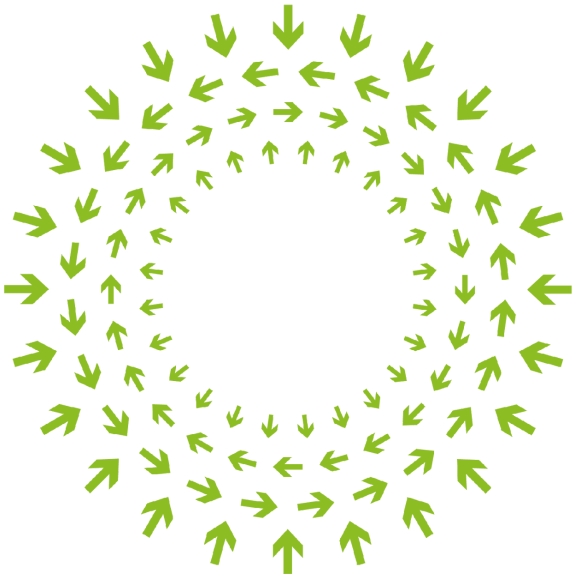


OPEN

Compute Project

Datacenter Design Guide : Next-Generation ML Infrastructure Design Principles 0.30.0

Specification

Effective Dec 2, 2025

Authors: Google

Version History

Date	Version #	Authors	Description
Dec 02	0.30.0	**Google:** Mohamed Abouelella Liwen Mao Ayan Majumdar Heather Huang	Initial Release

Table of Contents

Version History....2
1 License....3
1.1 Open Web Foundation (OWF) CLA....3
Acknowledgements....4
2. Compliance with OCP Tenets....4
2.1 Openness....4
2.2 Efficiency....4
2.3 Impact....5
2.4 Scale....5
2.5 Sustainability....5
3 Specifications....6
3.1 Serverhall Specification....6
3.1.1 Specifications for 2026 and Early 2027....6
3.1.2 Specifications for Late 2027/2028 and beyond (2028+)....12
3.2 CDU and Sidecar specification....17
4 DC Campus Layout Examples....17
4.1 : Cluster Layout, DC Serverhall Elevation....18

1 License

THE UPDATED DEFAULT CONTRIBUTOR LICENSE AGREEMENT (CLA) IS **OWFa 0.9**. PLEASE VERIFY THE CORRECT CLA/FSA IS USED AND EXECUTED FOR THIS CONTRIBUTION.

1.1 Open Web Foundation (OWF) CLA

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:

Google

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.

NOTWITHSTANDING 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.

Acknowledgements

The Contributors of this Specification would like to acknowledge the following for their feedback:

Google:

2. Compliance with OCP Tenets

2.1 Openness

Google is contributing Open Data Center specification to help accelerate deployment of standardized next generation AI/ML technology wide. This is the first Data Center specification submission to OCP, and is aligned with the recently launched OCP Open DC Initiative. The specification will include 2026 and 2027 versions of 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 Google's TPU 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

The Open Data center spec is targeted to support AI/ML systems targeting the 2026/2027 and beyond AI/ML systems. The data center spec will include 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 will enable 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 at Google. 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 the Open DC specification.

2.4 Scale

The Open DC is assembled from components that are sourced from multiple vendors that are widely known in the industry. Custom components are intended to be manufactureable by well-known processes. Beyond supply chain considerations, installation and maintenance procedures are expected to enable fast deployments of reliable equipment. The processes and best practices shared by Google are based on the 7+ years of scale liquid cooling data center experience.

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 defines the Data center requirements for the near term (2026-27) and the longer term (2027-28 and beyond). Its primary objective is to define the interface specifications for data center designs. The specification covers details pertaining to HAC (Hot aisle containment) design, cable conveyance attributes, availability prerequisites, and a comprehensive Bill of Design (BoD) guide.

The specification provides technical design details to enable the deployment of Google's TPU Systems and can be expanded to include industry GPU systems.

3.1 Serverhall Specification

- This section outlines the minimum facility and infrastructure requirements necessary to support high-density, mega-scale Machine Learning (ML) hardware deployments, focusing on physical dimensions, power distribution, and thermal management.

3.1.1 Specifications for 2026 and Early 2027

In this section we cover the specifications for the datahalls in both the "small" and "large" size datacenter designs. The delineation is made on capacity of the datahall <=20MW being considered "small" whereas anything >20MW is considered large. This is a rough delineation to help determine when to deploy the designs listed.

The "Small" DC design offers smaller cluster size options specifically in constrained regions and buildings. These buildings also might be power constrained and hence utilizing the space effectively to enable high dense ML deployments becomes a key capability.

The "Large" DC design provides the highly optimized large scale DC design option where space and power are considered less of a constraint. Enabling scale and optimizations is the key feature of these designs and they serve to enable a large scale ML deployments. These specifications can support DCs that are capable of supporting more than 20MW.

The specification provides technical design details to enable the deployment of Google's TPU Systems and can be expanded to include industry GPU systems.

	"Small" <=20MW	**"Large" >20MW**
Rack Depth	1219mm/48"	1321mm/52"
Thermal requirements- Air	Min. 1440 CFM/ft	Min. 1440 CFM/ft

cool		
Thermal requirements- Liquid cool	See Project Deschutes for interface	See Project Deschutes for interface
Liquid:Air Ratio	TBD	TBD
Rack Height	90” max height	95” max height
Floor loading	3300lbs (max rolling weight) CDU weight : 3500lbs	6000lbs (rolling load) max weight per rack 350 psf live load
Cold Aisle width	1372mm/54”	min. 1524mm/60”
Hot Aisle width	1220mm/48”	min. 1371mm/54”
Clear ceiling height	Min 14’	Min. 18’/5.47m
Row length	46’ a. 6ft/1800mm clearance in at least 1 side of the row for transit b. 4ft/1200mm clearance in other side of the row	2x 72 ft or 1 x- 96ft
Min # of rows	16 rows	Each DC hall should have a minimum of 12 machine rows to meet capacity and diversity requirements, 36 rows standard
Google CDU rack dimensions	Width : 56” (max) Height : 95” (max) Depth : 40.5” Weight : ~3500lbs	Refer to project Deschutes for thermal interface Width : 56” (max) Height : 95” (max) Depth : 40.5” Weight : ~4000lbs
# of feeds, redundancy	4 feeds with independent UPS support,	2N connection from 4N3DR topology (5n4DR maybe

	3+1 redundancy (aka “4 makes 3”)	acceptable with enough space in data hall for 10 rows/power plan)
Tap box protection	32A 3ph breakers (2 outlets per feed per rack spot)	Hard wire connection 100A or 200A (depends on the rack)
Row level power	750kW min per row	Power distribution shall be designed to support 1MW min per row and up to 19-20kW/ft (60kw/m) Data hall design space/row length shall be designed with 9kw/ft power density to provide a minimum of 50% additional rack space over average power density is recommended to enable future oversubscription
Facility -> Load power architecture - **2N feed redundancy** - **N+1 feed redundancy**	Feed redundancy at rack = 3+1 UPS redundancy >= Feed redundancy Generator redundancy>= Feed Redundancy	No central UPS with full generator space backup and No Li-Ion battery allowed for central UPS (Ni-zinc is ok)
ML Load profile	None.	Generator shall be capable of supporting fluctuation of ML loads up to 50% of generator rating routinely 30-45 second intervals with one sec duration
Minimum Data hall power	~6MW	12 MW

Reference design

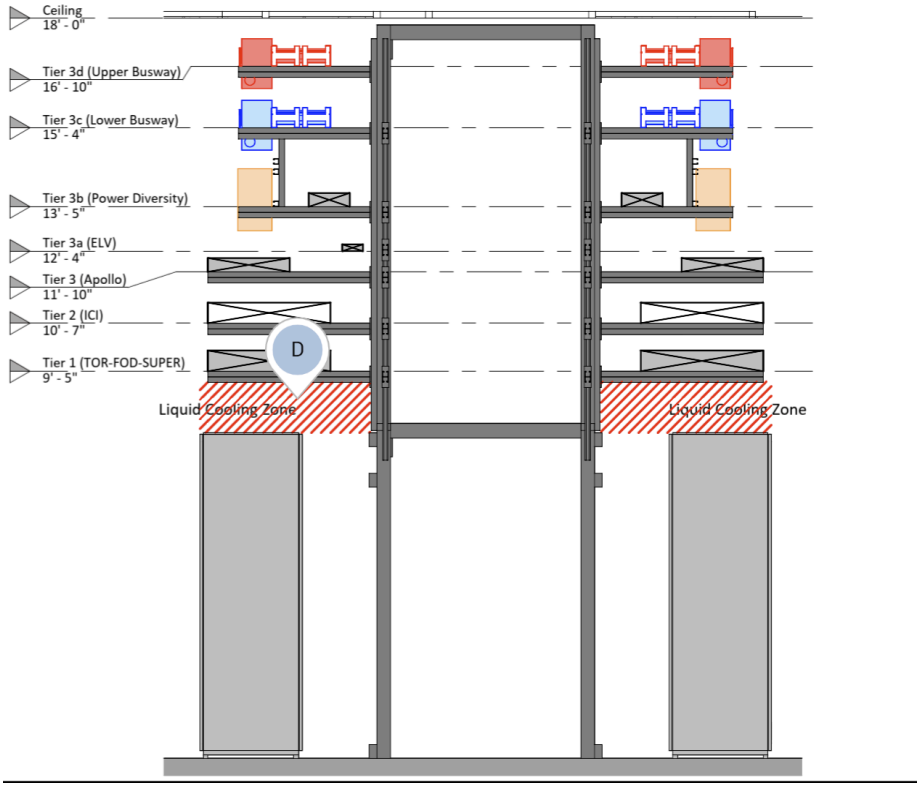


Figure 1:Example HAC cross section for the "Large" Datacenters for 2026 and early 2027 DCs

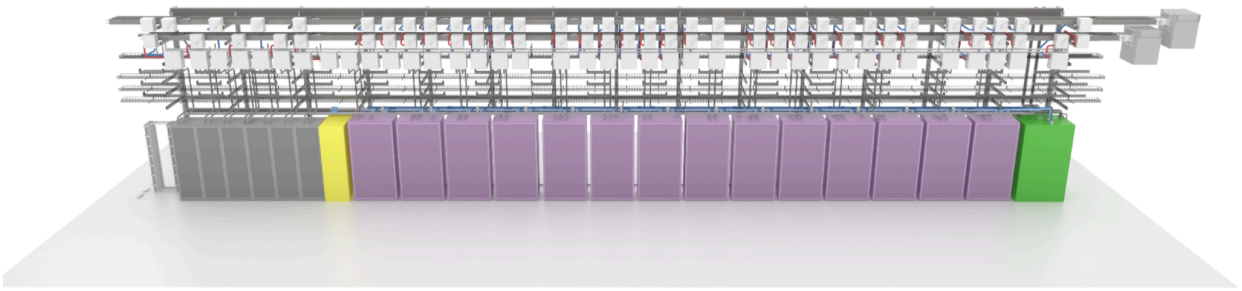


Figure 2: Example HAC cross section for the "Large" Datacenters for 2026 and early 2027 DCs

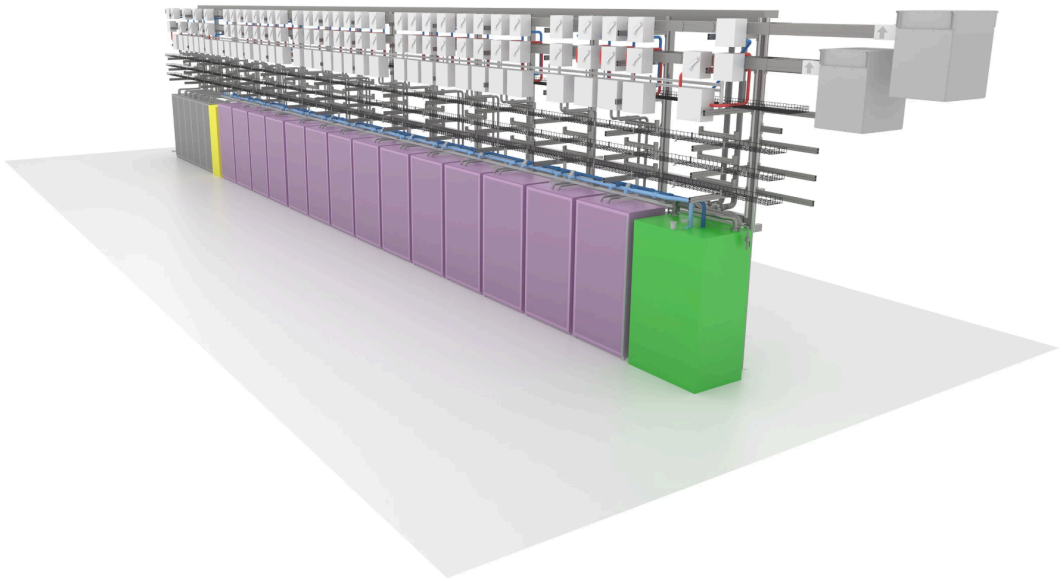


Figure 3: Example 3D renderings of the “Large” scale DC designs, showing the capability to deploy dense ML deployments at scale.



Figure 4: Ironwood deployment in one of the “large” scale DC designs

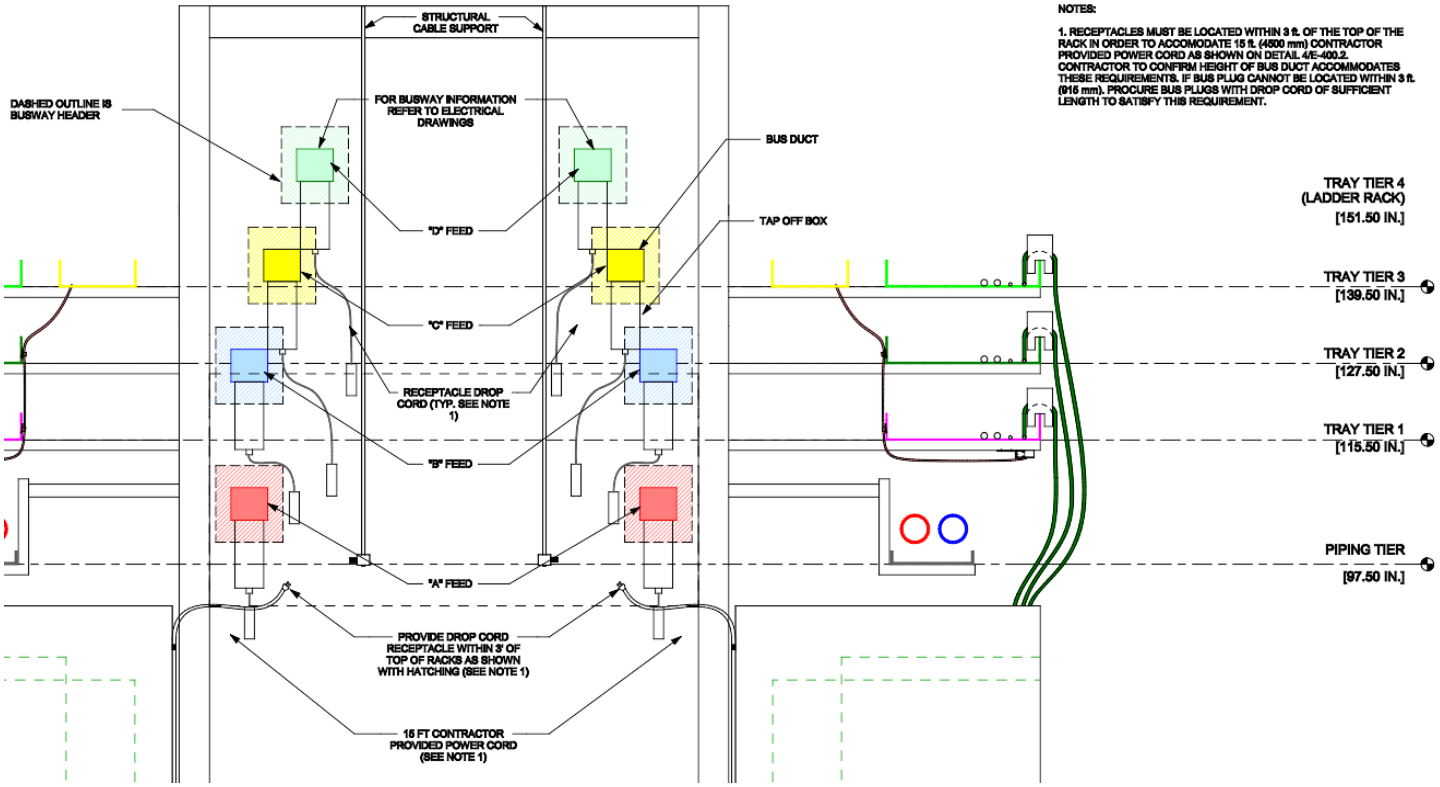


Figure 5: Example HAC cross section for the "Small" Datacenters for 2026 and early 2027 DCs

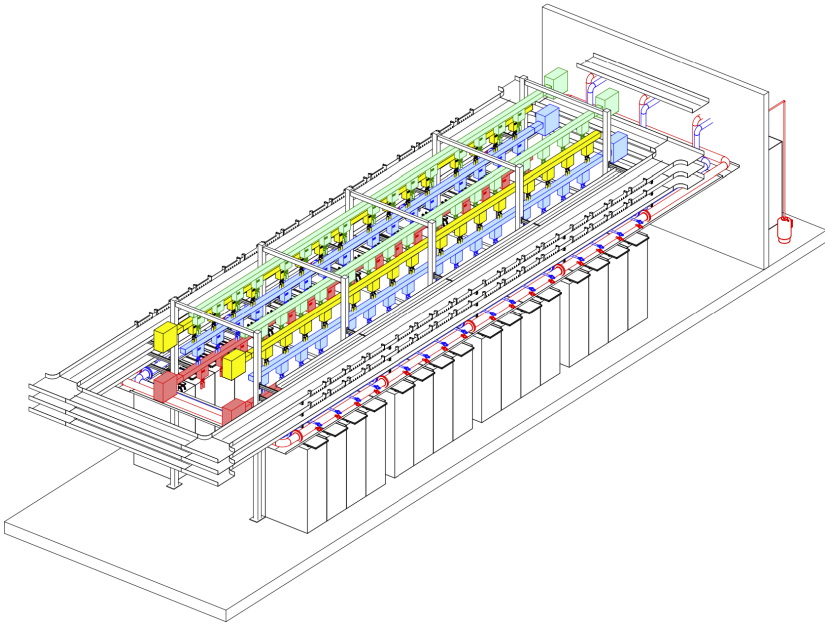


Figure 6: Example HAC cross section for the "Small" Datacenters for 2026 and early 2027 DCs

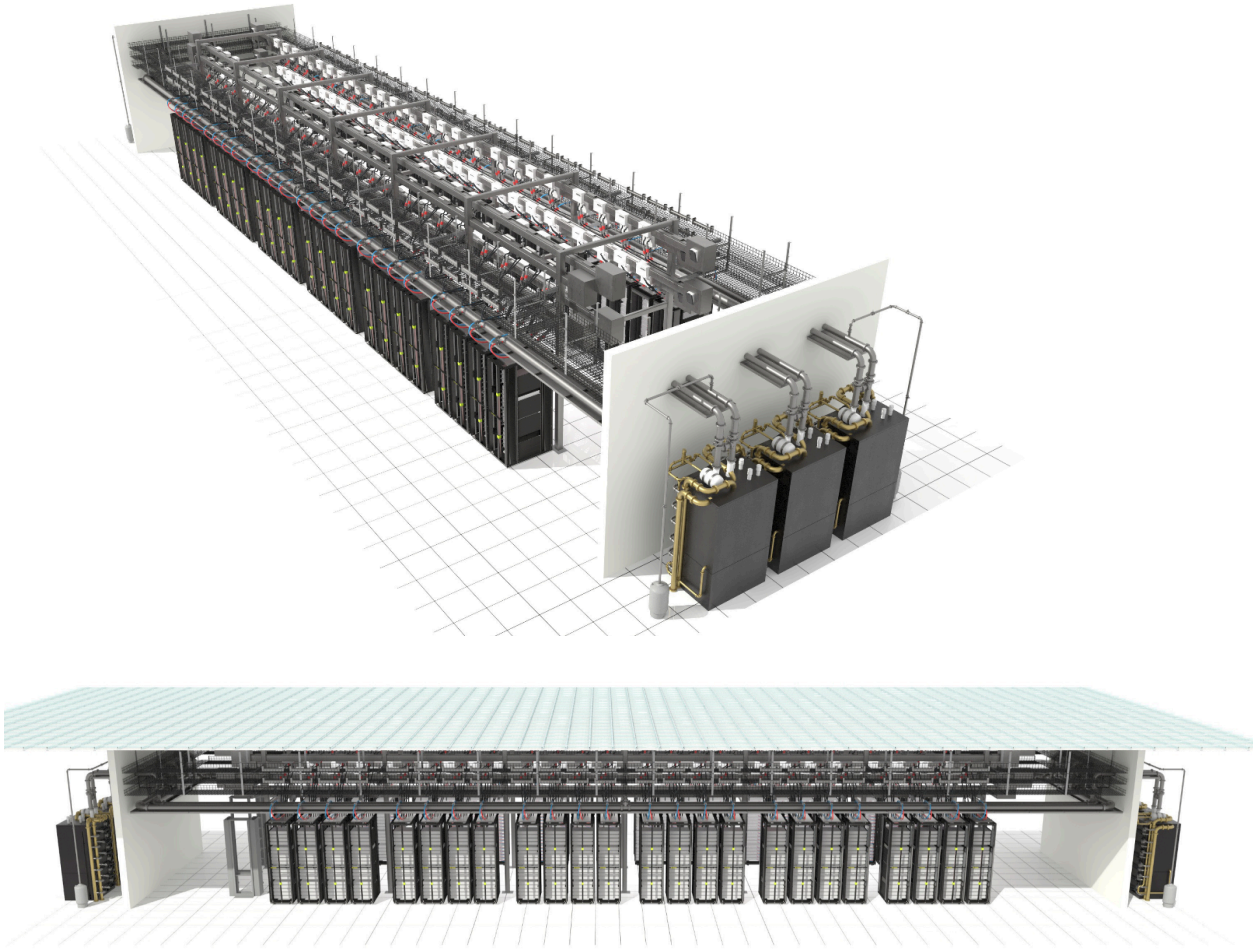


Figure 7:Example 3D renderings of the “Small” scale DC designs, showing the capability to enable high dense ML deployments in constrained spaces.

3.1.2 Specifications for Late 2027/2028 and beyond (2028+)

After 2027, with the continued increase in ML density and interconnected nature of deployments, we see the specifications needing to evolve to support denser platforms and enable more interconnected deployments including at the campus level to provide higher performance.

While these are designed with high dense ML platforms in mind, the specifications enable backwards compatibility as well i.e., platforms designed to 2026/2027 specifications can still be deployed in Datacenters designed with these specifications.

Detailed design guides, HAC cross-sections and conveyance attributes, floor and room utilization guidance will be added to the specification in a later revision. The floor and room utilization details are highly dependent on the room size. The scaling table below attempts to capture some of these details but additional guidance on kW/sqm and sizing will be shared in later revisions. A schematic version is shown below for reference on how these designs might work with different rack form factors and densities.

The specification provides technical design details to enable the deployment of Google's TPU Systems and can be expanded to include industry GPU systems.

Time frame	**Late 2027-2028+**
Rack Depth	1804mm/71" deep
Floor density (kW/ ft)	TBD, will publish guidance for space vs power guidance in next revision targeted for Q1 2026
Thermal requirements- Air cool	Min. 125 cfm/kW, continuous rear rack service support
Thermal requirements- Liquid cool	See Project Deschutes for interface
Liquid:Air Ratio	TBD
Rack Height	95" max rack height + 5" clearance for cable management
Floor loading	7000 lbs (max rolling weight) Min: 350 psf DCH design live load
Cold Aisle width	Min 2083mm/ 82" (*preferred)* *(54" is tolerable based on 16 kW/ft or less, will require enablement of parallel parking and longer row length)*
Hot Aisle width	Min 1372mm/ 54" for Standard parking
Hot Aisle height	N/A
Row length	72' (without columns) a. Minimal 10ft transport aisle clearance in at least 1 row.

	b. 6ft clearance in other transport row c. Preferred 12ft transit aisle + 10ft transit aisle
Clear Ceiling height	18ft (for 2N busway)
Min # of rows	26 rows (ML pod) *(based on a 50' row length)*
Google CDU rack dimensions	Refer to project Deschutes for thermal interface Width: 65.25" Height: 91.44" Depth: 40.9" Weight: ~7000 lbs
Sidecar	Refer to Diablo 400 Project
# of feeds, redundancy	4/3DR busducts per machine row Min: 2 busways with 200A MTS per row Preferred: 4 busways per row
Tap box protection	Hard wire connection 3ph - 200A with connectorized plug after FR - Q4, 2027
Row level power	2.6MW ~ 2.4MW min per row
Facility -> Load power architecture - **2N feed redundancy** - **N+1 feed redundancy**	Feed redundancy at rack = 3+1 Campus planning should always plan for sufficient space for full generators. Deploy Generator as default. Revisit if the generator can be removed case by case.
ML Load profile	BESS is required for load smoothing for every DC building with campus utility draws of greater than or equal to 600 MWs Min. 5min backup from BESS, UPS or generator Load profile(s) to be included in the next revision targeted for Q1 2026.
SLO	25%-50% of capacity with 99.998%

	Rest capacity min. 99% Details will be included in the next revision targeted for Q1 2026.
Dock Throughput	minimum throughput of 150 racks per day, ideal spec TBD
Minimum Cluster size	See scaling table below

The scaling table below outlines the power and rack space requirements for common numbers of chips, providing a quick reference for essential data. The power and rack quantity accounts not only for ML racks but ***all*** the supporting racks required to operate the ML machines, including power sidecar, CDU (if deployed in the serverhall), UPS (if required), Networking and Compute & Storage deployments required to serve ML.

The fiber distance between a network rack and server rack is preferred to be less than 20m, additional fiber guidance and server row placement, will be added in future revisions.

Power = Peak IT
Rack = 28” Ext. Width Rack Spot

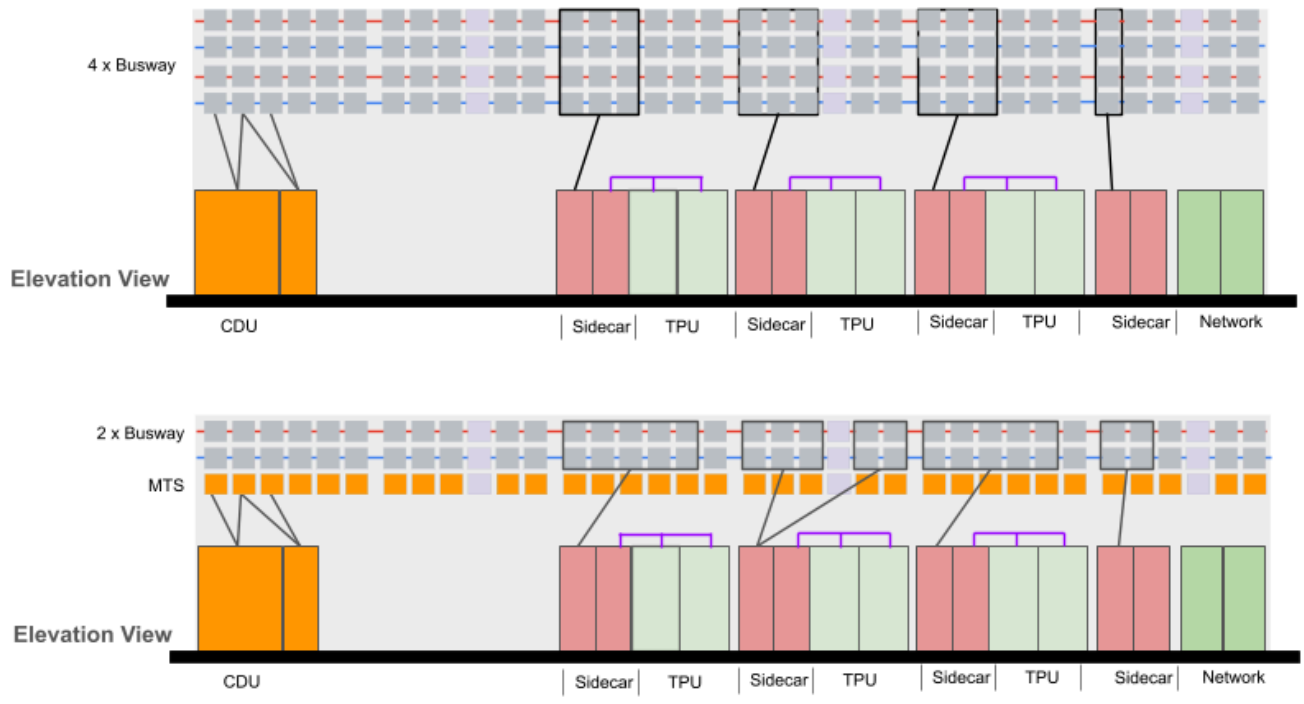
	GPU	
Chip Qty	Power	Rack
XXk	TBD	
ML:CS:NW		

	V7x (Now)		Future version (potential up to 2028/29)	
Pod Qty	Power	Rack	Power	Rack
1 Pod	17 MW	800	52 MW	1170
ML:CS:NW	7:1:2	4:2:3	8:1:1	4:3:3
2 Pods	29 MW	1300	99 MW	2100
ML:CS:NW	8:1:1	5:2:3	8:1:1	4:3:3
3 Pods	41 MW	1800	146 MW	2900
ML:CS:NW	8:1:1	6:2:2	8:1:1	5:3:3
4 Pods	53 MW	2300	193 MW	3800
ML:CS:NW	8:1:1	6:2:2	8:1:1	5:3:2

5 Pods	65 MW	2800	240 MW	4700
ML:CS:NW	8:1:1	6:2:2	8:1:1	5:3:2
6 Pods	77 MW	3300	287 MW	5500
ML:CS:NW	8:1:1	6:2:2	8:1:1	5:3:2
7 Pods	90 MW	3800	334 MW	6400
ML:CS:NW	9:1:1	6:2:2	8:1:1	5:3:2
8 Pods	102 MW	4200	381 MW	7200
ML:CS:NW	9:1:1	6:2:2	8:1:1	5:3:2

Notes:

1. % distribution and capacity are baseline estimation, NW% and C&S% could subject to change with IT bandwidth, customer C&S requirement
2. CDU and Side car are categorized as ML racks



Example layout of a HAC with the 2028+ TPU specifications

3.2 CDU and Sidecar specification

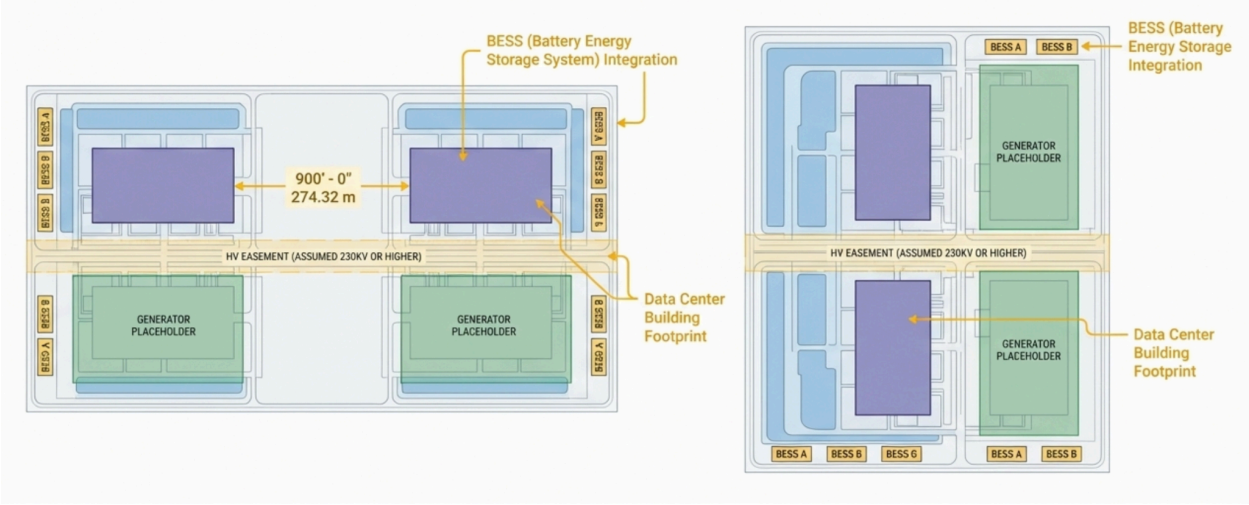
CDU specifications to support liquid cooling are noted in the Project Deschutes specification published in OCP and contain the minimum requirements for a CDU to be capable of supporting liquid cooling.

Similarly for products requiring native LVDC interface, a sidecar will be required. The sidecar specification is as per Diablo 400 project listed in OCP (Refer to Diablo 400 Project).

4 Data center Layout Examples

4.1 : Campus level layout for mega-scale

The campus layout must accommodate the mega-scale requirements of next-generation ML clusters, including total utility power needs and strict networking latency constraints that impact the physical distance between buildings.



Cluster/Building IT Requirement	Up to ~1.4GW IT load per campus in less than 4 buildings.
Campus Minimal Rack Spot Count	(13k) 28” Ext. Width Rack Spot
Worst case ML chip to chip fiber length	1.5km Minimal building distance will vary based on the actual layout

4.2 : Cluster Layout, DC Serverhall Elevation

This section defines the building blocks for ML clusters (rows, stamps, and pods) and the resultant power and rack space required for various deployment scales.

1. Data Halls
2. Electrical Rooms
3. Main Network room
4. Support spaces
5. Space for generators
6. Office space
7. Loading docks

