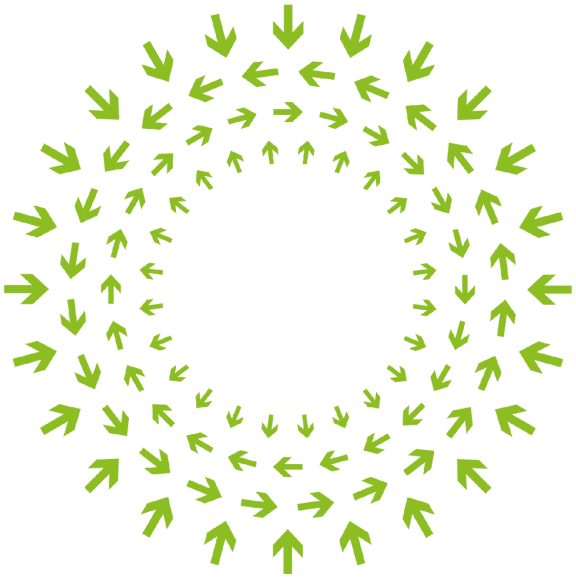


OPEN Compute Project

Datacenter: Next-Generation ML Infrastructure Design Principles

0.5.0

Design Specification
Effective Feb, 2026

Authors: Google

Version History

Date	Version #	Authors	Description
Dec 02, 2025	0.3.0	**Google:** Mohamed Abouelella Ayan Majumdar Heather Huang Scott Sharp	Initial Release
Feb 25, 2026	0.5.0	Mohamed Abouelella Ayan Majumdar Heather Huang Scott Sharp Mike Meakins	Updated min row length; Row qty guidance; Added row power density requirement; Indoor relative humidity; Liquid to air ratio

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
3.1 Serverhall Specification....7
3.1.1 Specifications for 2026 and Early 2027....7
3.1.2 Specifications for Late 2027/2028 and beyond (2028+)....13
3.2 CDU and Sidecar specification....18
4 Data center Layout Examples....18
4.1 : Campus level layout for mega-scale....18
4.2 : Cluster Layout, DC Serverhall Elevation....20
5. Appendix....21
5.1 Tradeoff analysis....21
● Row length & min # of rows....21
● Row power....22
● Aisle pitch....22
● Contiguous rows/Column positions....23
5.2 Design Process....24
6. Future Additions....26

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:

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.

2. Compliance with OCP Tenets

2.1 Openness

Google is contributing this specification to help accelerate deployment of standardized next generation AI/ML technology wide. 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 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

This spec is targeted to support AI/ML systems targeting the 2026/2027 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 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 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 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 specifications here are specific to ML related deployments. While they can be used for deploying regular Cloud C&S workloads, they are not optimized for it.

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 cool	Min. 1440 CFM/ft	Min. 1440 CFM/ft
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	3300 lbs (max rolling weight) CDU weight : 3500 lbs	6000 lbs (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 ft a. 6ft/1800mm clearance in at least 1 side of the row for transit b. 4ft/1200mm clearance in other side of the row	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 : ~3500 lbs	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, 3+1 redundancy (aka “4 makes 3”)	2N connection from 4N3DR topology (5n4DR maybe 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)
DCH indoor relative humidity	10% - 38%	10% - 38%
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

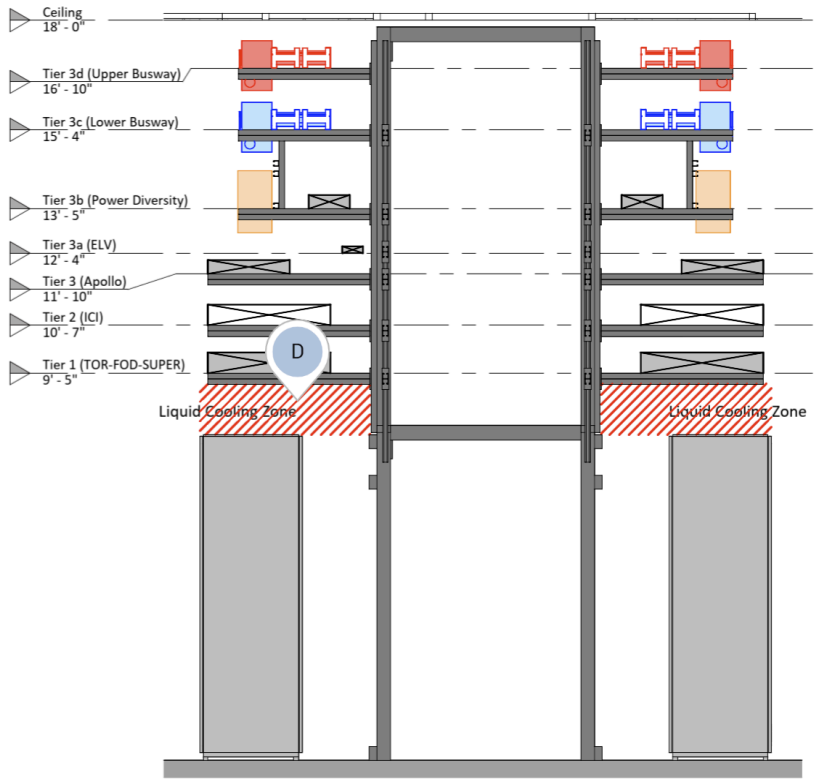


Figure 1:Example HAC cross section for the "Large" Datacenters for 2026 and early 2027 DCs

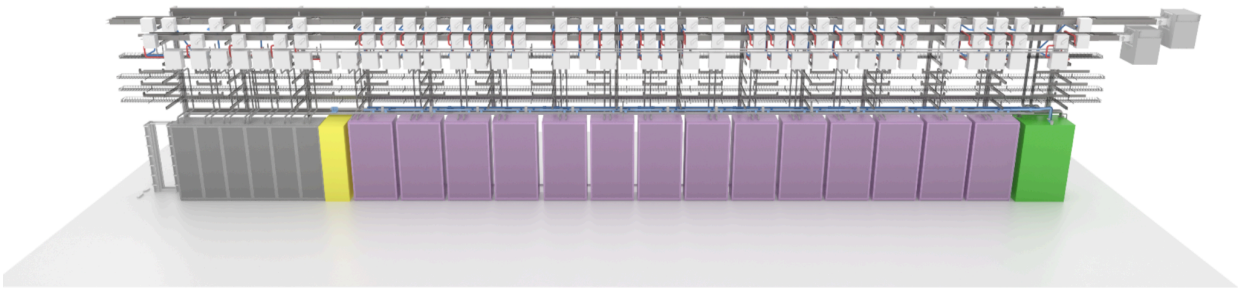


Figure 2: Example HAC cross section for the "Large" Datacenters for 2026 and early 2027 DCs

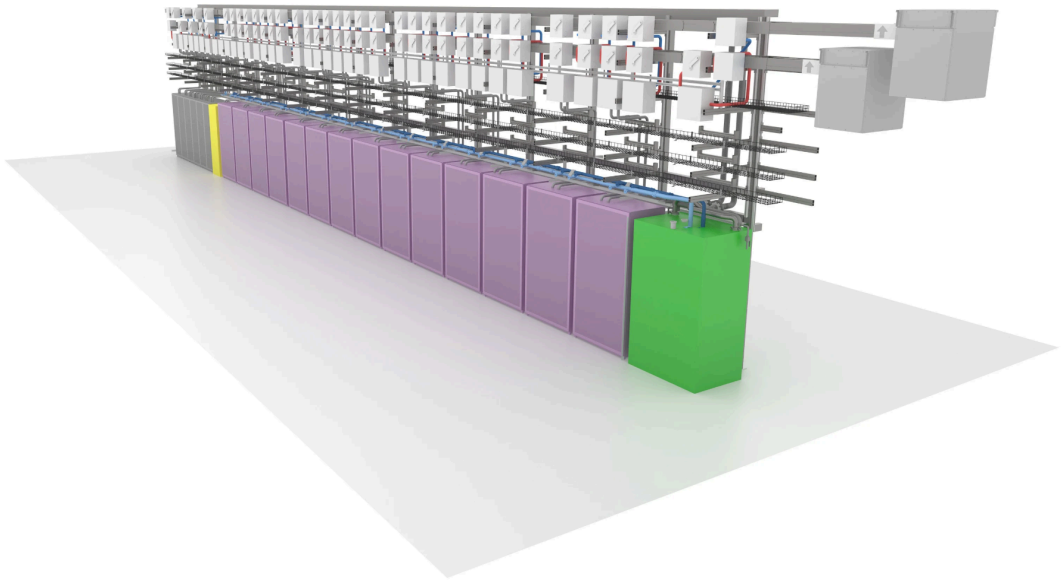


Figure 3: Example 3D renderings of the “Large” scale DC designs, showing the capability to deploy dense ML deployments at scale.



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

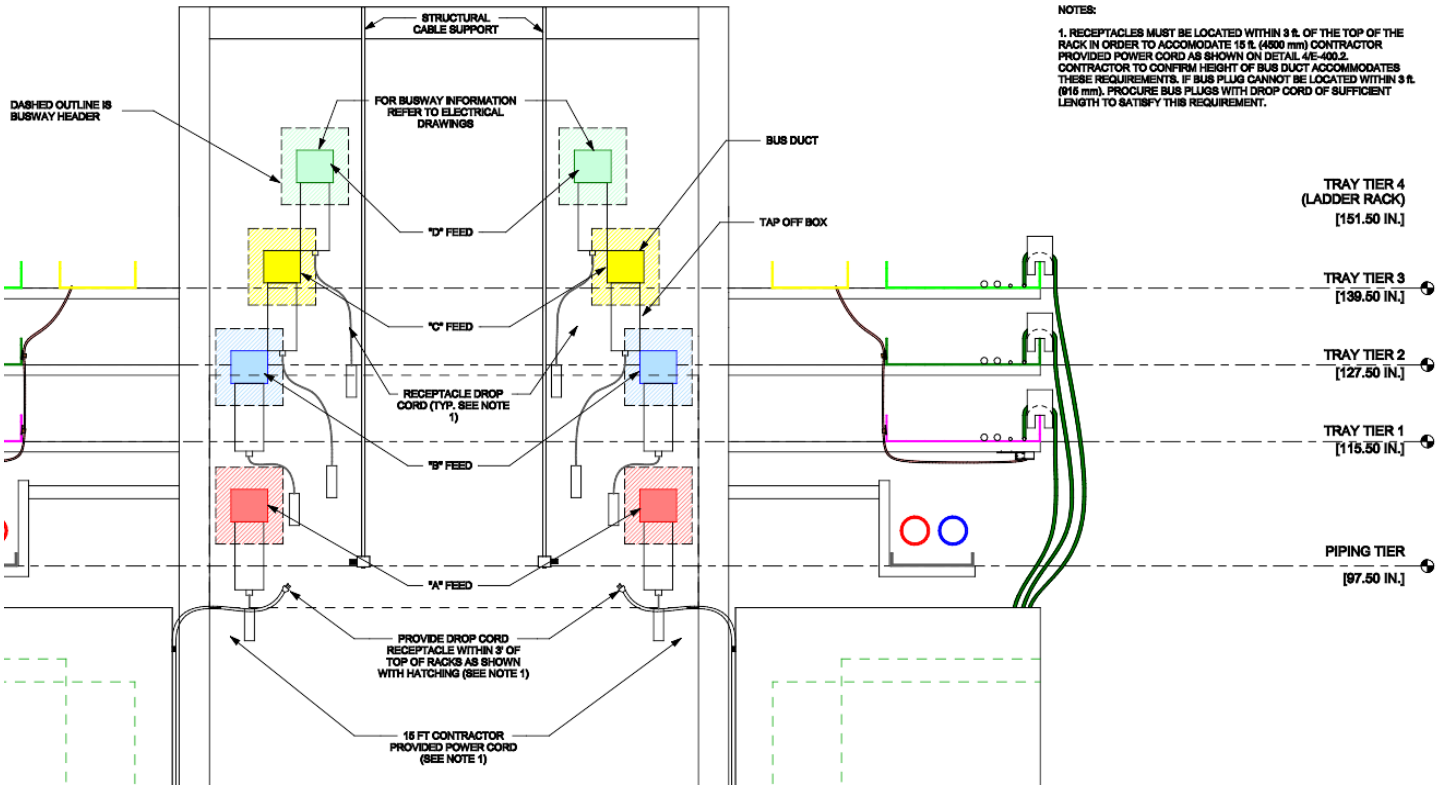


Figure 5: Example HAC cross section for the "Small" Datacenters for 2026 and early 2027 DCs

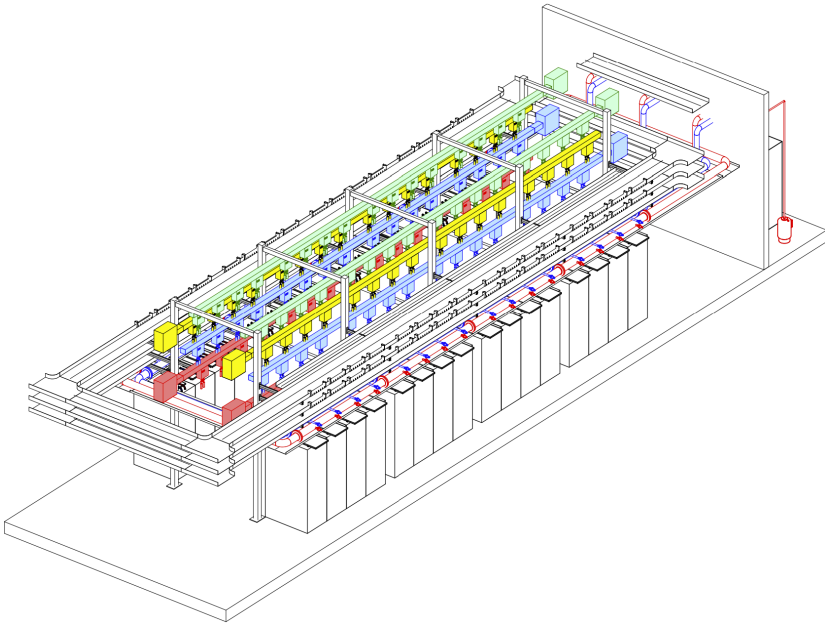


Figure 6: Example HAC cross section for the "Small" Datacenters for 2026 and early 2027 DCs

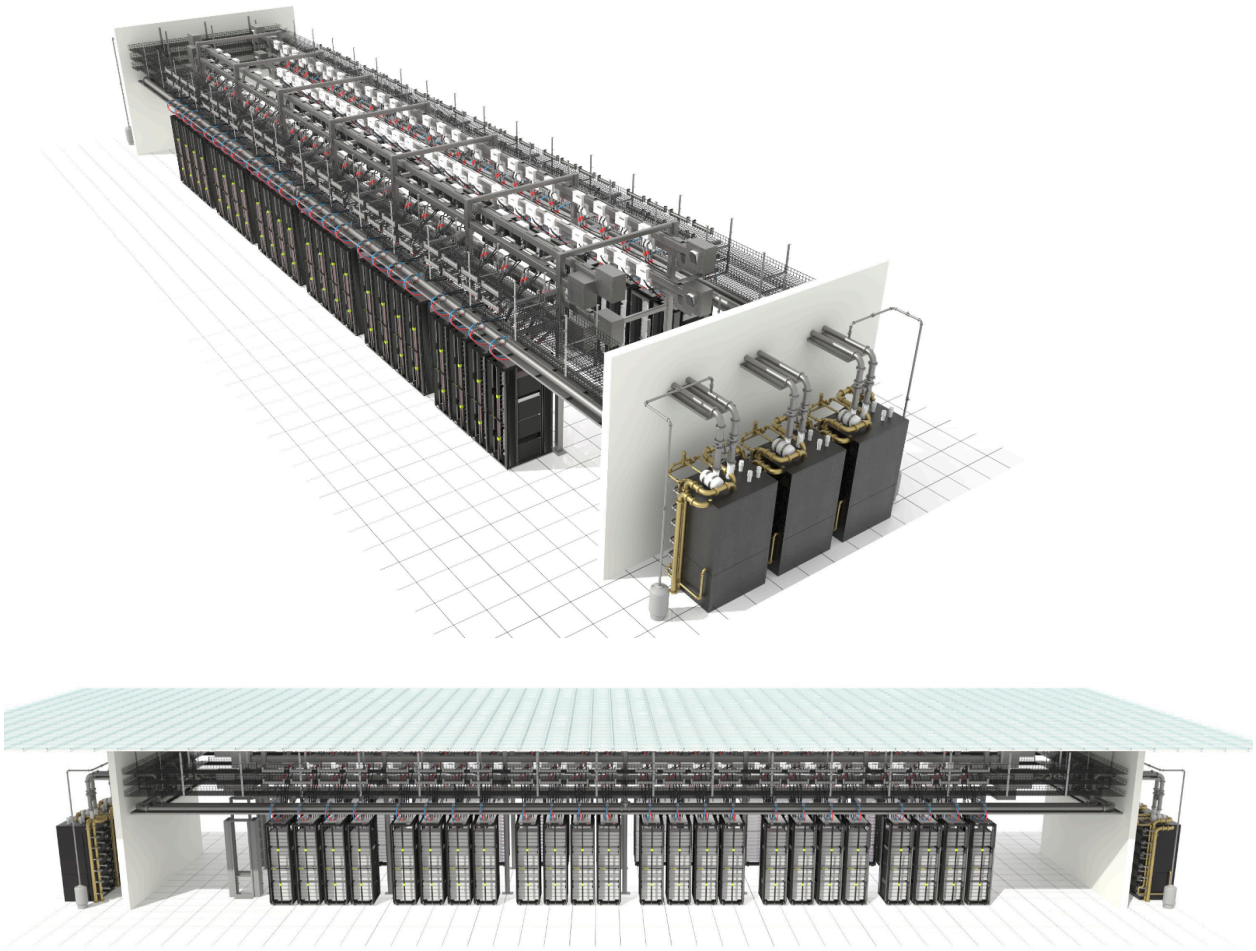


Figure 7:Example 3D renderings of the “Small” scale DC designs, showing the capability to enable high dense ML deployments in constrained spaces.

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	Preferred 96'

Transport aisle requirement	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
Linear Row density	For 2028 TPU deployment: Preferred (to avoid power stranding): < 26.6 kw/ft if CDU/UPS in row; < 29.1 kw/ft if CDU/UPS not in row
Row level power	2.4MW min per row
Clear Ceiling height	18ft (for 2N busway)
Min # of rows	26 rows for ML pods + Additional ~18 rows for NW, C&S and other cluster services.[1]
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 Data center ready - Q4, 2027

[1]If NW, C&S are in a separate room, <100m to the main room, no need to account for the full 18 rows in the DC hall.

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 specification 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.

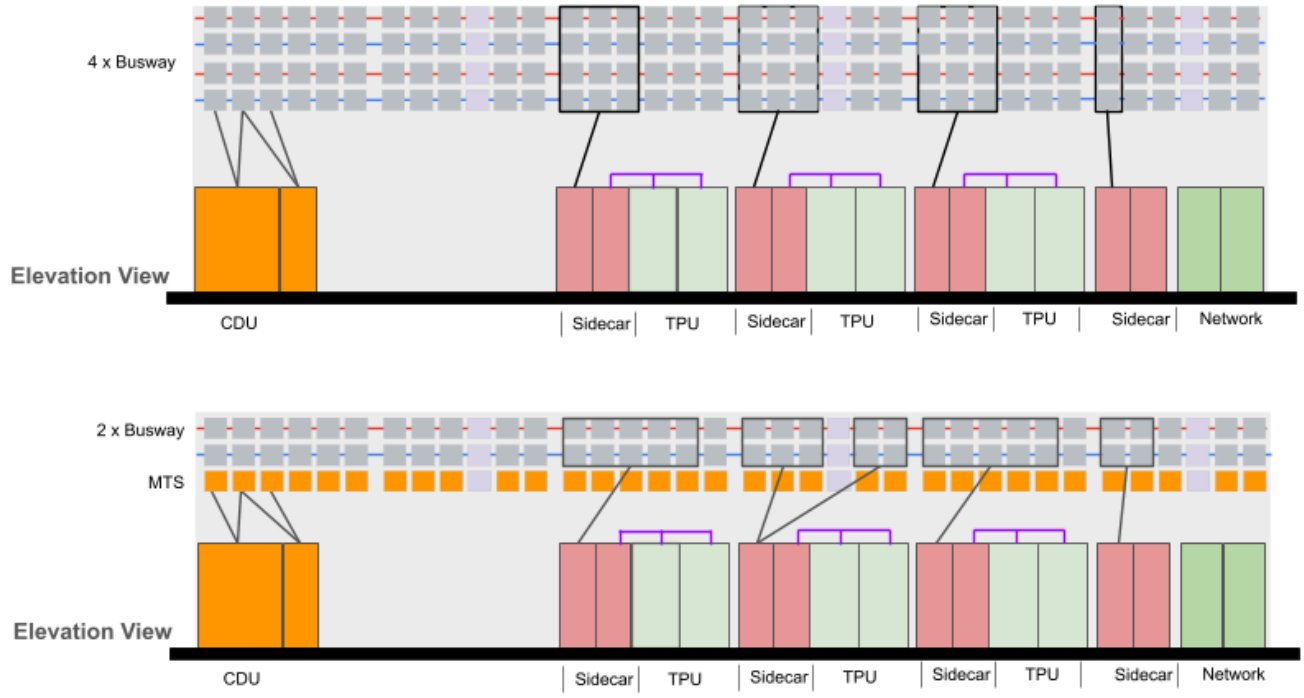
Power = Peak IT
Rack = 28” Ext. Width Rack Spot

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

	V7x (Now)		V8x		Future version (potential up to 2028/29)	
	Power	Rack	Power	Rack	Power	Rack
Starter block (Cluster turn up)	4.30 MW	300	4.40 MW	316	5.90 MW	373
Repeating block per pod	12.10 MW	444	15.69 MW	542	60.43 MW	733
Pod Qty	Power	Rack	Power	Rack	Power	Rack
1 Pod	17 MW	750	21 MW	860	67 MW	1110
ML:CS:NW:Aux	7:1:2:0	3:2:4:1	7:1:2:0	4:2:3:1	8:1:1:0	2:3:3:2
2 Pods	29 MW	1200	36 MW	1400	127 MW	1900
ML:CS:NW:Aux	8:1:1:0	4:2:3:1	8:1:1:0	4:2:2:1	8:1:1:0	2:3:2:3
3 Pods	41 MW	1700	52 MW	2000	188 MW	2600
ML:CS:NW:Aux	8:1:1:0	5:2:2:1	8:1:1:0	5:2:2:1	9:0:1:0	3:3:2:3
4 Pods	53 MW	2100	68 MW	2500	248 MW	3400
ML:CS:NW:Aux	8:1:1:0	5:2:2:1	8:1:1:0	5:2:2:1	9:0:1:0	3:3:2:3
5 Pods	65 MW	2600	83 MW	3100	309 MW	4100
ML:CS:NW:Aux	8:1:1:0	5:2:2:1	8:1:1:0	5:2:2:1	9:0:0:0	3:3:2:3
6 Pods	77 MW	3000	99 MW	3600	369 MW	4800
ML:CS:NW:Aux	8:1:1:0	5:2:2:1	8:1:1:0	5:2:1:2	9:0:0:0	3:3:2:3
7 Pods	89 MW	3500	115 MW	4200	429 MW	5600
ML:CS:NW:Aux	9:1:1:0	5:2:2:1	8:1:1:0	5:2:1:2	9:0:0:0	3:3:1:3
8 Pods	102 MW	3900	130 MW	4700	490 MW	6300
ML:CS:NW:Aux	9:1:0:0	5:2:2:1	9:1:1:0	5:2:1:2	9:0:0:0	3:3:1:3

Notes:

1. % distribution and capacity are baseline estimation, NW% and C&S% could subject to change with IT bandwidth, customer C&S requirement
2. Aux racks includes power, cooling racks like CDU and Sidecar



Example layout of a HAC with the 2028+ TPU specifications

3.2 CDU and Sidecar specification

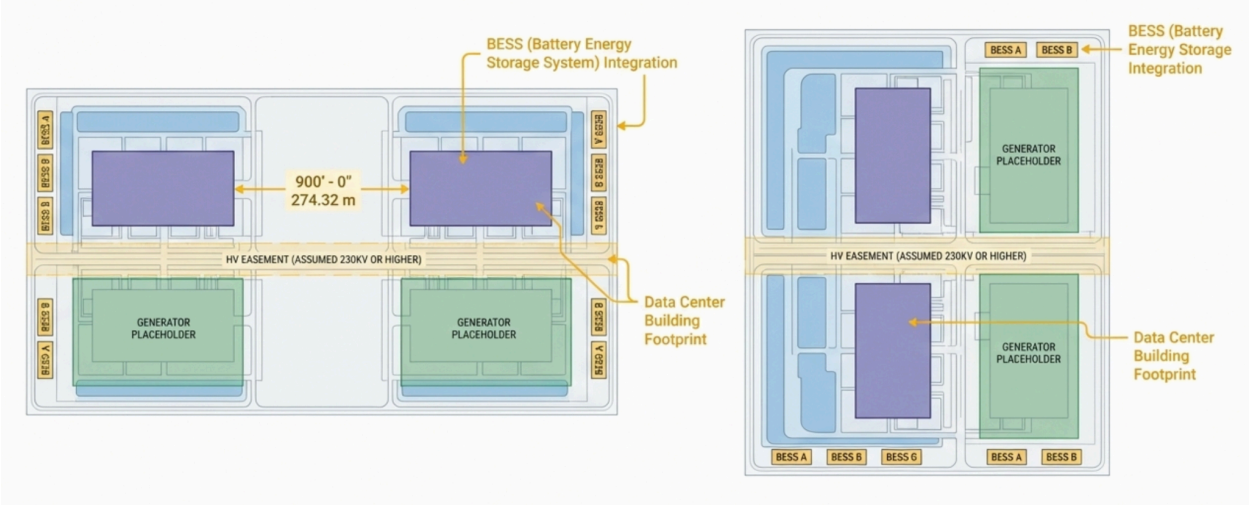
CDU specifications to support liquid cooling are noted in the Project Deschutes specification published in OCP and contain the minimum requirements for a CDU to be capable of supporting liquid cooling.

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.

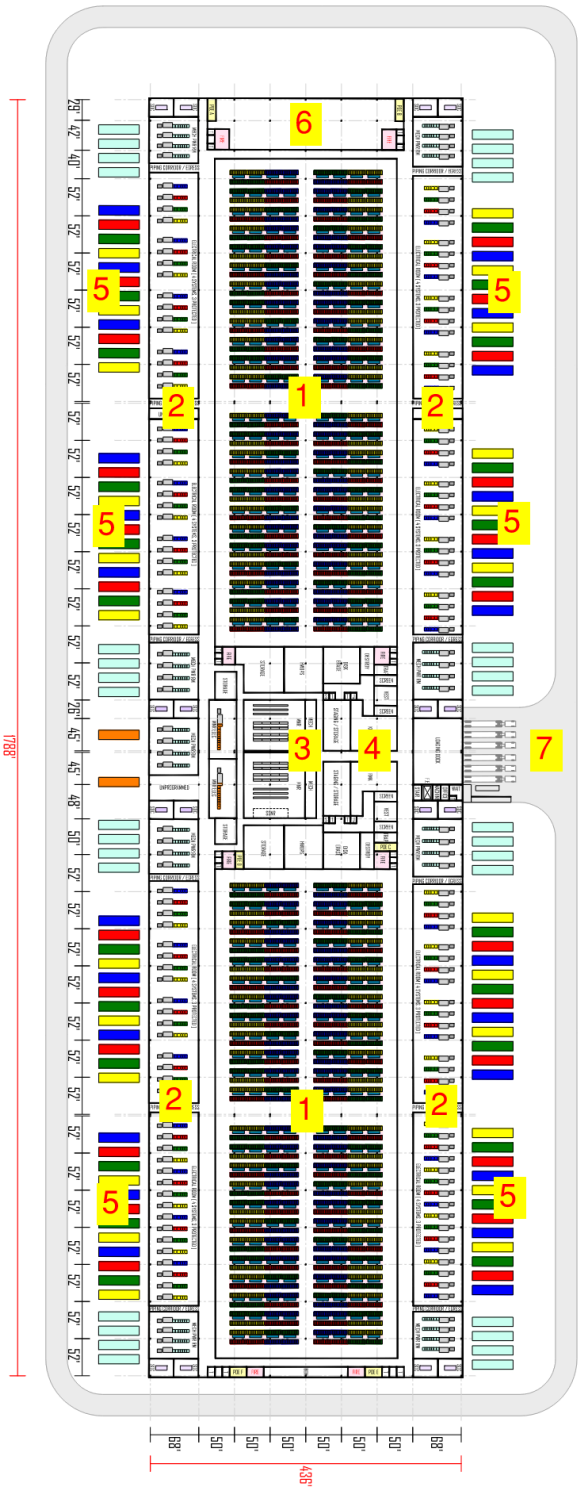


Cluster/Building IT Requirement	Up to ~1.4GW IT load per campus in less than 4 buildings.
Campus Minimal Rack Spot Count	(13k) 28” Ext. Width Rack Spot
Worst case ML chip to chip fiber length	1.5km Minimal building distance will vary based on the actual layout

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



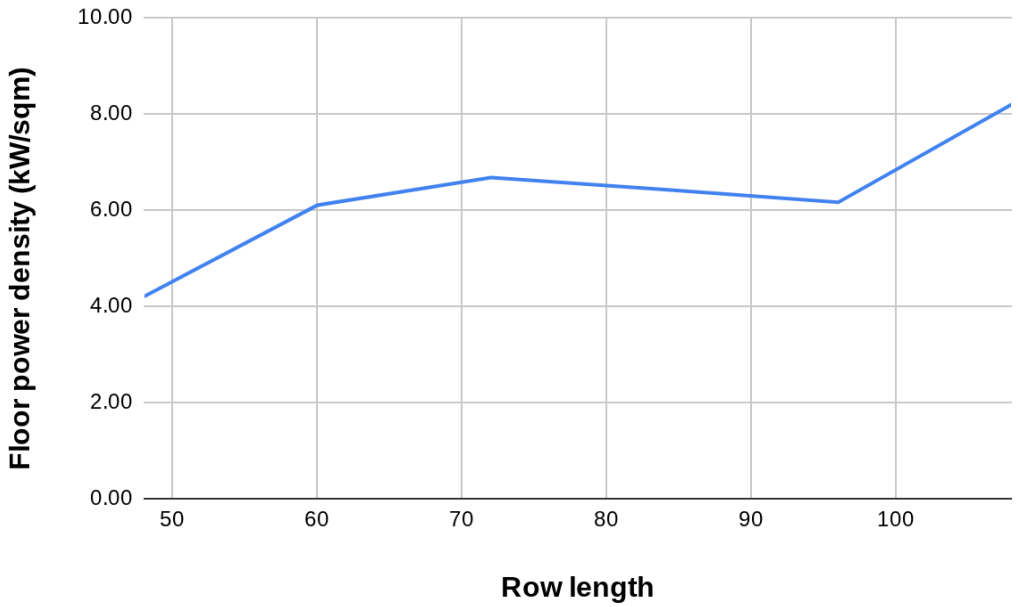
Date: Feb 25, 2026

5. Appendix

5.1 Tradeoff analysis

- **Row length & min # of rows**

  Recommended row length callout - 96’ (row length), and additional callout for “Additional ~18 rows for NW, C&S are required, if NW, C&S are in a separate room, <100m to the main room, no need to account for the full 18 rows in the DC hall, check scaling table on starter block”



Variation of the floor power density (kw/Sqm) with different row length (ft)

Shorter row length introduces inefficiencies in bin packing and requires additional rows for the same amount of power. This adversely affects the floor power density. While shorter rows allow for greater flexibility in available spaces in the market, beyond a certain point, they are not really suited for ML training pod deployment since the amount of space required for deployment increases significantly.

Shorter rows also have the downside of stranding CDU capacity since the same number of racks cannot be deployed in the row. To that end, a 96’ row length enables an optimized row deployment.

The minimum number of rows is recommended to be 44 to account for the rows required for starting fabric and C&S deployments. Not all of these rows scale with the ML chip

count. In order to ensure that the scaling factor is accounted for, we separate the ML rows (which scale with pod count) and the NW rows (which are fixed for the initial cluster) to better help with planning. The NW rows can be located in a separate datahall (since these racks also have a larger air cooling load requirement) as long as they are within 100m of the datahall where the ML racks are deployed.

- **Row power**

  Preferred - 2.4- minimum row power

  Row power increase helps in cases where row lengths are able to support a larger complement of ML training racks per row. This has a dependency on the ability to support liquid cooling as well. As a result, there is an inherent dependency between row power and row length. Hence it is recommended to add in a linear row density callout with some margins.

	CDU/UPS in row	CDU/UPS not in row
2028 TPU blended density	36.7 kw/ft	40.4 kw/ft
2027 TPU blended density (for reference)	16.14 kw/ft	19.10 kw/ft
Recommended design callout for 2028	**26.6 kw/ft**	**29.1 kw/ft**
Recommended design callout for 2027	**11.62 kw/ft**	**13.75 kw/ft**

  For 2028 TPU, Based on 26.4 kw/ft and 2.4MW minimum power per row it leads to 91-92 ft min row length, which is in line with the 96ft row length recommendation. This ensures sufficient “white space” is provided in the row to allow for a bias towards power consumption.

- **Aisle pitch**

  Preferred - 23’/278”
  (Cold Aisle + Hot Aisle + 2 * Rack Depth = 82” + 54” + 2* 71” = 278”)

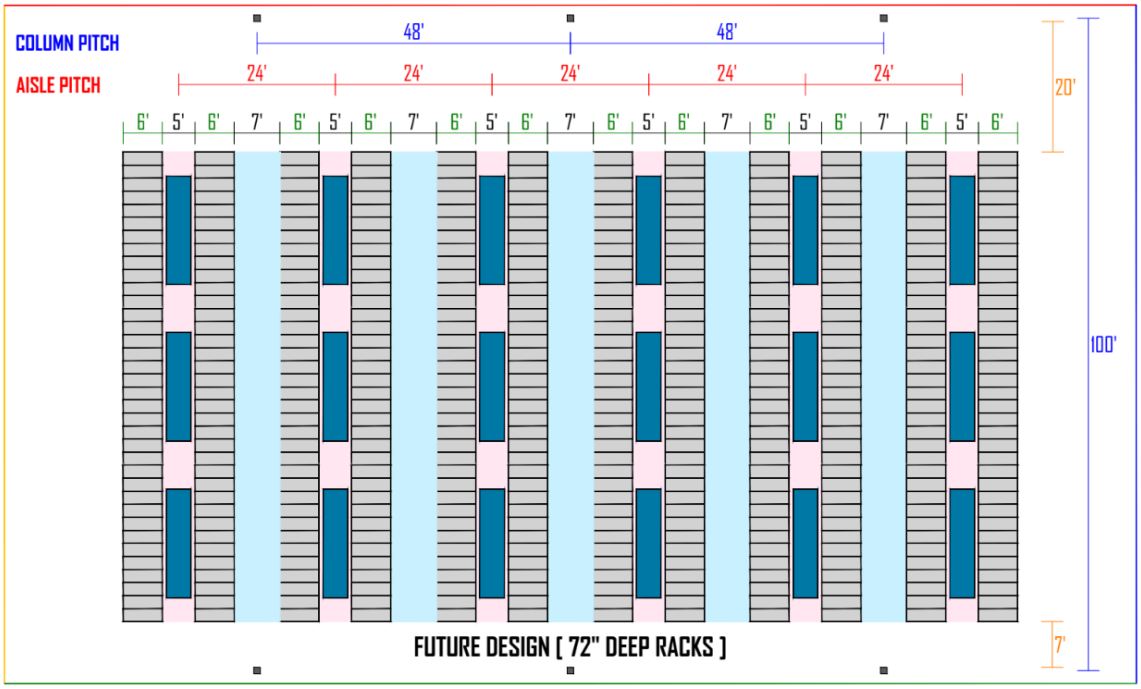
  An aisle pitch of 278” is preferred to be able to provide flexibility for GPU and TPU platform deployments as a minimum for columnless design. This allows for a rack depth of 71” and a cold aisle width of 82” to be able to deploy deeper platforms with ease.

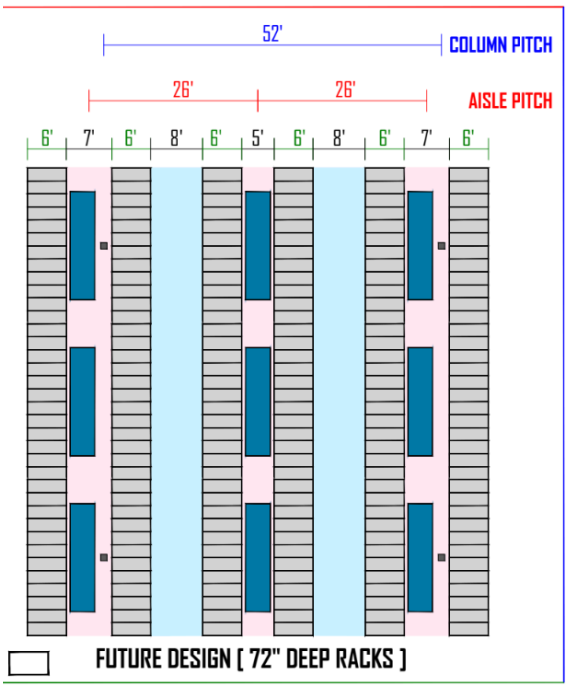
  This aisle pitch is the minimum required for performance and is optimal for fan wall design. See below for the tradeoff with in-row cooling.

- **Contiguous rows/Column positions**

  No explicit callout.

  Column-grid often drives the aisle pitch for the datahall with the columns located inside the HACs. The current aisle pitch of 23ft works well for fan wall design but for datahalls that have in-row cooling, a larger aisle pitch of 26ft might be required in order to account for the column grid structure within the datahall. While we don't specifically call out a 26' aisle pitch in the doc, the following diagrams are provided here for reference (24ft aisle pitch below still includes in row cooling units with 5' hot aisle instead of 4.5' which leads to the increase from 23ft to 24ft).





5.2 Design Process

Regular Cloud Compute & Storage deployments in a datacenter take Order of (years) to fill up since the goal is to optimize the cost of deployment over time. Given the lower fill rate, rack power densities (13-16kW/rack) and the system deployment being limited to the enclosure itself, the design criteria is relatively straightforward.

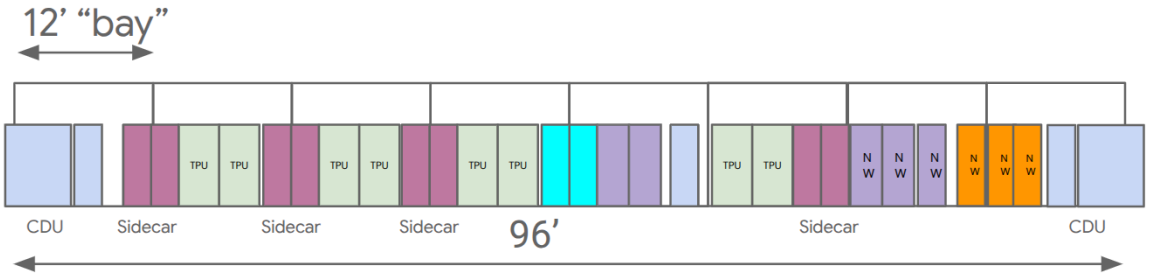
In contrast for ML systems, the entire system (which consists of several interconnected racks) needs to be deployed on day 1. This in addition to the higher rack power densities (>100kW/rack) requires a different way to layout the datacenter design.

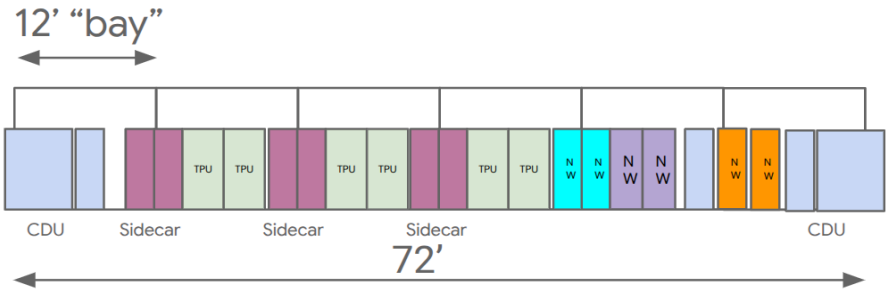
The datacenter design for an ML system requires an understanding of the following

a. The number of interconnected ML racks or pods required
b. The various power levels for the ML platforms
c. The networking equipment required to start the basic cluster services as well as to interconnect the ML racks
d. The Auxiliary Compute & Storage machines required for successfully operating and delivering ML capacity
e. Repair spaces required within the datahall and space required for transit and access
f. Availability target and its impact on backup generators
g. BESS requirements from utility for voltage ride through which impact space on campus based on the load

Based on the above, an iterative approach can be used to determine the datacenter design

1. A specific row length can be defined, in increments of “bays”. Where the notion of a “bay” exists only to account for the vertical posts in the serverhall structure and is often aligned to the building column-grid. E.g., a single bay can be 10’ wide and a row can be 4 or 8 bays.
    a. This is done mainly to the room dimensions, accounting for the transit aisle spaces at the end of the room.
    b. A few options can be considered here. Single row vs multiple row sections etc.
2. Once a few row length options have been determined, then a potential rack layout is done to ensure that the row is fully utilized.
    a. CDU in row vs out of row is one consideration at this stage
    b. Rack counts are impacted by potential sidecar locations and quantity. Example row layouts are shown below for 72’ and 96’ rows.





3. A given rack layout allows for defining the power level required in each row in order to meet the deployment needs.
4. Depending on the number of ML, Network (NW) and Command and Control racks for the overall ML system, the total number of racks and hence the total number of rows required can be calculated.
5. The total power in the datahall can then be balanced against the total number of rows and power/row.

These steps are often iterated upon in order to ensure the optimal design point. Given that the variables here include row length, row power, # of rows and ML rack power itself, there are various optimizations that can be done at all interface levels to achieve the ideal design point.

6. Future Additions

Additional material to be added beyond Version 0.5 into Version 1.0:

- ML Rack liquid cool thermal interface (inlet temperature, gpm, gpm/kw) plot
- Row length and relationship to inference vs training pod deployments relative to floor density requirements
- Power efficiency flexibility around space
- ML Load profile management
- SLO requirement guidance covering different protected states
- PUE optimization guidance
- Operational practices required to support high-density data center environments