




OPEN

Compute Project

Open Data Center: Next-Generation ML Infrastructure Design Principles

Revision 0.7.0

Design Specification

Effective Aug 27, 2026

Authors: AMD, Google, Meta, Nvidia

Version History

Date	Version #	Authors	Description
Dec 02, 2025	0.3.0	Google: Mohamed Aboueilella Ayan Majumdar Heather Huang Scott Sharp	Initial Release
Feb 25, 2026	0.5.0	Google: Mohamed Aboueilella Ayan Majumdar Heather Huang Scott Sharp Mike Meakins Rakesh Jain	Updated min row length; Row qty guidance; Added row power density requirement; Indoor relative humidity; Liquid to air ratio
June 23, 2026	0.7.0	Google Mohamed Aboueilella Ayan Majumdar Heather Huang Moe Ghamary Scott Sharp Mike Meakins Rakesh Jain Meta Steve Mills Mike Edie Nvidia Wade Vinson Mike Sabotta AMD Paul Artman	Collaborative 2027–2028 guidance unified across AMD, Google, Meta, and Nvidia, generalizing the scope for industry AI/ML systems. Removed Google-specific near-term tables, TPU pod scaling, campus layouts, environmental limits (e.g., 10%–38% RH), and appendices. Row layouts now use modular kW/ft density instead of rigid row lengths. Standardized rack height to 100", weight to 8,000 lbs max, and unified power and cooling limits

Table of Contents

Version History.....	2
1 License.....	4
2. Compliance with OCP Tenets.....	5
2.1 Openness.....	5
2.2 Efficiency.....	5
2.3 Impact.....	5
2.4 Scale.....	5
2.5 Sustainability.....	6
3. Specifications.....	7
4. Facility and Infrastructure Specifications.....	7
4.1 Structural and Weight requirements.....	7
4.2 Row Layout and Clearance.....	7
5. Power and Density Requirements.....	8
5.1 Cluster and Power Density Targets.....	8
5.2 Power Architecture and Interfaces.....	8
6. Thermal Management and Cooling.....	9
6.1 Data Center Requirements.....	9
6.2 Connectivity and Infrastructure.....	9
6.3 Rack Cooling Density Caps.....	9
Appendix.....	10

1 License

Contributions to this Specification are made under the terms and conditions set forth in **Modified Open Web Foundation Agreement 0.9 (OWFa 0.9)**. (As of October 16, 2024) (“Contribution License”) by:

AMD, Google, Meta, Nvidia

Usage of this Specification is governed by the terms and conditions set forth in **Modified OWFa 0.9 Final Specification Agreement (FSA)** (As of October 16, 2024) (“**Specification License**”).

You can review the applicable Specification License(s) referenced above by the contributors to this Specification on the OCP website at

<https://www.opencompute.org/contributions/templates-agreements>.

For actual executed copies of either agreement, please contact OCP directly.

Notes:

The above license does not apply to the Appendix or Appendices. The information in the Appendix or Appendices is for reference only and non-normative in nature.

NOT WITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

2. Compliance with OCP Tenets

2.1 Openness

Google, Nvidia, AMD and Meta are contributing this specification to help accelerate deployment of standardized next generation AI/ML technology. This contribution is aligned with the recently launched OCP Open DC Initiative. The specification includes data center design to accommodate design collateral for multiple generations of AI/ML systems.

The specification provides technical design details to enable the deployment of Machine Learning (ML) Systems and can be expanded to include industry GPU systems. One of the objectives of opening the specification to the industry is to collect end-to-end feedback to align to a standardized data center design that can be used widely in first and third party data centers.

2.2 Efficiency

This specification is targeted to support AI/ML systems targeting the 2027-2028 and beyond AI/ML systems. This specification includes high efficiency cooling (e.g. Project Deschutes) and power (e.g. Mt. Diablo) equipment. The design and sharing are intended to enable more efficient scaling of the next generation AI/ML workloads.

2.3 Impact

The specification enables any Colo or 3rd party supplier in the industry to develop, manufacture and improve upon the Open Data Center design which is informed by multiple years of scale experience. Furthermore, 1st party hyperscalers and 3rd party data center owners alike all over the world would be able to design/build a data center based on this specification.

2.4 Scale

A data center based on this specification is assembled from components that are sourced from multiple vendors that are widely known in the industry. Custom components are intended to be manufacturable by well-known processes. Beyond supply chain considerations, installation and maintenance procedures are expected to enable fast deployments of reliable equipment.

2.5 Sustainability

AI/ML systems utilize liquid cooling that generally requires less electrical power and physical data center space to cool IT hardware compared to air cooling. Also, high efficiency power equipment such as Mt. Diablo is designed to increase the data center power distribution efficiency through the use of higher voltage power supply (direct current) at the high power AI/ML rack. Furthermore, since the specification is available to Colo and 3rd party DC manufacturers globally, data center owners would be able to optimize the manufacturing based on sustainability metrics.

3. Specifications

This specification outlines the unified data center design standards developed in collaboration with Google, Meta, Nvidia, and AMD. This iteration focuses on the 2027-2028 timeline and emphasizes flexible, first principles guidelines to support next-generation AI and ML workloads. The standards herein govern structural, power, and cooling infrastructure to ensure industry wide adaptability and scalability. The specification covers details pertaining to HAC (Hot aisle containment) design, power density, cooling density and physical attributes of the data center.

4. Facility and Infrastructure Specifications

The physical constraints of the data center have been updated to accommodate high-density ML deployments while providing flexibility for different vendor requirements.

4.1 Structural and Weight requirements

- **Maximum Rack Weight:** The design target for the heaviest single rack is established at 8,000lbs (3629 kg) over 28" (711mm) W x 48" (1219mm) D. Racks would require specific handling and moving equipment whose weight needs to be accounted for in the total moving load.
- **Peak IT Weight:** Peak weight would be AC/DC sidecar + ML rack combination.
 - 7,000lbs (3175kg) (sidecar) + 8,000lbs (3629kg) (ML rack) = 15,000lbs (6804kg) per (56" (1422mm) W x 48" (1219mm) D)

4.2 Row Layout and Clearance

- **Transport Aisle Clearance:** A 12 ft (3.7m) aisle clearance is set as a design recommendation.
- **Row Length:** Rigid row length requirements have been replaced with modular designs using power density (kW/ft) to better support evolving power and fluid connectivity needs.
- **Ceiling Height:** An 18ft (5.5m) minimum ceiling height is required for 2-busway configurations. This is measured from the floor to the bottom of the beam and includes air plenum space. Configurations utilizing 4-busway setups with different configuration and additional tiers of tray may require additional height beyond 18ft (5.5m). Custom design variations are necessary for facilities unable to meet the 18ft (5.5m) standard.
- **Aisle Pitch:** The recommended aisle pitch based on the below minimums is 24ft (7.3m). Racks are front justified in the cold aisle.

Additional requirements and row characteristics in Table 1.

Table 1: Minimum dimensions for data hall rows

Metric (minimum specified below)	Target
Hot Aisle	54" (1372mm)
Cold Aisle (IT rack face to face)	82" (2083mm)
Rack Height	100" (2540mm)
Rack Depth	71" (1803mm)
Material Transit/Transport Aisle	12' (3658mm)

5. Power and Density Requirements

To support rapidly scaling ML models and hardware generation shifts (2027 through 2028), specific power and cluster size benchmarks have been unified. Peak IT density represents the highest powered ML racks. The blended density incorporates other workloads including network, general compute, and storage systems.

5.1 Cluster and Power Density Targets

Table 2: Cluster and Power Density Targets

Metric	Target
Blended Data Center Capacity Baseline	25 kW/ft (82 kW/m) blended average, which will require 15 kW/ft (49 kW/m) air cooling capability and 25 kW/ft (82 kW/m) liquid cooling capacity
Peak IT Density	350 kW/ft (1148 kW/m). In 2027, peak IT for racks is expected to be 125 kW/ft (410 kW/m)
Cluster Size Baseline	12 MW Minimum, 24 MW preferred

5.2 Power Architecture and Interfaces

- **Power Lineups:** Each data hall must provide at least 4 power blocks. The specific number of busways per row can vary depending on the design and final row power density desired.

6. Thermal Management and Cooling

Efficient cooling mechanisms are critical for the elevated IT densities projected in the upcoming hardware cycles. Standards have been set for both air and liquid cooling integrations. This is divided in order to distinguish between data center row/ cluster requirements from rack level interface and cooling density caps

6.1 Data Center Requirements

In line with section 5.1 blended data center capacity, cluster cooling should be designed for up to 25 kW/ft (82 kW/m) liquid cooling and 15 kW/ft (49 kW/m) of air cooling by 2028. In order to achieve this design team must appropriately size the wall manifolds to support the entire length of the machine row.

6.2 Connectivity and Infrastructure

- **Manifold Design:** Wall/ row manifold is to be sized to meet the cooling requirements for the row. Maximizing number of drops for multiple connections in racks with multiple cooling loops per rack.
- **Interface Optimization:** Manifold interface windows have been optimized to enhance equipment accessibility and maintenance operations.

Table 3: Maximum cooling limit in kW for various dimensions

	Speed (m/sec)=	3 ¹
Diameter (inches)	1lpm/kW	1.5 lpm/kW
1	91	61
2	365	243
3 ²	820	547
4 ²	1459	972

2. Industry available quick connects currently available at maximum 2"

6.3 Rack Cooling Density Caps

As peak IT is expected to be at 350 kW/ft (1148 kW/m) for high density machine learning racks, the liquid cooling caps are expected to be 335 kW/ft (1100 kW/m) +15kW/ft (49 kW/m) of air cooling. Liquid cooling hence needs to be sized to support cooling requirements in high density zones of the floor where ML racks are positioned.

Appendix

Typical row layout with multiple rack depths (Google)




Figure 1: Example layout of a HAC with the 2028+ TPU specifications

In order to fill the gaps for shallower racks and seal against the hot aisle containment panels, gaps are filled with filler panels as shown in Figure 2.




Figure 2: Filler Panels installed in HAC gaps

Typical row layout TPU versions (Google)

TPU row layout v9




Figure 3: Example layout of a HAC with the 2028+ TPU specifications

TPU row layout v8




Figure 4: Example layout of a HAC with the 2026-2027 TPU specifications

Reference Next Gen HAC design (Google)

Two Busbar per row configuration

The below design shows the HAC structure with 2 busbars per row/ 4 busbars per HAC.




Figure 5: Example HAC cross section for 2026 and early 2027 DCs with fanwall cooling.




Figure 6: Example HAC cross section for 2026 and early 2027 DCs with In-row cooling and building columns




Figure 7: Example HAC cross section and isometric view of HAC for 2026 and early 2027 DCs with 2 busway configuration

Four Busbar per row configuration




Figure 8: Example HAC cross section and isometric view of HAC for 2027-2028 DCs with a 4 busway configuration

Reference Next Gen HAC design (Nvidia)




Figure 9: Example HAC configuration concept from Nvidia with 4 busways

Multi – Row (12.5MW)

- Multiple Configured heights: (examples)
- 18' - 4 Busways, 3 Cable tray, Local secondary loop
 - 17' - 4 Busways, 3 Cable tray,
 - 16' - 4 Busways, 2 Cable tray

Close coupled Air Cooling




Figure 10: Example HAC cross section concept from Nvidia with 4 busways




Figure 11: Example HAC concept cross section concept from Nvidia with in-row cooling




Figure 12: Example concept power network module and Row Manifolds




Figure 13: Example Row Manifolds showing 4" tap location for high density ML

Support Both AC and 800VDC on Common Structure




Figure 14: Example concept Nvidia busway with row to row connectors

Hall Views




Figure 15: Example concept busway configuration across 2 HACs / 4 rows