AI OPEN
DATA CENTER



Next-Generation ML Infrastructure: The Open DC Architecture

Ayan Majumdar *HW Lead - Google*
Mike Sabotta *Solutions Architect for Cloud and HPC system design - NVIDIA*
Moe Ghamary *System Architect, Data center engineering-Google*
Paul Artman *Director System Design - AMD*
Mike Edie *Lead System Engineer - Meta*

OCP PLATINUM Member

Panel Discussion



Ayan Majumdar

HW Lead - Google



Maureen Tate

Senior Solutions Architect for Datacenters - NVIDIA



Moe Ghamary

(Moderator)

System Architect, Data center engineering-Google



Paul Artman

Director System Design - AMD



Mike Edie

Lead System Engineer - Meta

Next- Generation ML infrastructure

- Rapid ML advancement requires evolving data center infrastructure. Google, Meta, NVIDIA, and AMD established the Open DC working group to develop a minimum set of standards for AI factories for 2028 and beyond

bolt **Power**

Managing power densities from **125 kW/ft to 350 kW/ft** with unified interfaces for redundancy.

dns **Physical**

Architectural guidelines for liquid-cooled environments, supporting **>70-inch** rack depths and **>80-inch** cold aisles.

ac_unit **Thermal**

Scaling advanced liquid cooling systems to support high-TDP accelerators.

- This initiative empowers the ecosystem to efficiently deploy mega-scale AI factories through collaborative, open-source standards.

Unified Power & Cluster Density

Establishes a unified timeline target spanning hardware deployments from 2027 through 2028.

Metric Baseline	2027 Baseline Target	2028+ Target
Blended DC Capacity Baseline	12 kW/ft 8 kW/ft Air-cooled 8kW/ft liquid-cooled	25 kW/ft 15 kW/ft Air-cooled 25 kW/ft liquid-cooled
Peak IT Density	125 kW/ft 115 kW/ft liquid-cooled	**350 kW/ft** 335kW/ft Peak liquid-cooled
Cluster Size Baseline	12 MW Minimum threshold	**24 MW** Preferred design baseline

Facility & Infrastructure Specifications

- Cluster capacity baselines (12-24 MW)
  - Minimum of 4 power blocks per data hall.

- Supports high weight distributions for heavy high-density racks:
  - Maximum Rack Weight Limit:
    - 8,000 lbs (3,629 kg) over 28” W x 48” D (711 x 1219 mm).
  - Peak IT Weight (ML Rack + Sidecar):
    - 15,000 lbs (6,804 kg) over 56” W x 48” D (1422 x 1219 mm).

Row Layout & Clearance Dimensions

Aisle Pitch Recommendation

- Maintains recommended 24’ aisle pitch with front-justified racks in the cold aisle.

Transport Aisle Clearance

- 12’ transport aisle clearance to a design recommendation.

Ceiling Height

- 18 ft (5.5m) minimum for 2-busway configurations. 4-busway configurations/additional trays may require additional height.

Minimum Physical Boundaries

Hot Aisle Width	**54” (1372mm)**
Cold Aisle Width (Face-to-Face)	**82” (2083mm)**
Rack Height	**100” (2540mm)**
Rack Depth	**71” (1803mm)**
Material Transit Aisle	**12’ (3658mm)**

Thermal Management & Cooling Limits

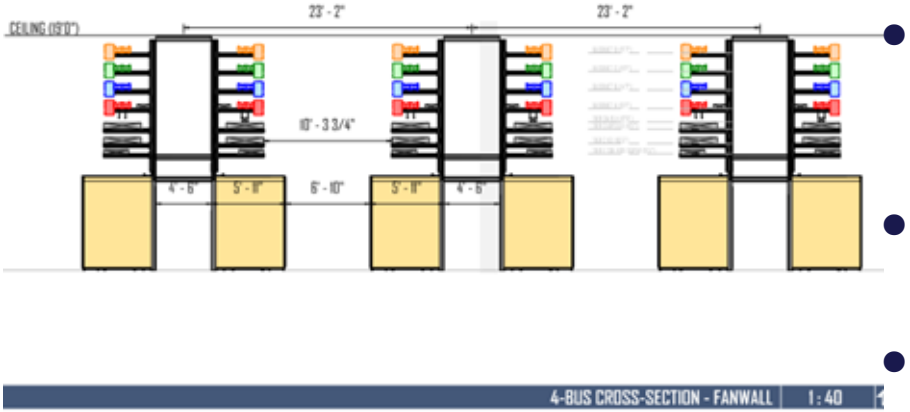
- Recommend a Peak IT Density power for sizing cooling for pod for ML racks
  - 335kW/ft liquid cooling and 15 kW/ft of air cooling
- Recommend a Blended power approach for sizing cooling for pod for network, storage, xPU
  - 25 kW/ft liquid cooling and 15 kW/ft of air cooling
- Recommend max fluid velocity 3 m/s based on ASHRAE standard
  - Multiple industry workstreams to test to enable support for higher velocities
- Power support is a function of fluid velocity, inlet TCS temperature, and delta T

Diameter (inches)	1lpm/kW (delta T of 15°C)	1.5 lpm/kW (delta T of 10°C)
1	91 KW	61 KW
2	365 KW	243 KW
3	820 kW	547 KW
4	1459 KW	972 KW

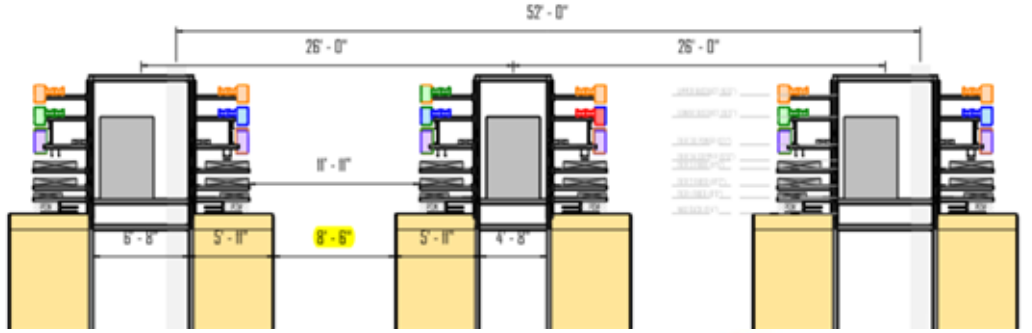
Note: 1. Fluid speed based on ASHRAE Handbook—Fundamentals, Chapter 33: Pipe Sizing specify the limit as 10fps (3m/sec)
2. Industry available quick connects currently available at maximum 2” and preferred connection type currently.

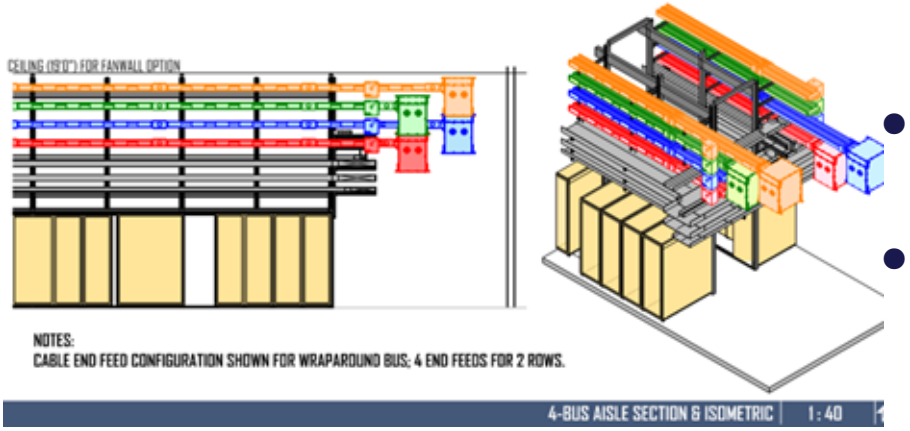
Potential Configurations - Google

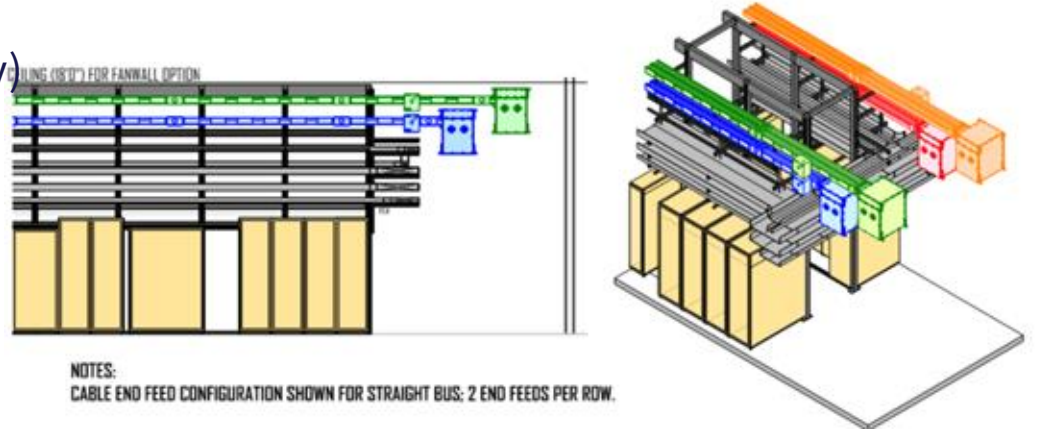
- AI standards designed to support different design manifestations



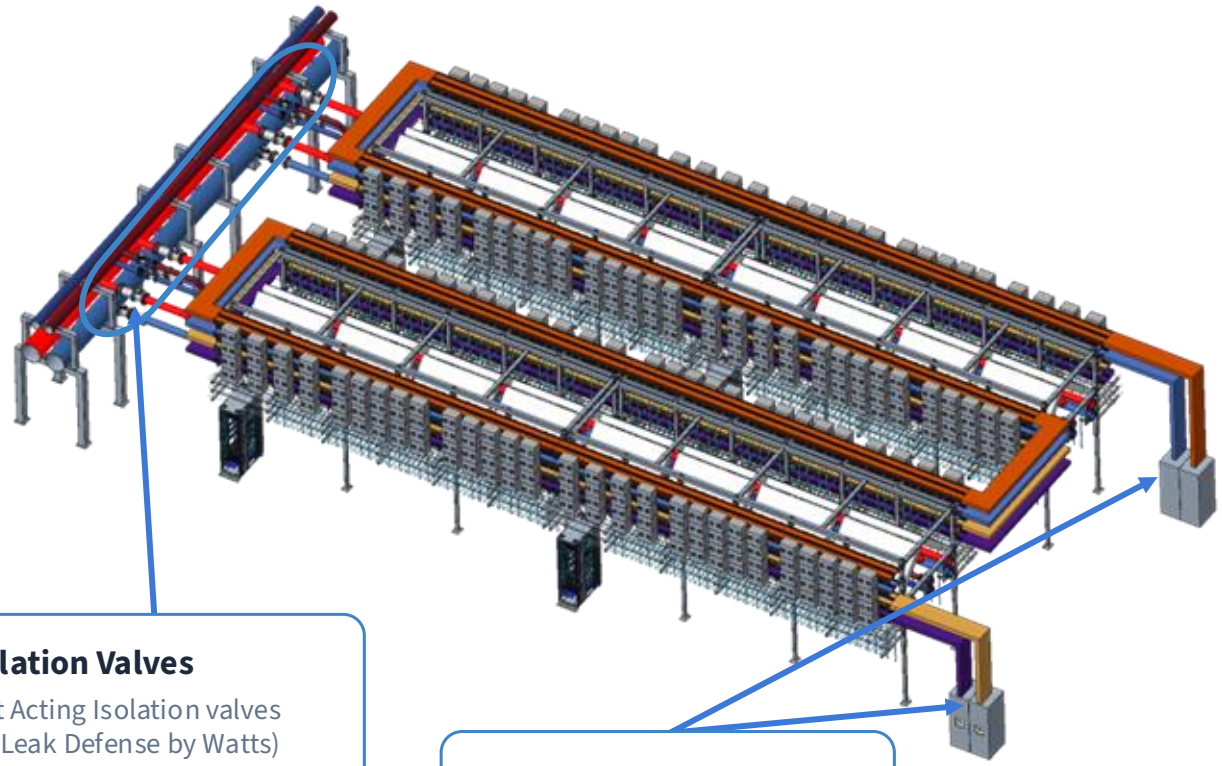
- 4-busway or 2-busway configuration → driven by ceiling height restrictions
- 23’ or 26’ aisle pitch depending on room constraints
- Row power density matching blended DC density of 25 kW/ft (blended ML, C&S, NW, Auxiliary)
- Rows Capable of supporting upto 800kW per IT racks
- Cooling capacity right sized to the power delivered within the row







Potential Configurations - NVIDIA (4 row, 12.5MW block)



Block Attributes

- Up to 120 – 24” Racks or 96-32” Racks
- Up to 4 MW of cooling per row
- 4- 5000 AMP Busways (800VDC transition)
- Rows are 68’ Long
- Space shown 98’ x 48’ (5000 SQFT)
- Pairs of rows are on 24’ pitch
- In-Row close coupled air cooling



Isolation Valves

Fast Acting Isolation valves
(IE: Leak Defense by Watts)

Power Feeds

4 Power Feeds
5000 Amp AC or Future 800VDC

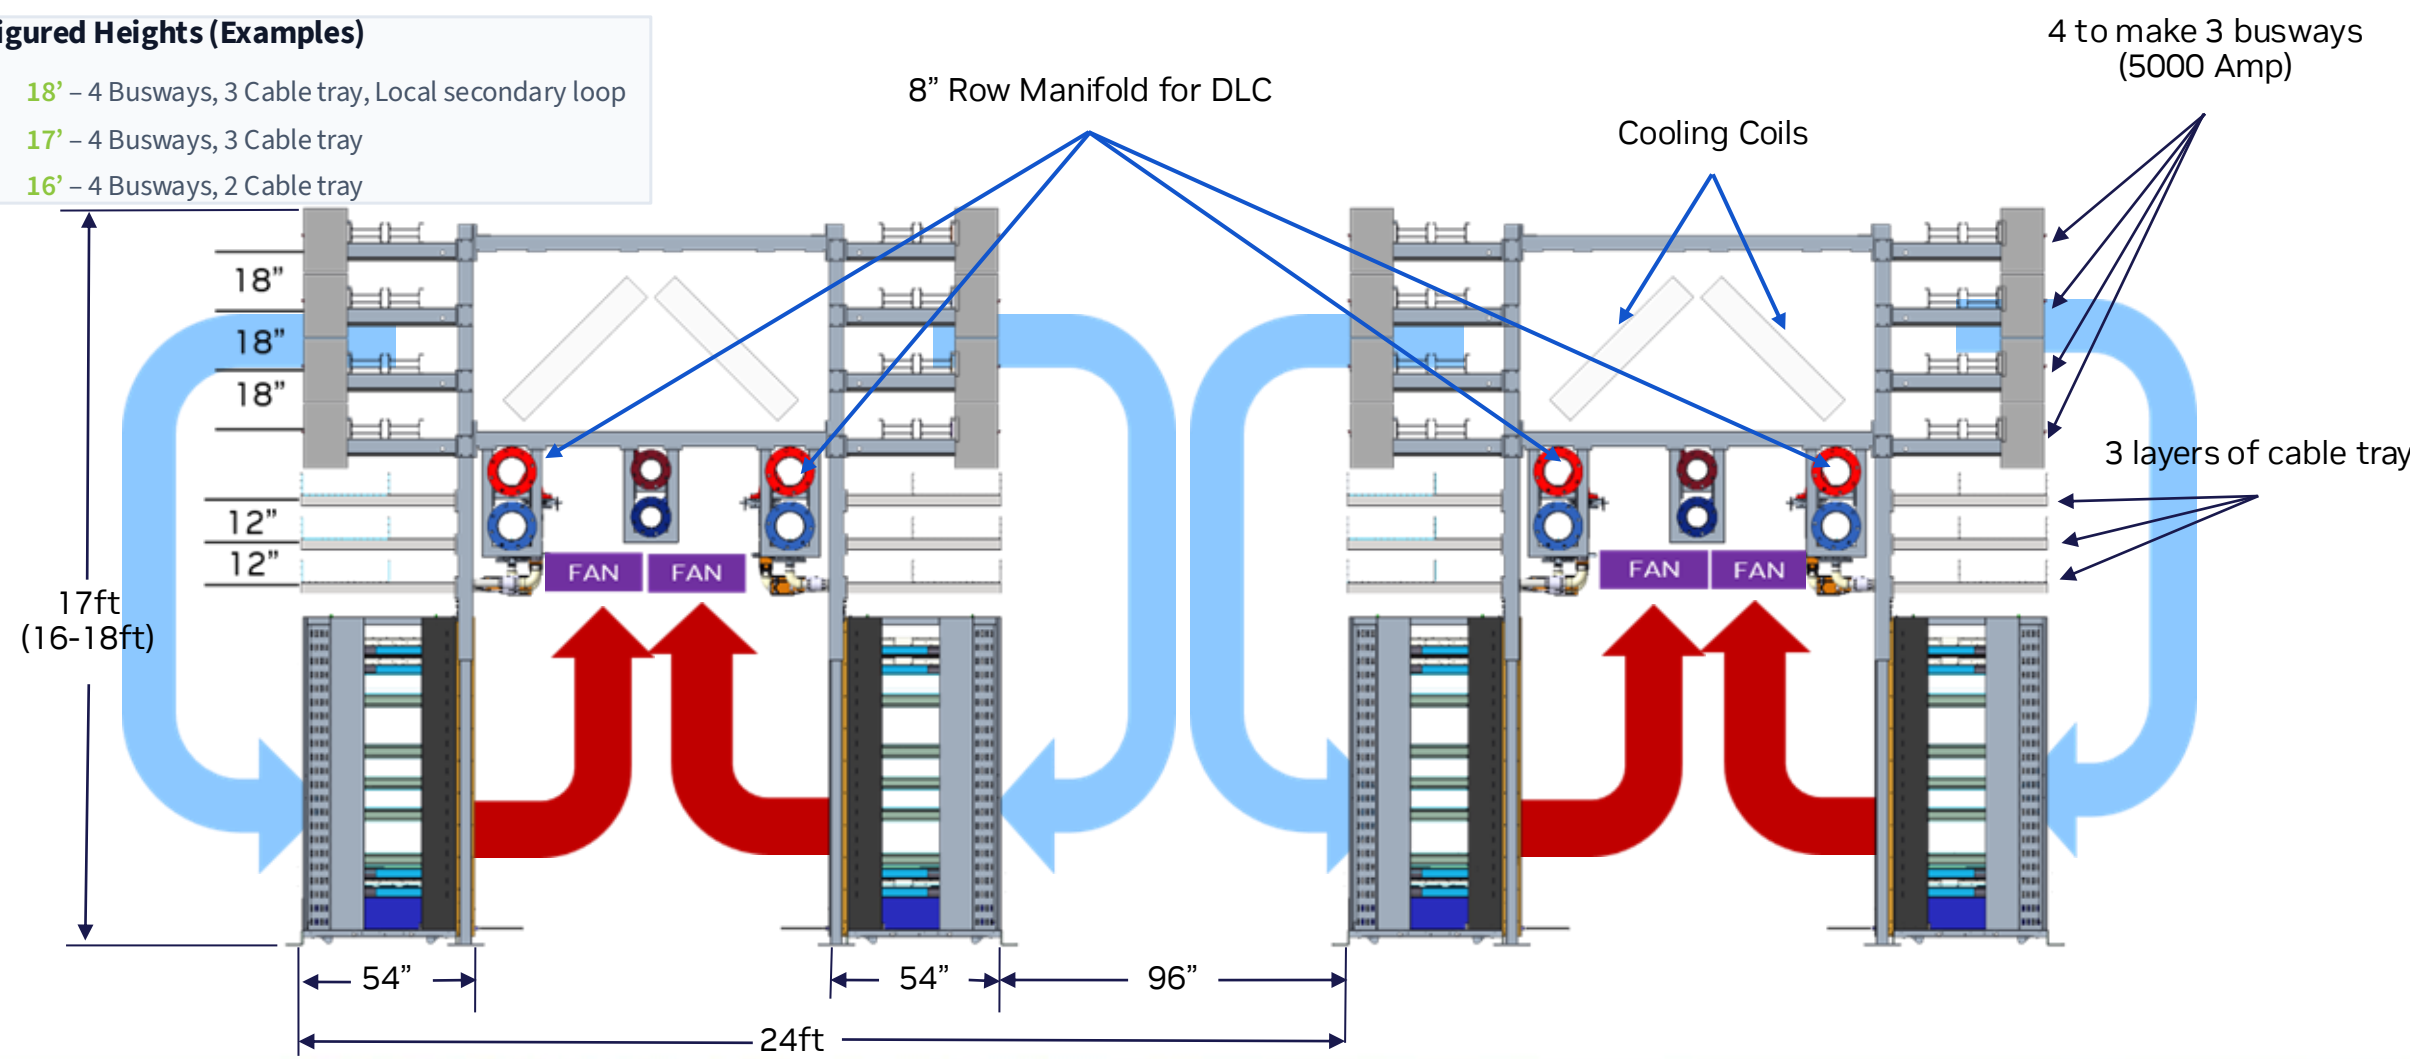
Water Feeds

Two Water Feeds:
- 12” – For Air-cooling
- 24” – Warm Water

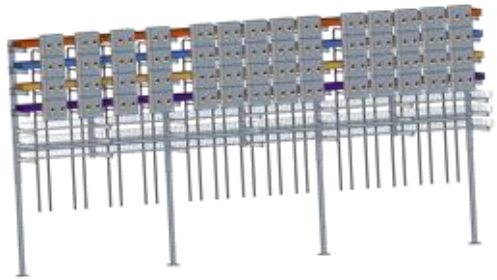
Close Coupled Cooling - NVIDIA

Configured Heights (Examples)

- 18’ – 4 Busways, 3 Cable tray, Local secondary loop
- 17’ – 4 Busways, 3 Cable tray
- 16’ – 4 Busways, 2 Cable tray



Building Blocks - Recommendations



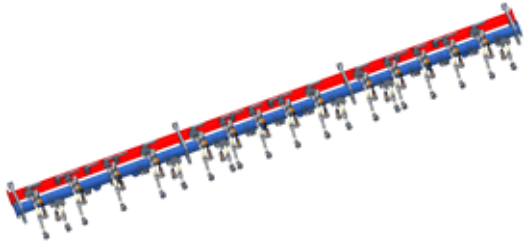
Row Infrastructure

- 16’–18’ high
- 34.6’ Long (example)
- 4 Busways (4 to make 3) or 2 Busways
- 3 network trays
- Supports secondary manifold



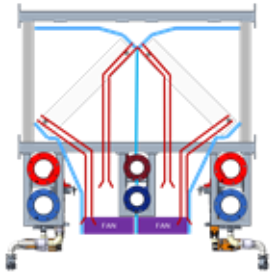
TAP Boxes

- Dual outputs
- Fits busways 18” apart
- Supports 415/480V AC (60A, 100A, 200A, 250A)
- 800VDC support (125A DC and 250A DC)
- Shunt trip support
- Direct power shelf connect



Row Manifolds

- Up to 4” connection (6 locations)
- Configurable for 2” or 4” sizes
- Integrated isolation valves
- Built-in drip pans
- In-rack-CDU support option



Hot Aisle Containment (NVIDIA)

- Supports 2 water temperatures
- 15kW Air Cooling / foot / side
- 6” main for Air (2.4 MW/pair)
- Dual 8” mains (4 MW/Row)
- Fully serviceable from underside
- 4200 CFM fans per foot

Call to Action

- Adopt the minimum AI standards in base building design to support next generation ML payloads
  - Help create specifications for the building blocks that meet these standards
  - Help standardize the sub-systems and components!
- Review v0.7 Open DC standard and provide feedback
- Monthly feedback sessions starting in November.
- v1.0 has future additions coming up - partnering with other workstreams to resolve next generation specifications
  - LVDC Workstream
  - Cooling Workstream

Thank You!

