR5610				
AMD	MI300X OAM	X	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
AMD	MI210	X	X	X	✓	X	X	X	X	X	X	✓	X	X	X	X	✓	✓	X	X	X	X	X	X	X	X	X				
Intel	Gaudi3 PCIe	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Intel	Gaudi3 OAM	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Intel	Max1550 OAM	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Intel	Max1100	X	X	X	X	X	X	X	X	X	X	✓	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Intel	Flex 140	✓	✓	X	X	X	✓	✓	X	X	X	✓	✓	X	X	✓	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H200 NVL	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H100 NVL	X	X	X	X	X	X	X	X	X	X	✓	X	X	X	X	✓	✓	X	X	X	X	X	X	X	X	X				
Nvidia	H200 SXM5 (x8) DLC	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H200 SXM5 (x8)	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H100 SXM5 (x8)	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H100 SXM5 (x4)	X	X	X	X	X	X	X	X	✓	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	A100 SXM4 (x8)	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	A100 SXM4 (x4)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	A100 SXM4 (x4)	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				
Nvidia	H100	✓	X	X	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	X	X	X	X	X	X	X	X	X				
Nvidia	A100	✓	X	X	✓	✓	X	X	X	X	X	✓	✓	X	X	X	✓	✓	X	X	✓	X	X	X	✓	X	X				
Nvidia	L40S	X	X	X	X	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	X	X	✓	X	X	X	✓	X	X				
Nvidia	L40	✓	X	X	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	X	X	X	✓	X	X	✓	X	X				
Nvidia	A16	✓	X	X	✓	X	X	X	X	X	X	✓	✓	X	X	X	✓	✓	X	X	✓	X	X	X	✓	X	X				
Nvidia	L4 (HHHL)	✓	✓	X	✓	X	✓	✓	X	X	X	X	X	X	X	✓	X	X	X	X	X	X	X	X	X	✓	X				
Nvidia	L4 (FHHL)	✓	✓	X	✓	✓	✓	✓	X	X	X	✓	✓	X	✓	✓	✓	✓	✓	✓	X	✓	X	✓	✓	✓	✓				
Nvidia	A2 (v2)	✓	✓	✓	✓	X	X	✓	X	X	X	X	X	✓	X	✓	X	X	X	X	X	X	X	✓	X	✓	X				
Nvidia	A2 (v2)	✓	✓	✓	✓	✓	✓	✓	X	X	X	✓	✓	X	✓	✓	✓	✓	✓	✓	X	✓	X	✓	✓	✓	✓				
Nvidia	A10	✓	X	X	✓	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X				

17 <sup>th</sup> Gen Server Support											
Brand	Model	XE7740	XE7745	R6715	R6725	R7715	R7725	R570	R470	R670	R770
AMD	MI300X OAM	X	X	X	X	X	X	X	X	X	X
AMD	MI210	X	X	X	X	X	X	X	X	X	X
Intel	Gaudi3 PCIe	✓	X	X	X	X	X	X	X	X	X
Intel	Gaudi3 OAM	X	X	X	X	X	X	X	X	X	X
Intel	Max1550 OAM	X	X	X	X	X	X	X	X	X	X
Intel	Max1100	X	X	X	X	X	X	X	X	X	X
Intel	Flex 140	X	X	X	X	X	X	X	X	X	X
Nvidia	H200 NVL	✓	✓	X	X	X	X	X	X	X	X
Nvidia	H100 NVL	✓	✓	X	X	✓	✓	X	X	X	✓
Nvidia	H200 SXM5 (x8) DLC	X	X	X	X	X	X	X	X	X	X
Nvidia	H200 SXM5 (x8)	X	X	X	X	X	X	X	X	X	X
Nvidia	H100 SXM5 (x8)	X	X	X	X	X	X	X	X	X	X
Nvidia	H100 SXM5 (x4)	X	X	X	X	X	X	X	X	X	X
Nvidia	A100 SXM4 (x8)	X	X	X	X	X	X	X	X	X	X
Nvidia	A100 SXM4 (x4)	X	X	X	X	X	X	X	X	X	X
Nvidia	A100 SXM4 (x4)	X	X	X	X	X	X	X	X	X	X
Nvidia	H100	X	X	X	X	X	X	X	X	X	X
Nvidia	A100	X	X	X	X	X	X	X	X	X	X
Nvidia	L40S	X	✓	X	X	✓	✓	✓	X	X	X
Nvidia	L40	X	X	X	X	X	X	X	X	X	X
Nvidia	A16	✓	X	X	X	✓	✓	X	X	X	✓
Nvidia	L4 (HHHL)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Nvidia	L4 (FHHL)	X	X	X	X	X	X	X	X	X	X
Nvidia	A2 (v2)	X	X	X	X	X	X	X	X	X	X
Nvidia	A2 (v2)	X	X	X	X	X	X	X	X	X	X
Nvidia	A10	X	X	X	X	X	X	X	X	X	X

PLAT-FORM	NVIDIA										
	H100 80GB PCIe	B200 DLC	H100 NVL	H200 NVL	H200 SXM5 (x8)	A100 80GB PCIe	H100 SXM5 (X8)	H100 SXM5 (X4)	L40S	L4	A16
XE7745				Shipping (8)					Shipping (8)		
XE7740	Shipping (8)										
R770			Shipping (2)						Shipping (2)	Shipping (4)	
R670	Shipping (4)										
R570										Shipping (4)	
R470	Shipping (2)										
R7725									Shipping (2)	Shipping (6)	Shipping (2)
R7715	Shipping (3) Shipping (6) Shipping (3)										
R6725										Shipping (3)	
R6715	Shipping (3)										
XE9685L		Shipping in EMEA & APJ									
XE9680L	Shipping										
XE9680					Shipping		Shipping				
XE9640	Shipping										
XE8640								Shipping			
R760xa	Shipping (4 <sup>3</sup> )		Shipping (4 <sup>2</sup> )			Shipping (4 <sup>3</sup> )			Shipping (4 <sup>3</sup> )	Shipping (8 <sup>3</sup> )	Shipping (4 <sup>3</sup> )
R760	Shipping (2)		Shipping (2)			Shipping (2)			Shipping (2)	Shipping (4)	Shipping (2)
R760xs											
R760xd2										Shipping (2)	
R660	Shipping (3)										
R7625	Shipping (2)		Shipping (2)			Shipping (2)			Shipping (2)	Shipping (4)	Shipping (2)
R7615	Shipping (3)		Shipping (3)			Shipping (3)			Shipping (3)	Shipping (4)	Shipping (3)
R6625										Shipping (3)	
R6615	Shipping (2)										
R960						Shipping (4)			Shipping (4)		Shipping (4)
T560	Shipping (5)										
R360											

(qty) - max number of GPUs allowed, maximum number of GPUs allowed might differ in different configurations on the same platform

PLATFORM	NVIDIA										
	H100 80GB PCIe	B200 DLC	H100 NVL	H200 NVL	H200 SXM5 (x8)	A100 80GB PCIe	H100 SXM5 (x8)	H100 SXM5 (x4)	L40S	L4	A16
C6620											
XR7620						Shipping (2)			Shipping (2)	Shipping (5)	
XR5610										Shipping (2)	
XR8620t										Shipping (3)	
R750	Shipping (2)					Shipping (2)				Shipping (6)	Shipping (2)
R650										Shipping (3)	
R7525	Shipping (3)					Shipping (3)				Shipping (6)	Shipping (3)
R6525											
XR4520c										Shipping (2)	
XR12						Shipping (2)				Shipping (3)	
XR11										Shipping (3)	

(qty) - max number of GPUs allowed, maximum number of GPUs allowed might differ in different configurations on the same platform

PLATFORM	AMD				INTEL		
	MI210	MI300X OAM (X8)	MI355X	MI350X	Flex 140	Gaudi3 PCIe	Gaudi3 OAM
XE7745							
XE7740							
R770							
R660							
R570							
R470							
R7725							
R7715							
R6725							
R6715							
XE9680		Shipping					Shipping
XE9640							
XE8640							
R760xa	Shipping (4 <sup>3</sup> )				Shipping (10 <sup>3</sup> )*		
R760					Shipping (6)*		
R760xs							
R760xd2							
R660					Shipping (3)*		
R7625	Shipping (2)						
R7615	Shipping (3)						
R6625							
R6615							
R960							
T560							
R360							
T360							
C6620							
XR7620							
XR5610							
XR8620t							
R750					Shipping (6)		
R650					Shipping (2)		
R7525	Shipping (3)						
R6525							
XR4520c							
XR12					Shipping (2)		
XR11					Shipping (3)		

(qty) - max number of GPUs allowed, maximum number of GPUs allowed might differ in different configurations on the same platform

1. Suggested ideal workloads, but can be used for other workloads
2. Different SKUs are mentioned because different platforms might support different SKUs. This sheet doesn't specifically call out platform-SKU associations
3. up to 600GB/sec for H100 when NVLink bridge is used, up to 200GB/sec when NVLink bridge is used, up to 112.5GB/sec when NVLink bridge is used, up to 300GB/sec for MI210 when Infinity Fabric Link bridge is used
4. w/NVLink bridge is supported on R760xa, w/NVLink bridge is supported on R760xa, MI210 w/Infinity Fabric Link bridge is supported on R760xa
5. DW - Double Wide, SW - Single Wide, FH- Full Height, FL - Full Length, HH - Half Height, HL - Half Length

\* Currently available as Customer Install only



Learn more about  
Dell solutions



Contact a Dell  
Technologies Expert



View more resources



Join the conversation  
with #PowerEdge