Advancements in AI Factory with NVIDIA Blackwell Ultra and Future Outlook with NVIDIA Vera Rubin

Dr. Arts Yang, Principal Architect, Systems

NVIDIA Blackwell

The Engine of the New Industrial Revolution Now in Full Production

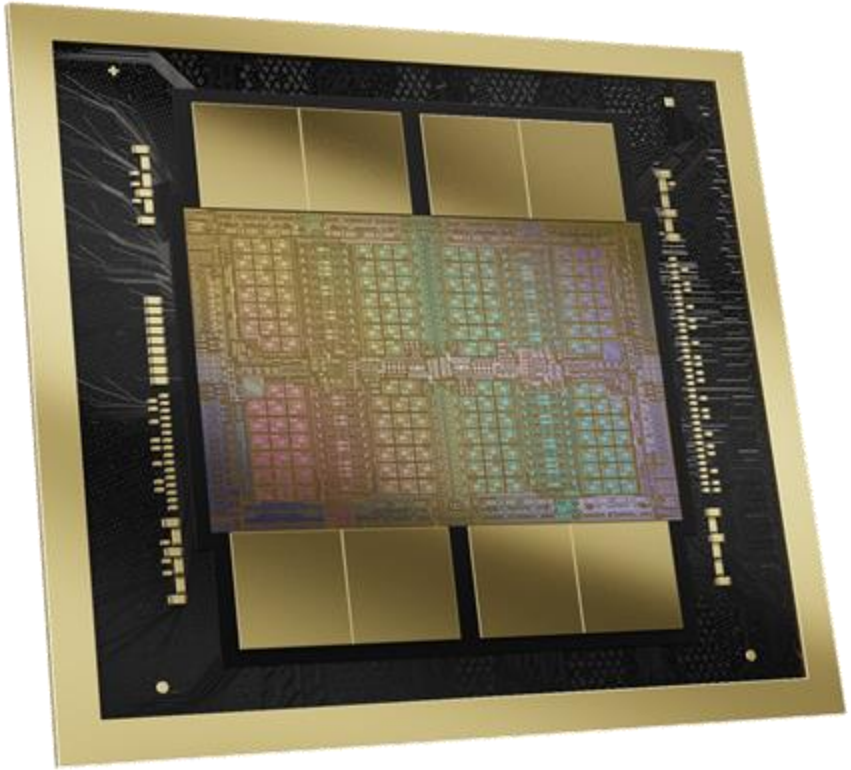

Built to Democratize Trillion-Parameter AI

20 PetaFLOPS of AI performance on a single GPU

4X Training | 30X Inference | 25X Energy Efficiency & TCO

Expanding AI Datacenter Scale to beyond100K GPUs

AI SUPERCHIP
208B Transistors

2nd GEN TRANSFORMER ENGINE
FP4/FP6 Tensor Core

5th GENERATION NVLINK
Scales to 576 GPUs

RAS ENGINE
100% In-System
Self-Test

SECURE AI
Full Performance
Encryption & TEE

DECOMPRESSION ENGINE
800 GB/s

DGX GB300 (Blackwell Ultra)

Accelerate Real-Time State of the Art Inference Models

- Based on the NVIDIA GB300 NVL72 rack architecture
- Provides the foundation for NVIDIA DGX SuperPOD with DGX GB300
- Powered by Grace Blackwell Ultra Superchips connected with fifth-generation NVIDIA NVLink
- 36 Grace CPUs and 72 Blackwell Ultra GPUs
- 1.4 exaFLOPS of AI performance and 38TB of fast memory
- Massive shared memory space to accelerate the most data-intensive workloads

DGX B300

Accelerated Infrastructure for the Era of AI Reasoning

10U Chassis | ~14.5 kW system

Designed for the modern data center

- Newest air-cooled DGX system with NVIDIA Blackwell Ultra GPUs
- All new system design seamlessly integrates into NVIDIA MGX or traditional enterprise racks
- **2.3TB of GPU memory**, enabling training and inference of complex models
- Equipped with NVIDIA ConnectX-8 high speed networking at **800Gb/s**
- Delivers **72 PFLOPS AI training** and **144 PFLOPS AI inference** performance
- Purpose-built platform for the era of AI reasoning, setting a new bar for LLM inference

SuperPOD Energy Components – Learnings from GB200 NVL72

Power Ratio (576 GPU Example)

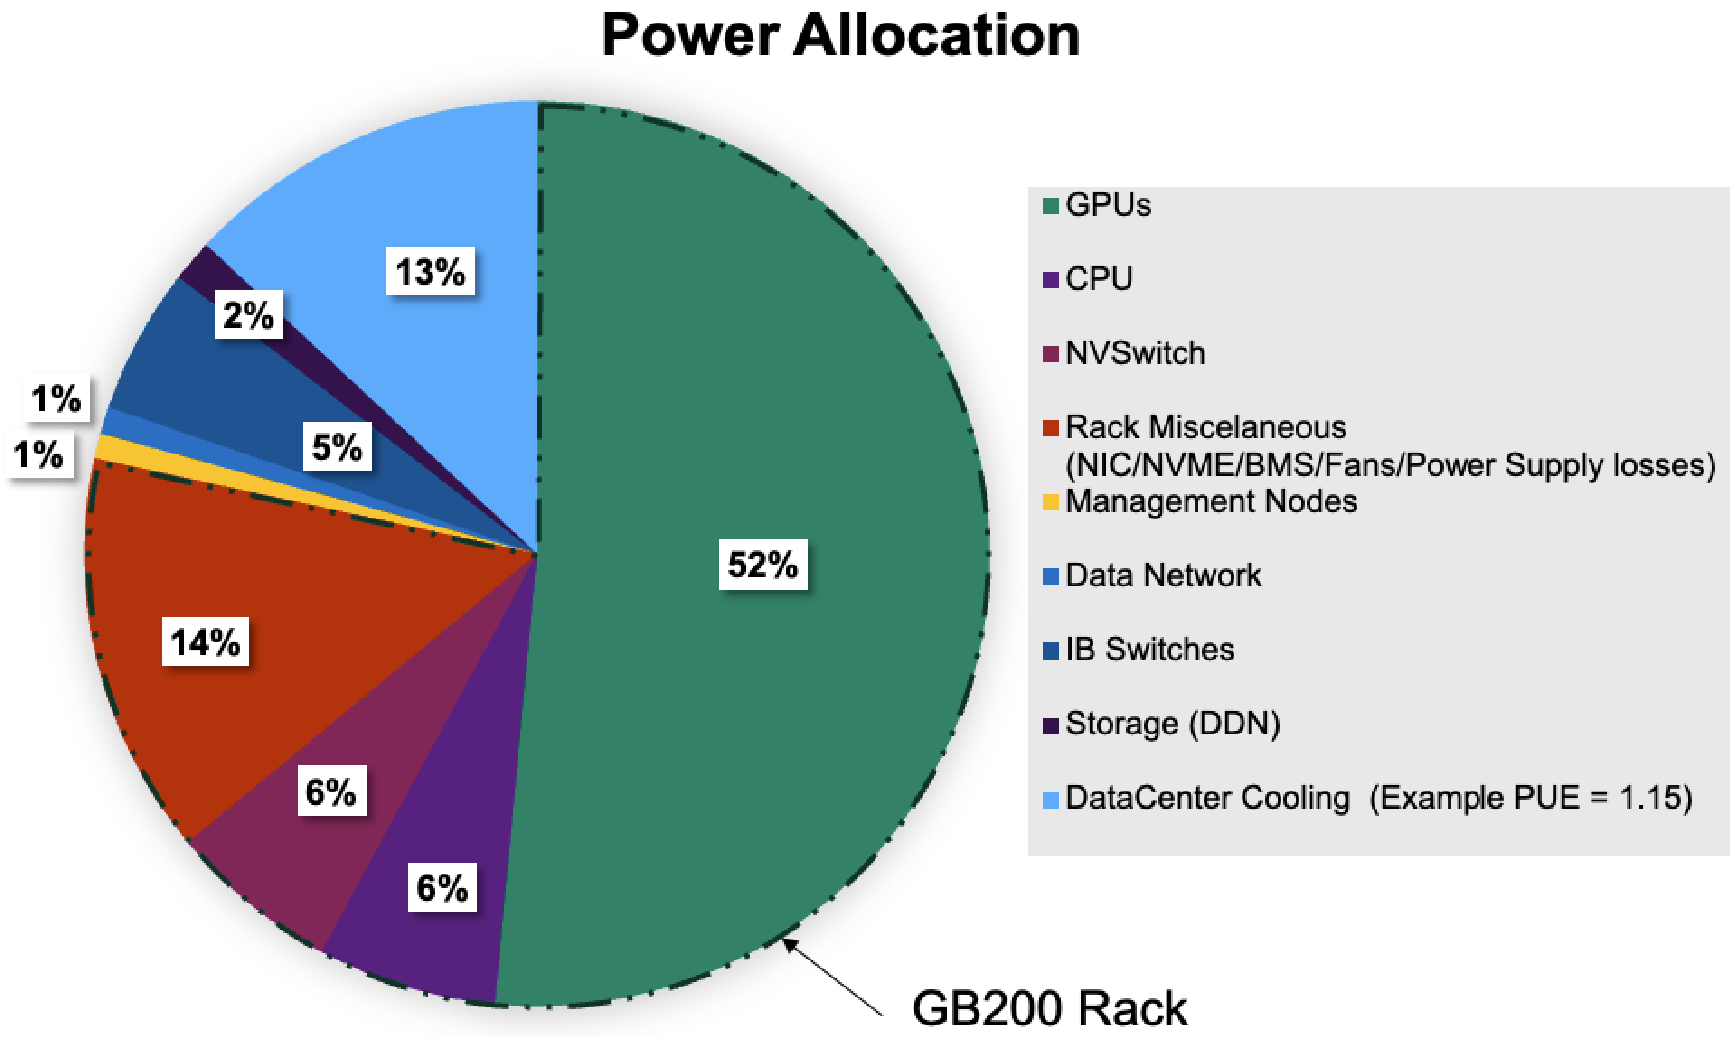

DGX GB300 NVL72 and B300 NVL8

GB300 NVL72 System Design

DC Powered, Rack-scale, DLC+Air Hybrid cooled (90%/10%)

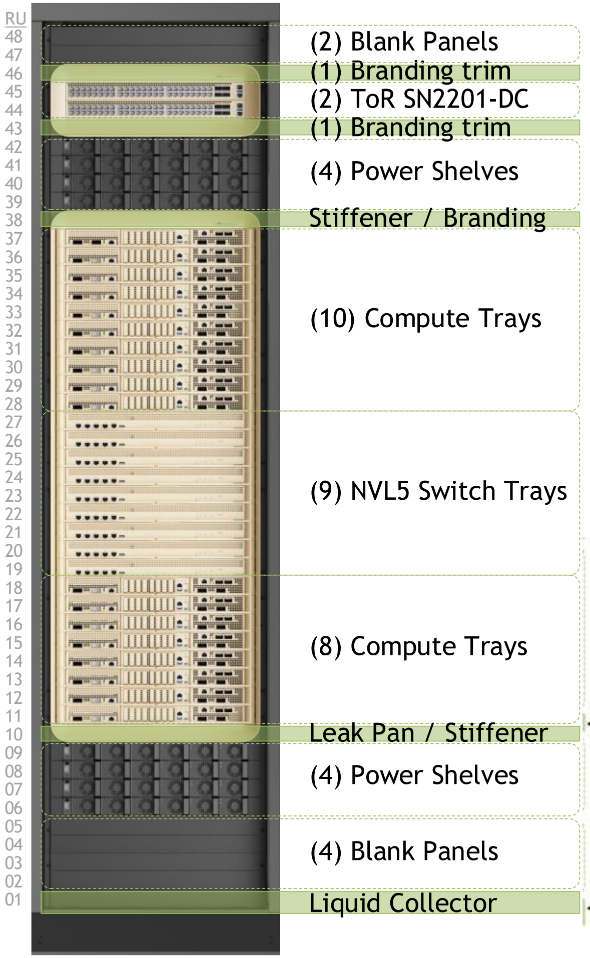

<u>Rack Configuration</u>

- DGX rack with 1400A busbar
- (18) DGX 1RU Compute Tray (1C:2G):
  - (2) Bianca
    - (1) Grace TH500 + (2) GPUs B110
- (9) DGX 1RU Switch Tray.
  - NSO NVL5 switch tray: 100% LC
  - (4) Cable cartridges for NVL
- (8) DGX Power Shelf:
  - 33kW 1RU EIA power shelf
  - Each power shelf will have 6x 5.5kw PSUs
- (2) DGX SN2201-DC management switches

<u>Leak Pan</u>

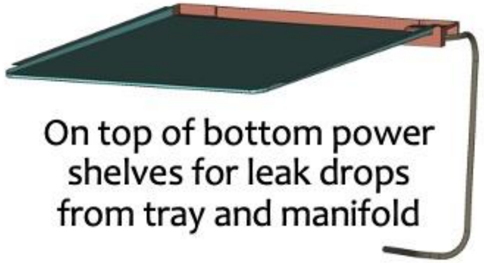

On top of bottom power shelves for leak drops from tray and manifold

<u>Liquid Collector</u>

- Placed at rack bottom to collect liquids from the drain tube of top leak pan
- Spot leak sensor deployed by DC facility and connected to BMS

DGX GB300 SuperPOD 1SU HAC

Updated Design with higher-rack density on non-compute Racks

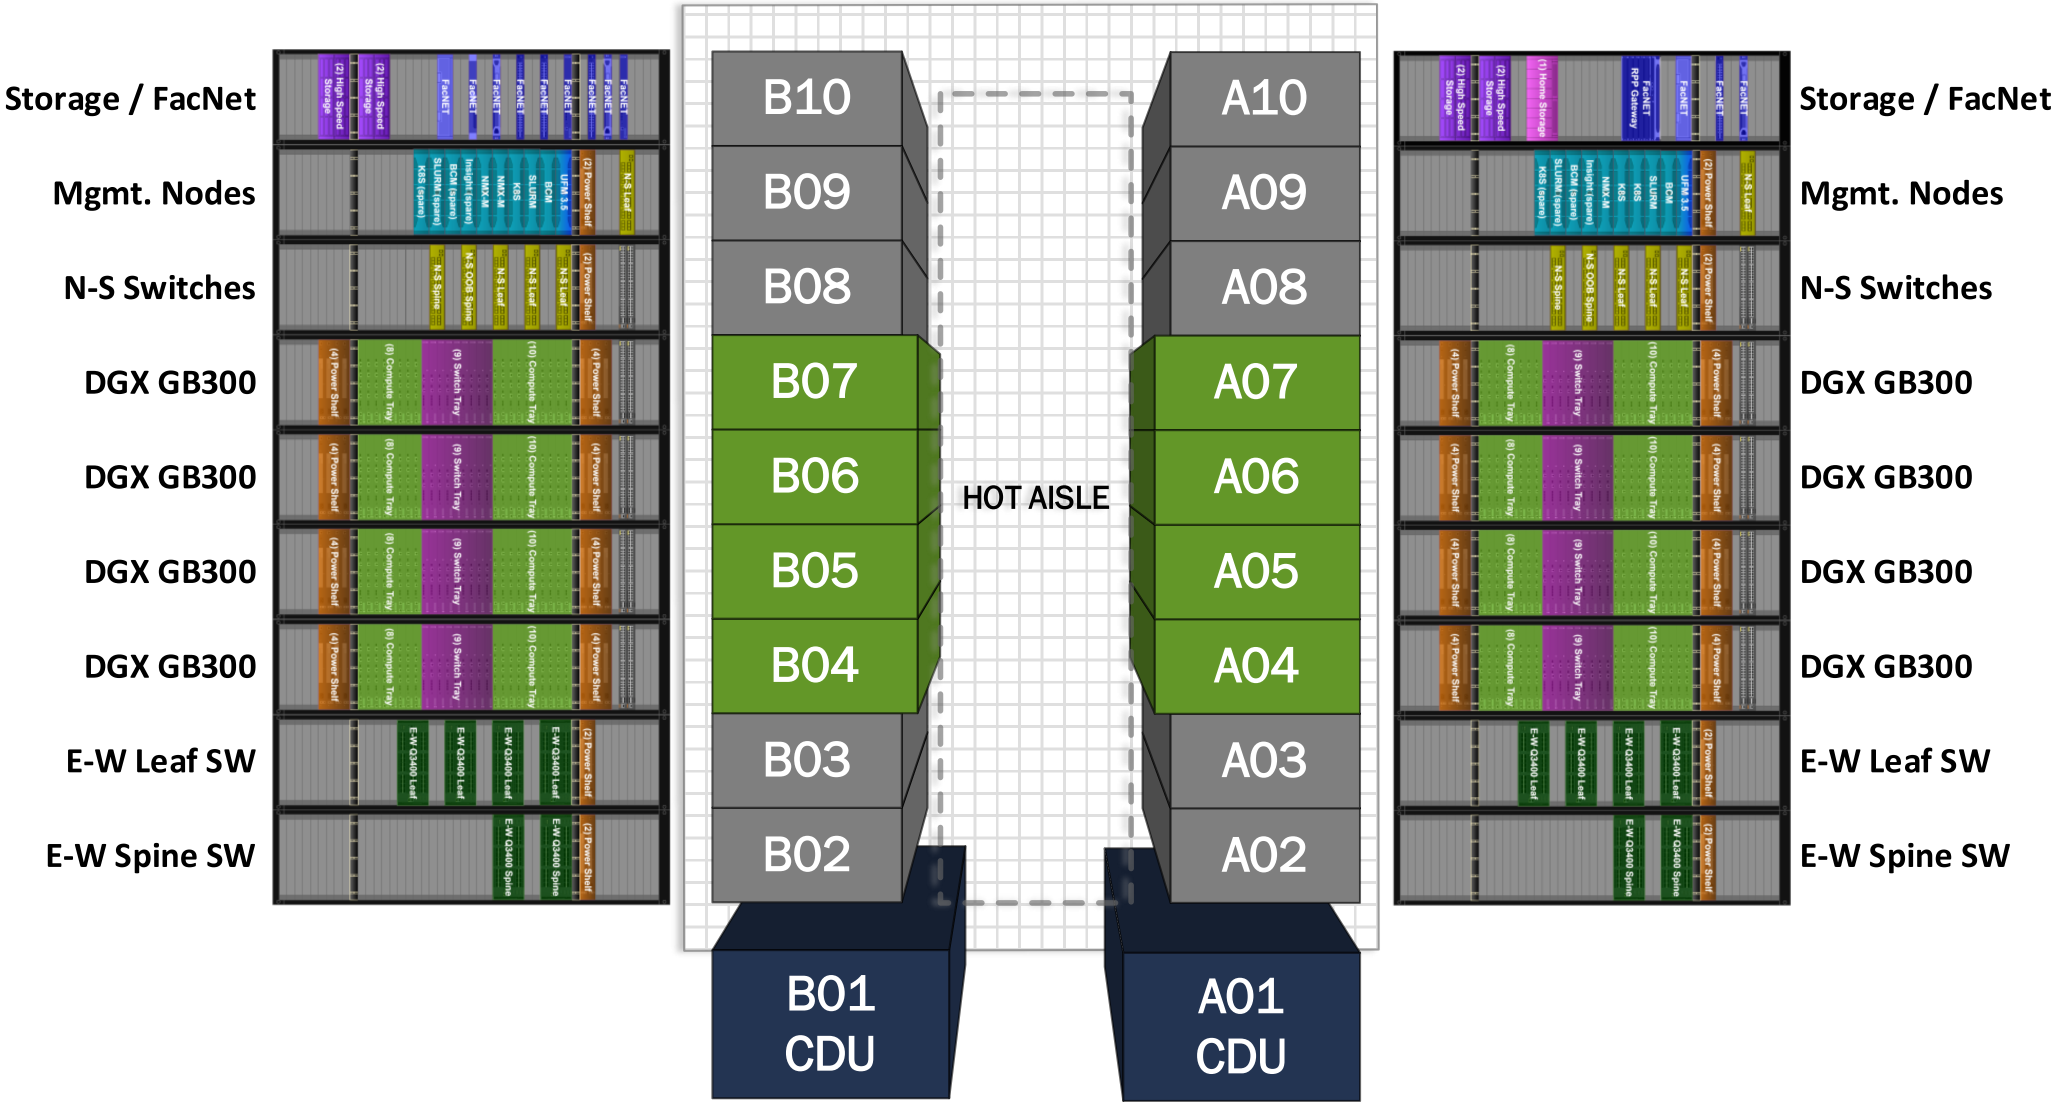

NVIDIA

Programmable Power Smoothing and EDP Attributes

Repetitive Load Profile on UPS Outputs

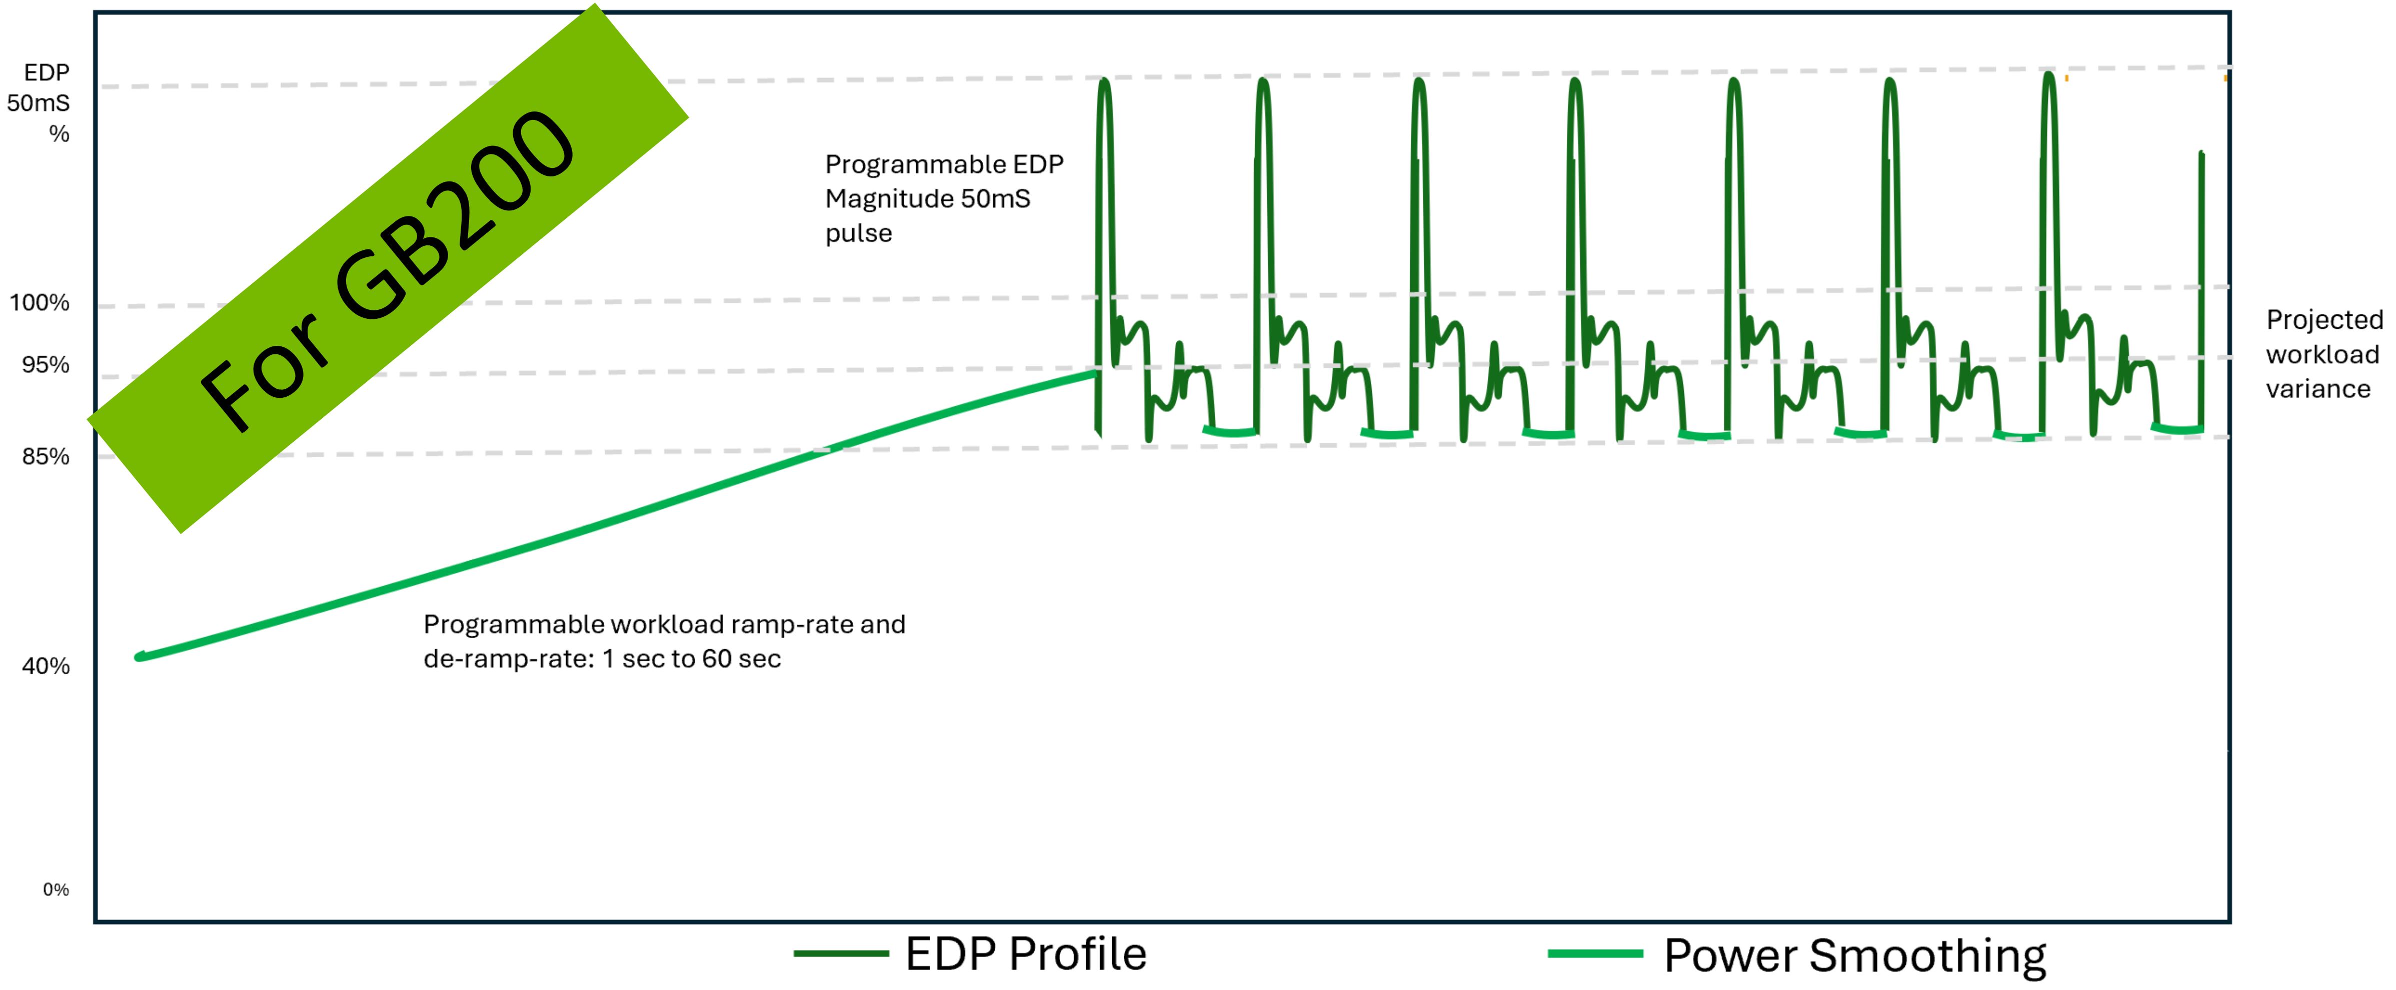

Power & Cooling Requirements for GB300

HAC / SU	GPU	Racks	[1]EDPc (MW)	[1]EDPp1 (MW)	[2]EDPc LPM @40C TCS	[2]EDPc CFM @35C Inlet
1	576	8+10	1.25 (**1.1**+0.15)	1.35 (**1.2**+0.15)	1K per HAC	32K per HAC
2	1152	16+18	2.47 (**2.2**+0.27)	2.67 (**2.4**+0.27)		
4	2304	32+34	4.92 (**4.4**+0.52)	5.32 (**4.8**+0.52)		
8	4608	64+66	9.82 (**8.8**+1.02)	10.62 (**9.6**+1.02)		
16	9216	128+130	19.60 (**17.6**+2.0)	21.20 (**19.2**+2.0)		

[1] *Increase in power due to GPU upgrade from GB200 to GB300*

[2] *Increase in LPM and decrease in CFM due to cooling changes of 87% LC/13% air vs. 80% LC/20% air*

DGX GB300 SuperPOD 1SU HAC – Cable Tray and Facility Water

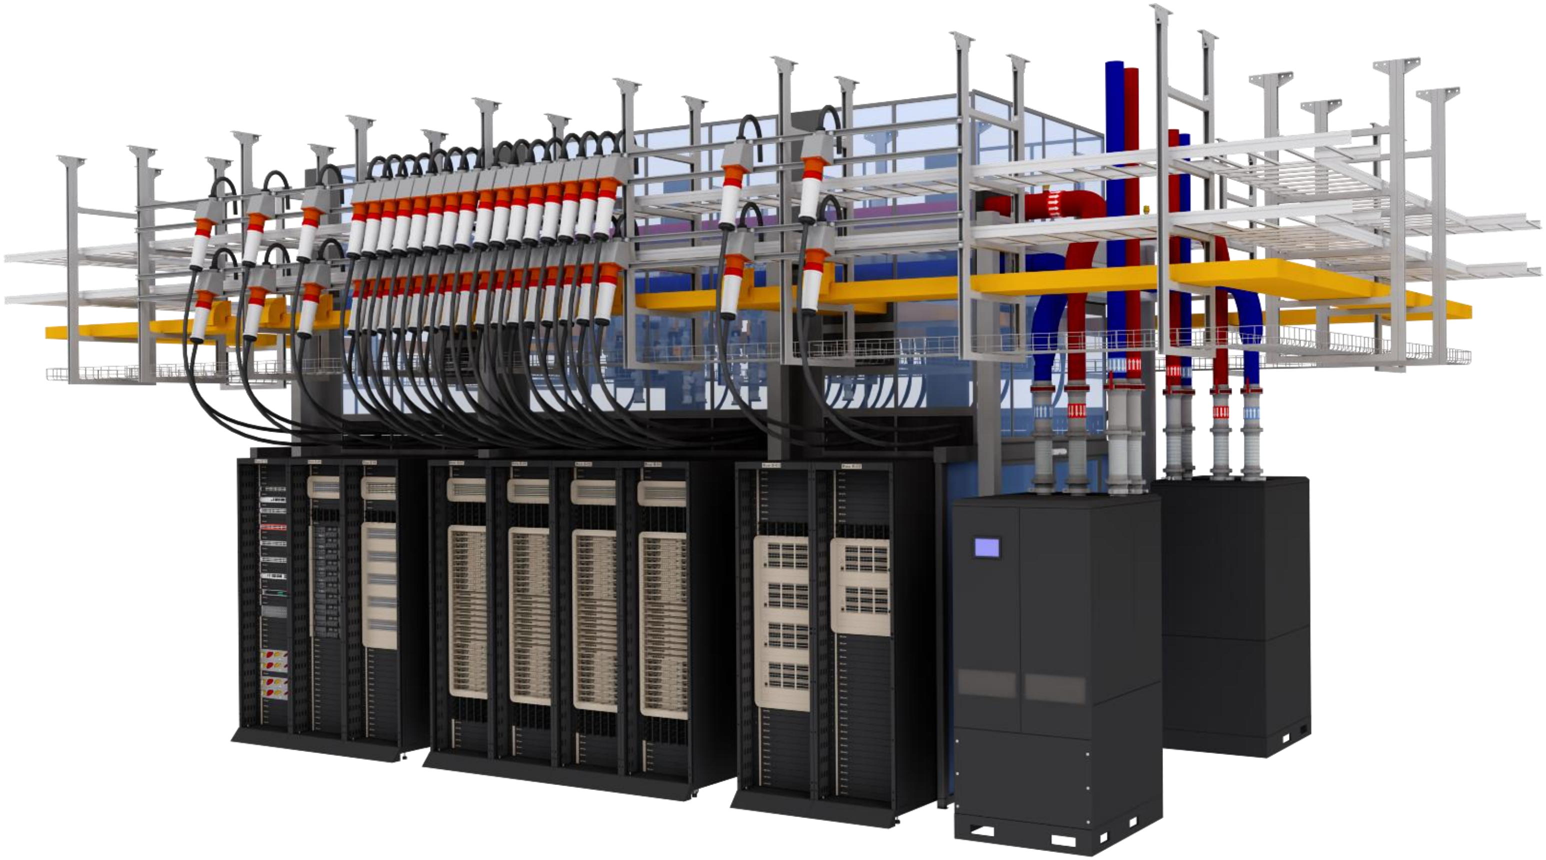

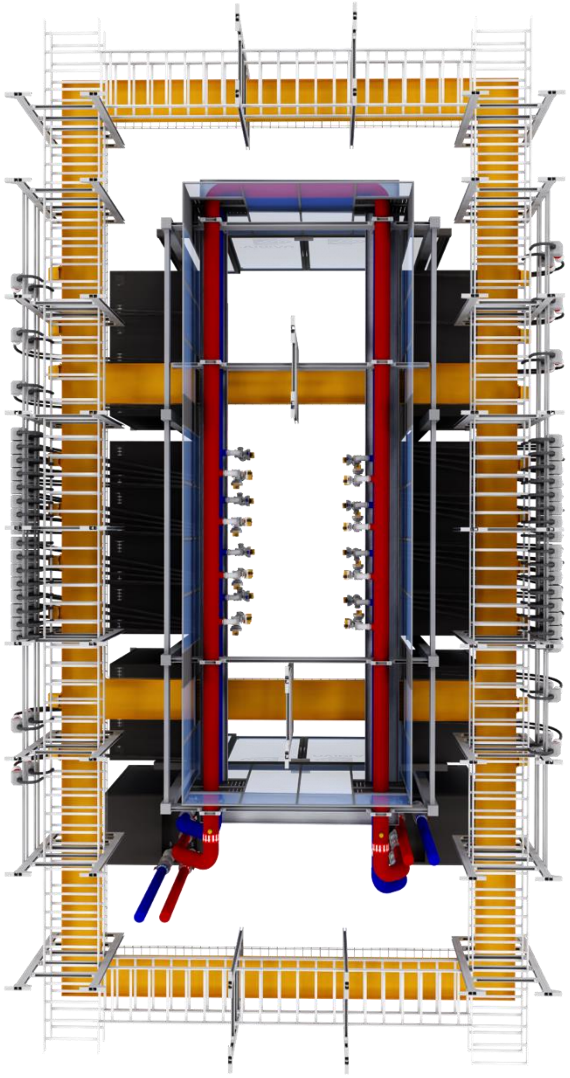

DGX B300

10U Busbar

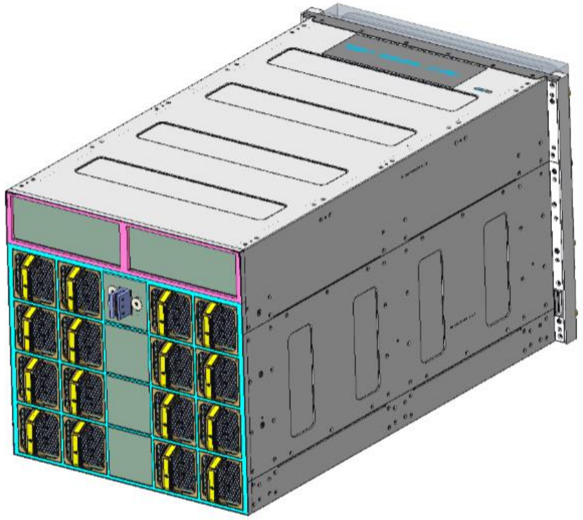

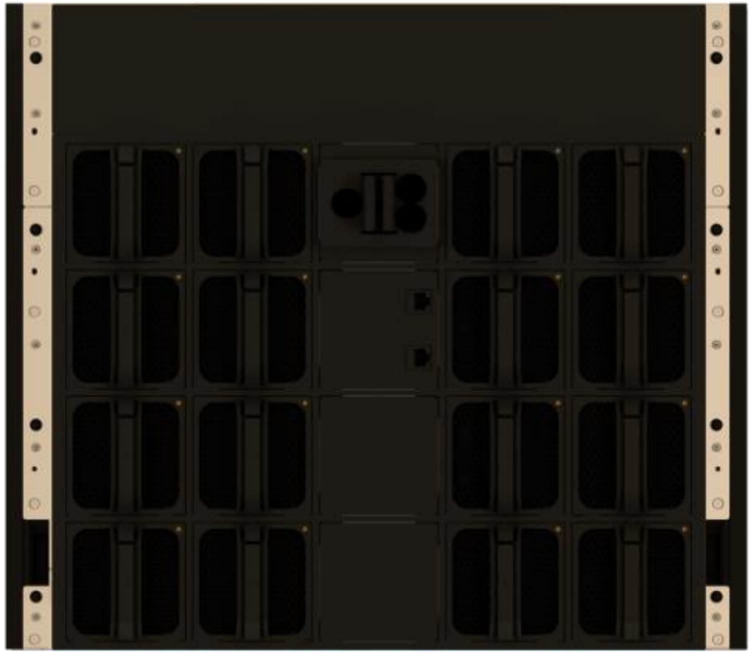

Different Options for DGX SuperPOD B300

4 Total Options (DC/AC, XDR/Ethernet)

	Spectrum-X (Minimum Order 64n)		XDR InfiniBand	
	DC	AC	DC	AC
Compute Fabric Switch	SN5600D	SN5610	Q3400-RD	Q3400-RA
Compute Fabric Management	NetQ 3x shared with Admin K8s		2x UFM3.5 2U DC	2x UFM 3.5 1U AC
Compute Fabric Transceiver & Cables	Single or Multi-Mode, 2x400G		Single Mode, 1x 800G	
SU Size	64 Nodes, 512 GPUs		72 Nodes, 576 GPUs	
Storage Fabric	Spectrum Ethernet (SN5600D or SN5610) NDR InfiniBand (Q9700DC or Q9700)			
In-band Management Fabric	Spectrum Ethernet (SN5600D or SN5610)			
Out-of-band Management	Spectrum Ethernet (SN2201 or SN2201M)			
Power Supply	54V DC 300A	12x C19	54V DC 300A	12x C19
TDP (kW)	14.5kW			
EDPp1	15.34 kW	16.32 kVA	15.34 kW	16.32 kVA
Cooling	Air			
Software	NVIDIA Mission Control			

Figure 2: Power Shelf #

DC or AC version DGX B300 SuperPOD can be both IB Compute Fabric or Spectrum-X Compute Fabric!!

Different Options for DGX BasePOD B300

Only XDR IB with DC or AC

	Spectrum-X		XDR InfiniBand	
			DC	AC
Compute Fabric Switch	SN56	5610	2x Q3400-RD	2x Q3400-RA
Compute Fabric Management	3x shar	min K8s	1x UFM3.5 2U DC	1x UFM 3.5 1U AC
Compute Fabric Transceiver & Cables	Singl	ulti-Mod 400G	Single Mode, 1x 800G	
Supported Size	64 Nodes, 512 GPUs		1-18 Nodes 2 Compute Switches	
Storage Fabric	Spectrum Ethernet (SN5600D or SN5610) NDR InfiniBand (Q9700DC or Q9700)			
In-band Management Fabric	Spectrum Ethernet (SN5600D or SN5610)			
Out-of-band Management	Spectrum Ethernet (SN2201 or SN2201M)			
Power Supply	54V DC 300A	12x C19	54V DC 300A	12x C19
TDP (kW)	14.5kW			
EDPp1	15.34 kW	16.32 kVA	15.34 kW	16.32 kVA
Cooling	Air			
Software	NVIDIA Mission Control			

Figure 2: Power Shelf #

BasePOD (partner lead install **only Supports XDR Infiniband Compute Fabric)**

DGX B300 NVL8 DC Air-cooled – High Density Racks

AC/DC-powered, air-cooled, Liquid Rear Door Heat Exchanger

- 32 GPUs
- 8x Intel Granite Rappids CPUs
- 32x ConnectX-8 XDR IB or Spectrum-X NICs
- 8x BlueField-3 DPUs
- 4x DGX B300 NVL8 per Rack

- MGX-v1.1 compatible racks
- Rack TDP <= 60kW
- 4x Power Shelves with large capacitors for transient power smoothing
- Rear-door heat exchange support with same liquid cooling connectivity as GB300 NVL72 racks

Liquid Control

Leak Monitoring and Control

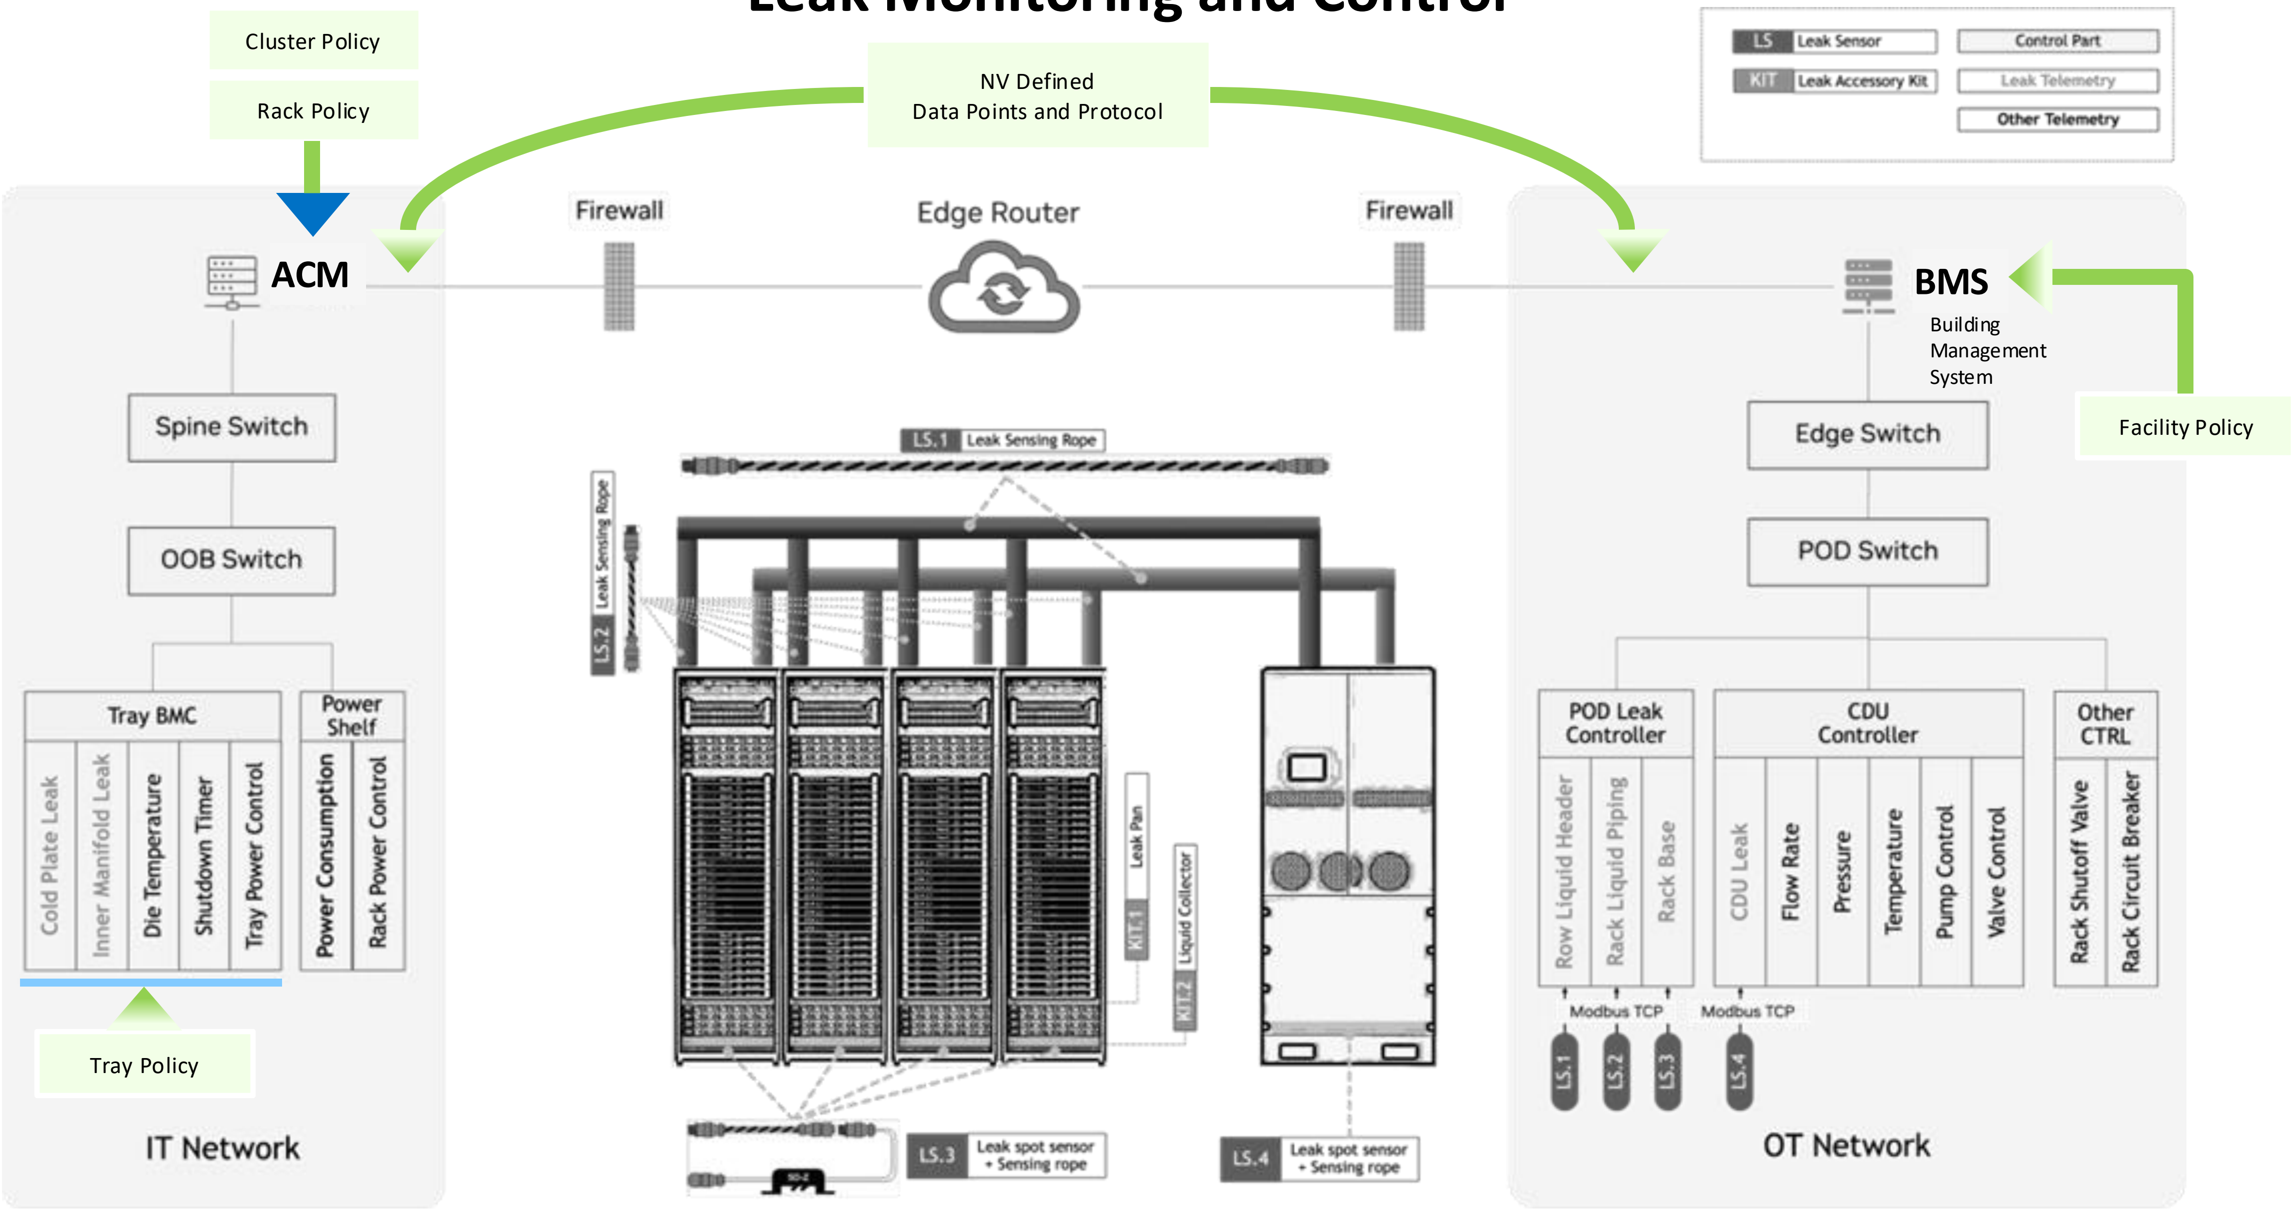

NVIDIA

BCM-BMS Integration

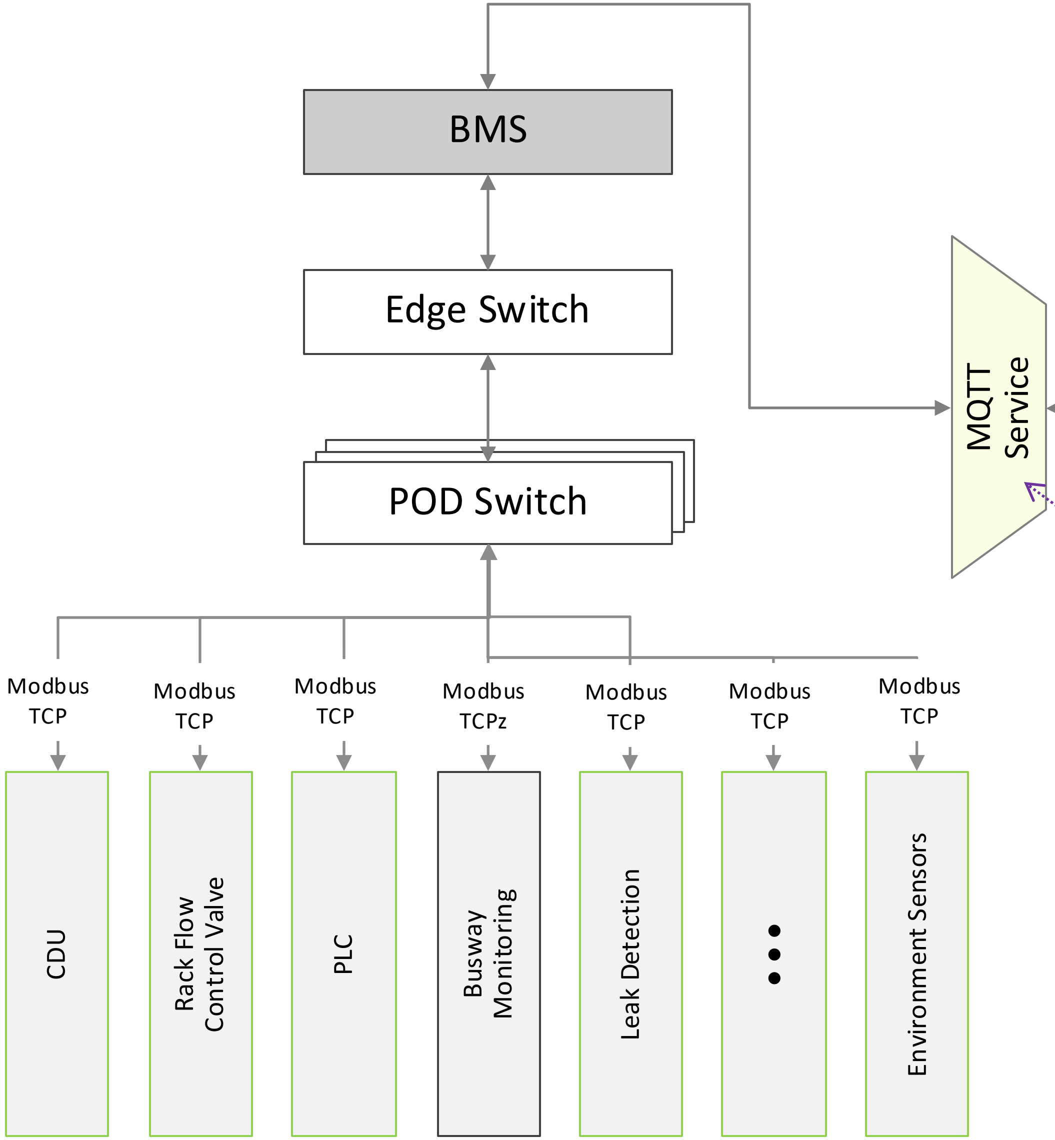

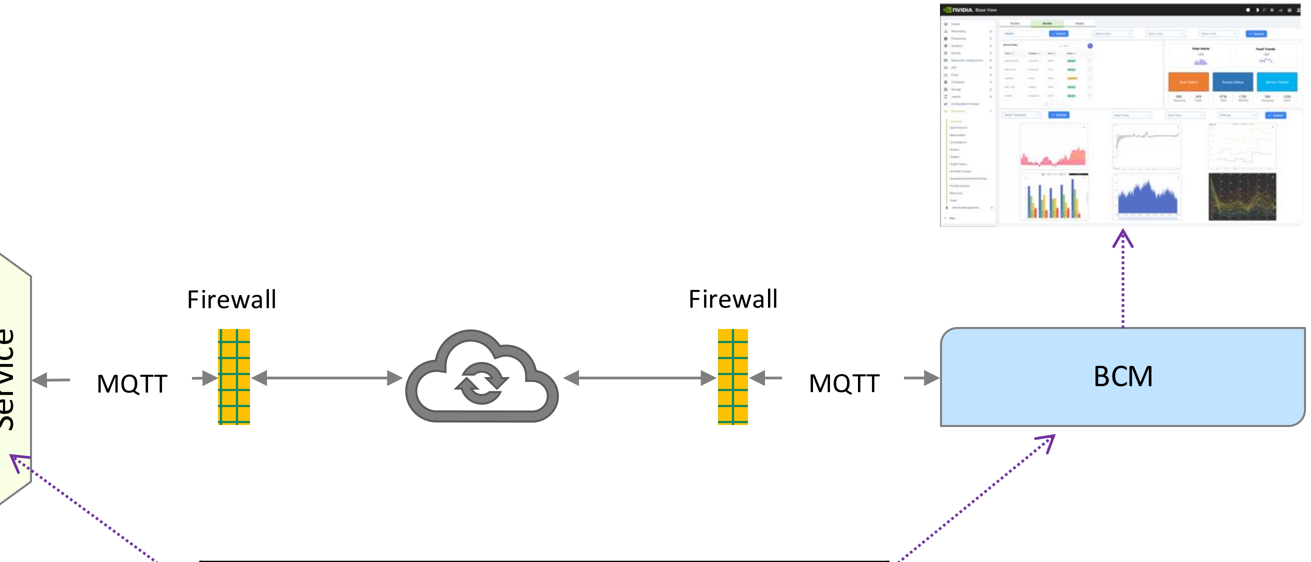

MQTT Data Points
BCM/TPE01/(RackID)/LEAK/DetectStatus/Value
BCM/TPE01/(RackID)/LEAK/TrayDetectStatus/Value
BCM/TPE01/(RackID)/LEAK/LiquidIsolationStatus/Value
BCM/TPE01/(RackID)/LEAK/ElectricalIsolationStatus/Value
BCM/TPE01/(RackID)/LEAK/LiquidIsolationReqeust/Value
BCM/TPE01/(RackID)/LEAK/ElectricalIsolationRequest/Value
BCM/TPE01/(CDU#)/LIQUID/SupplyTemperature/Value
BCM/TPE01/(CDU#)/LIQUID/ReturnTemperature/Value
BCM/TPE01/(CDU#)/LIQUID/DP/Value
BCM/TPE01/(CDU#)/LIQUID/Flow/Value
BCM/TPE01/(CDU#)/LIQUID/SystemPressure/Value
BCM/TPE01/(CDU#)/Status/Value
BCM/TPE01/(CDU#)/Available/Value

- NVIDIA defined data points and structure for BCM-BMS integration
- Leverage MQTT messaging framework for lightweight and reliable implementation

Outlook: Vera Rubin

54V DC
5000A Liquid Cooled Busbar

Oberon Rack
Liquid Cooled

Rubin NVL144

Second Half 2026

3.6 EF FP4 Inference
1.2 EF FP8 Training
3.3X GB300 NVL72

13 TB/s HBM4
75 TB Fast Memory
1.6X

260 TB/s NVLink6
2X

28.8 TB/s CX9
2X

Vera

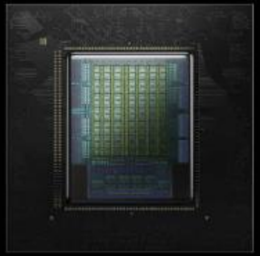

88 Custom Arm Cores
176 Threads
1.8 TB/s NVLink-C2C

Rubin

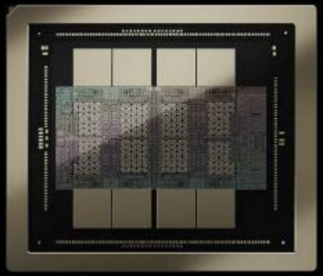

2 Reticle-Sized GPUs
50PF FP4 | 288GB HBM4

800V HVDC
Higher PUE

Kyber Rack
Liquid Cooled

Rubin Ultra NVL576

Second Half 2027

Vera

88 Custom Arm Cores

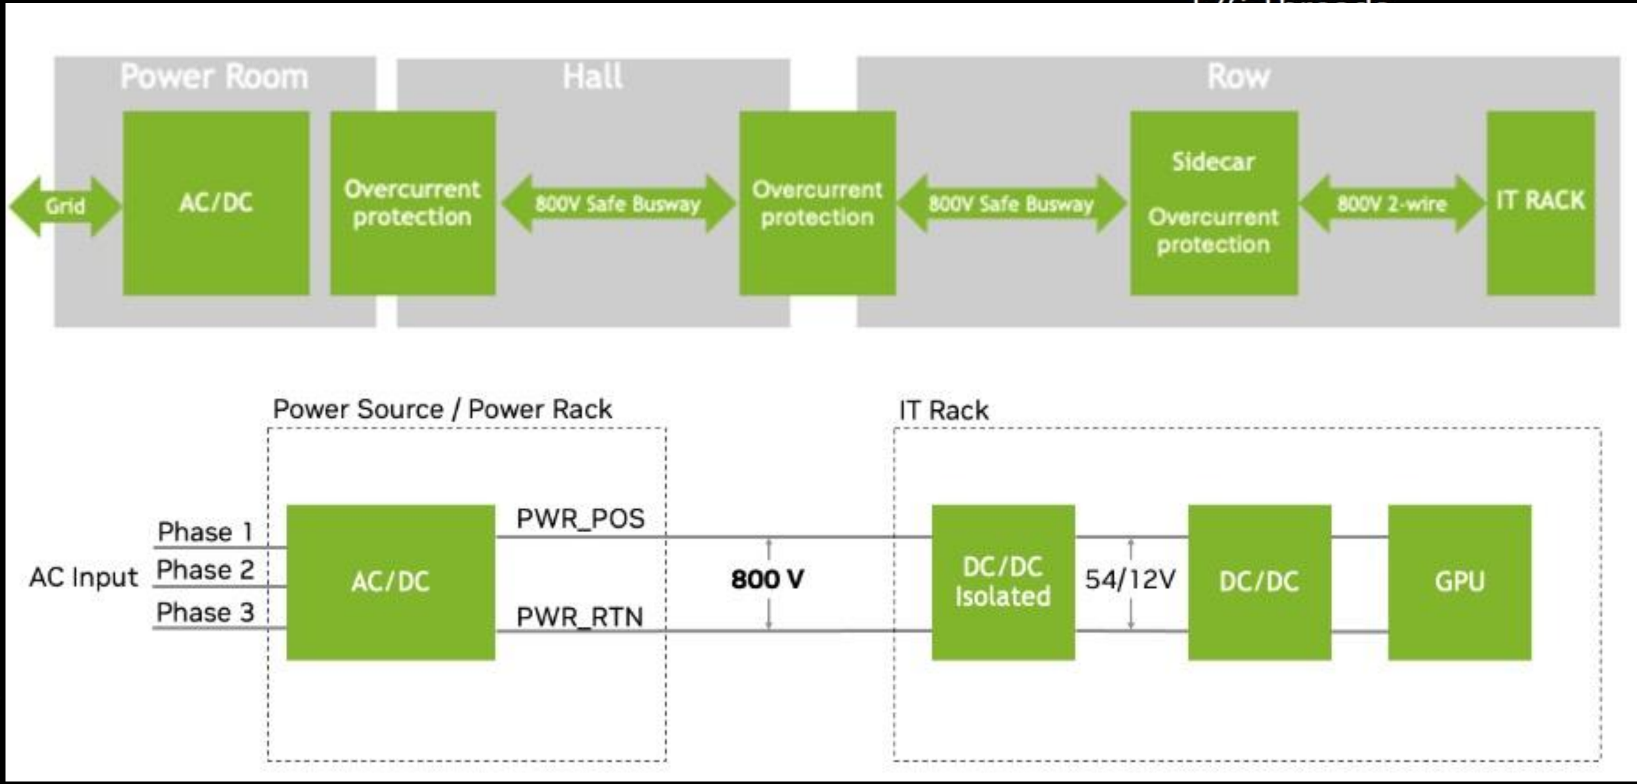

NVIDIA