From Chip to Grid:
Scaling Holistic AI Data Center Co-Design Through Open Standards

Tom Garvens
VP Data Center Technology & Systems

Amber Huffman
Principal Engineer
OCP Board of Directors

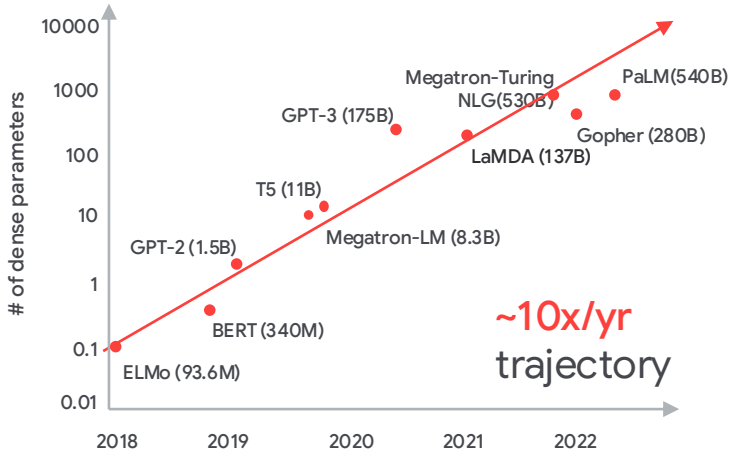
Google Cloud

OCP Global Summit · October 2026

AI Exponential Demands

AI is not only a powerful enabler, but a platform shift.

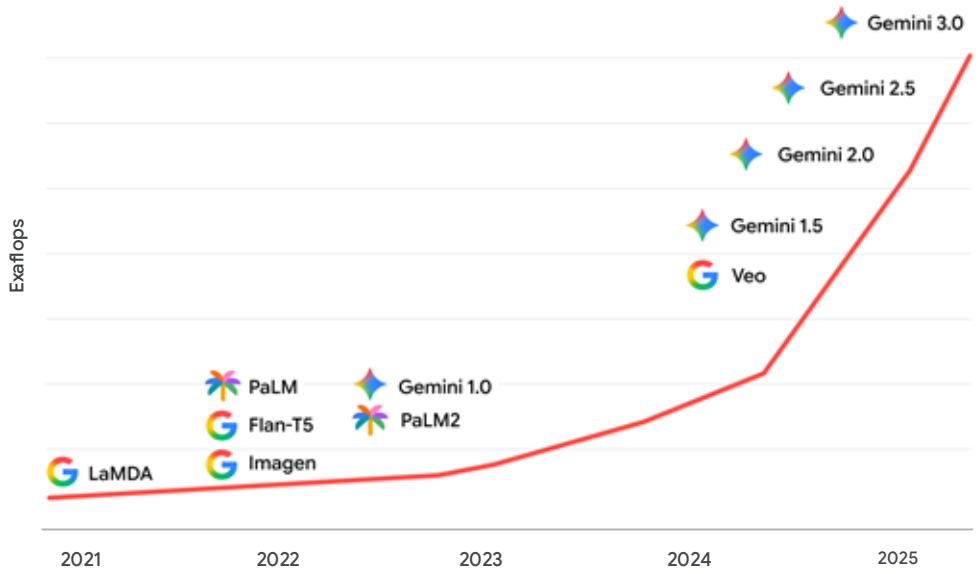
Innovating on compute to power the AI-driven future



Parameters

Data source: "Computation used to train notable intelligence systems", Our World in Data

AI Model Momentum



2026 Milestones

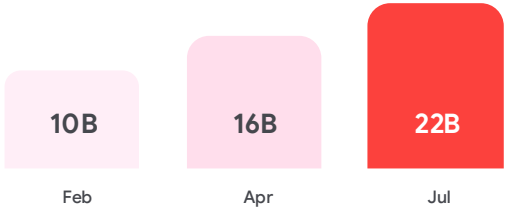
Gemini 4 Argon

Our new frontier model · introduced Sep 30

Shipped Gemini 3.5 | 3.6 | 3.7 | 3.8

>22 billion tokens/min

via direct API use by our customers



(100 tokens roughly equals 75 words)

AI Stack

Services

Models

Gemini

AI Infrastructure Software

PyTorch

AI Infrastructure Hardware





Data Centers




Energy




TPU Supercomputing

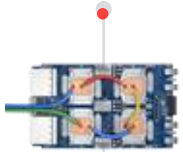
TPU v1
2015

Internal inference accelerator



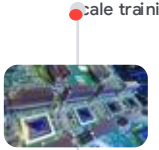
TPU v3
2020

Liquid cooled 1k chips distributed shared memory



TPU v5e
2023

Purpose-built cost-efficiency and performance for medium/large-scale training and inference



Trillium
2024

Exceptional performance & efficiency. Enabling the next frontier of AI models



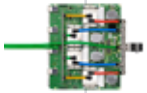
TPU 8t & TPU 8i
2026





TPU v2
2018

Domain-specific AI supercomputing 256 chips distributed shared memory



TPU v4
2022

Optically reconfigurable 3D Torus 4k chips with distributed shared memory



TPU v5p
2023

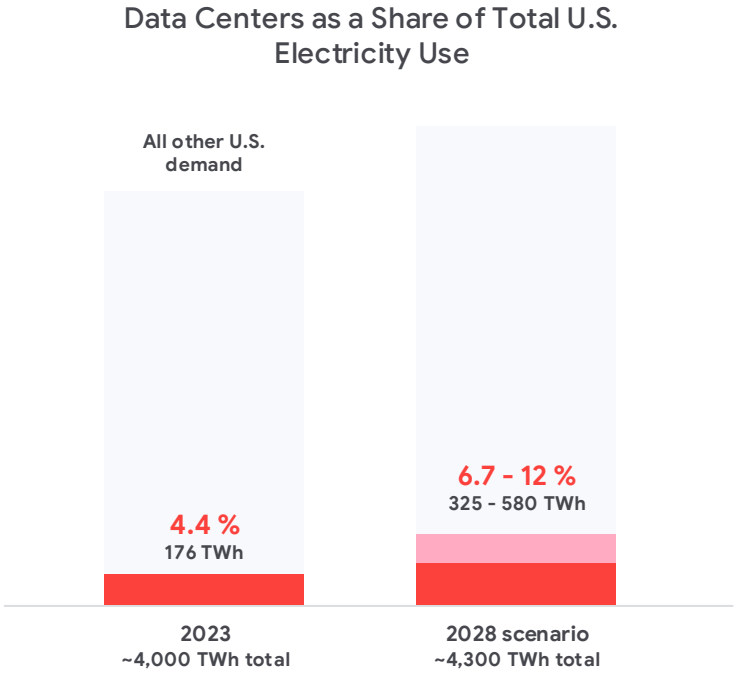
A powerful, scalable, and flexible AI accelerator



Ironwood
2025

Over 9k chips & designed to meet the growing demands of the most complex thinking models

A Small Slice Becoming a Grid Resource



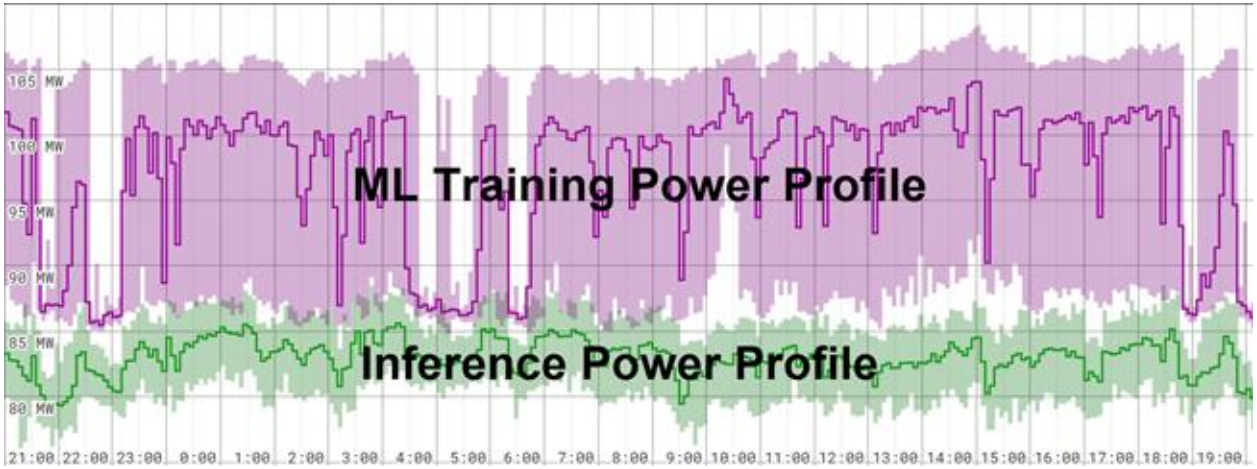
Source: LBNL, 2024 United States Data Center Energy Usage Report (Dec 2024) for usage
Source: EIA, Electricity Explained and Annual Energy Outlook for totals

Impacts due to ML workloads

Ramp Rate

Intra-Minute Fluctuations

Oscillations

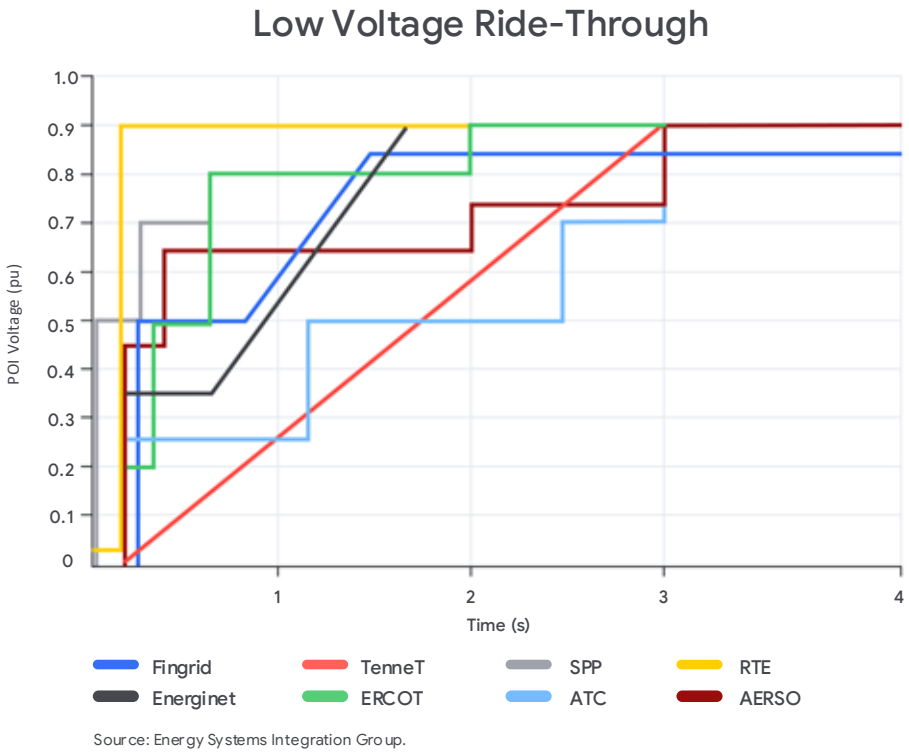


Example 24 hour Load Profile for Training and Inference Workloads

New Utility Requirements

Large load disconnect ride-through and restore requirements

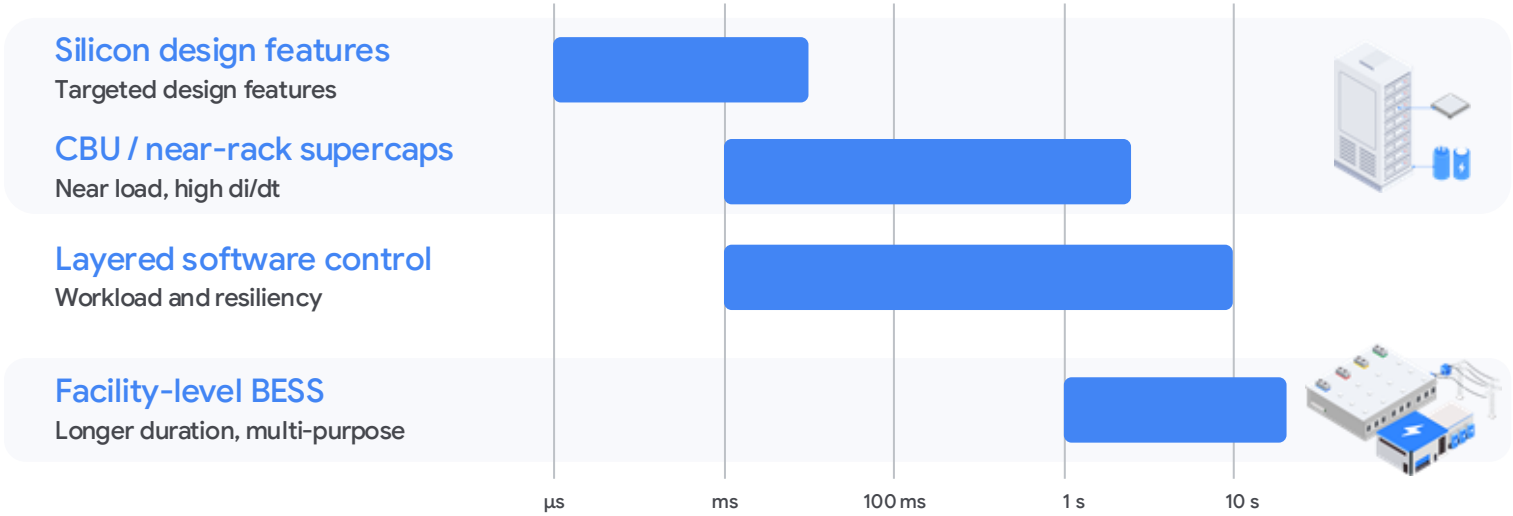
Region / Grid Operator	Required Timeframe
FinGrid (Finland)	90% within 1s
ERCOT (Texas, USA)	90% within 2s
AESO (Alberta, Canada)	100% within 1.0 to 10.0 seconds (configurable)
Energinet (Denmark)	80% within 5 seconds, then 90% within 20 seconds, 95% within 30s



Source: Energy Systems Integration Group.

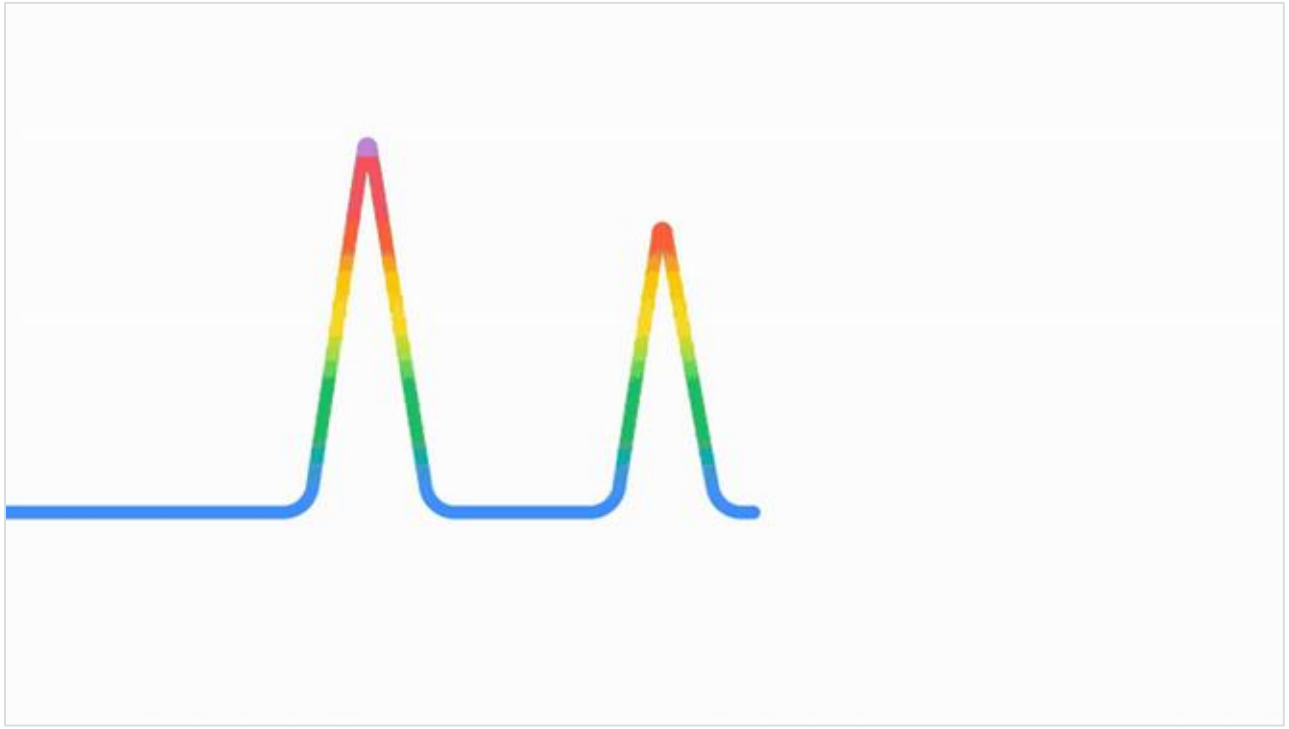
Mitigation Across Every Timescale

No single layer covers the full frequency range of an AI load. Coverage has to be built up by timescale.



Holistic Solutions Required

Mitigation in Action



▶ Mitigation on:
the swing goes flat.
See it at **Booth C1**.

Battery Energy Storage Systems (BESS)

Buffer

Smooths the load the grid sees

Backup

Reduces diesel runtime

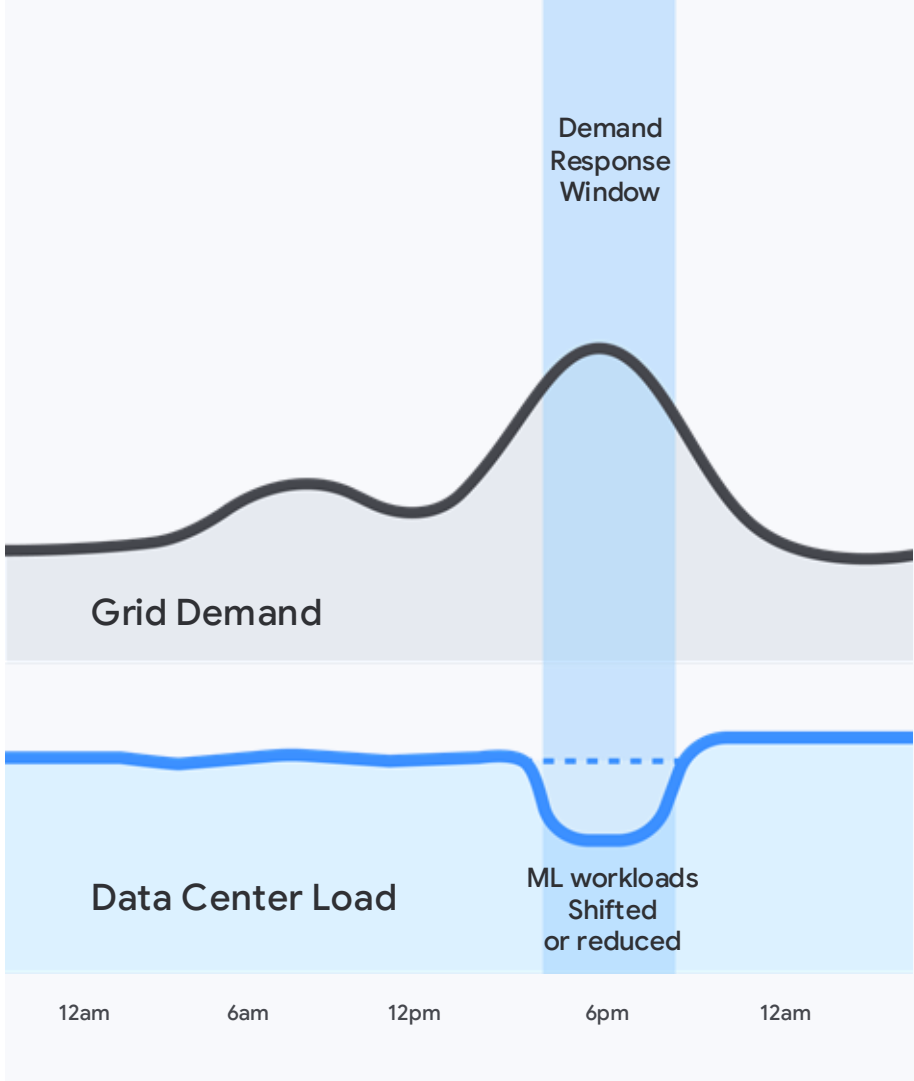
Grid Service Participation

- A Peak Shaving (covering local demand spikes)
- B Demand Response (ex. grid stability during extreme weather)
- C Supplying Power Back (if regulations allow)



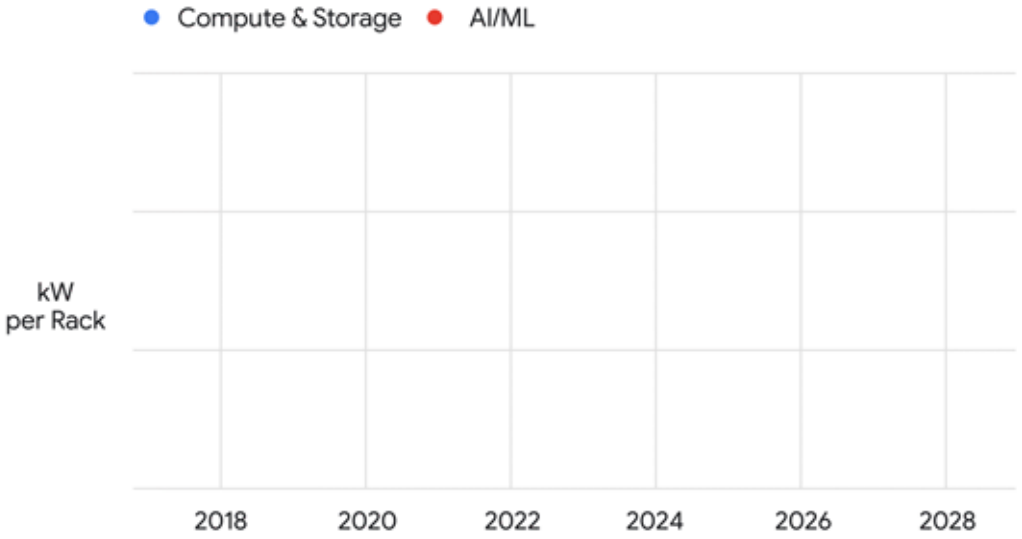
Google has integrated 1 GW of demand response capacity into our contracts with multiple U.S. utilities

On March 19, 2026, we announced a milestone in demand-side flexibility in the United States.
Along with other new resources we're bringing to the system, demand response enables valuable capacity while helping our utility partners support reliability.




Data Centers
For The AI Era

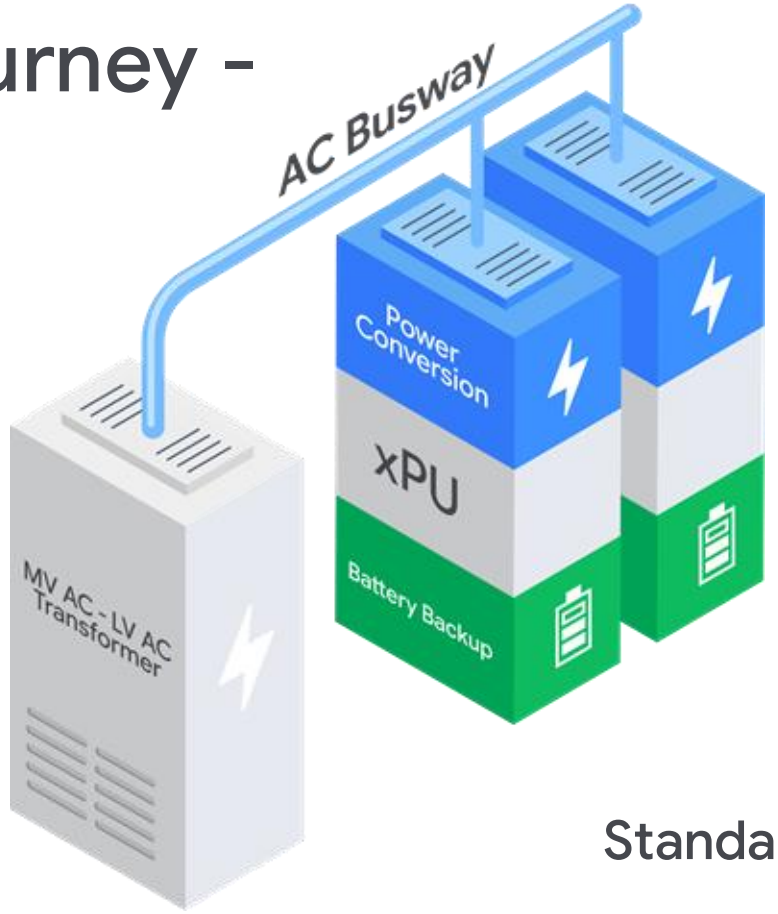
ML Rack Density Challenge





TPU rack

LVDC Journey - Existing

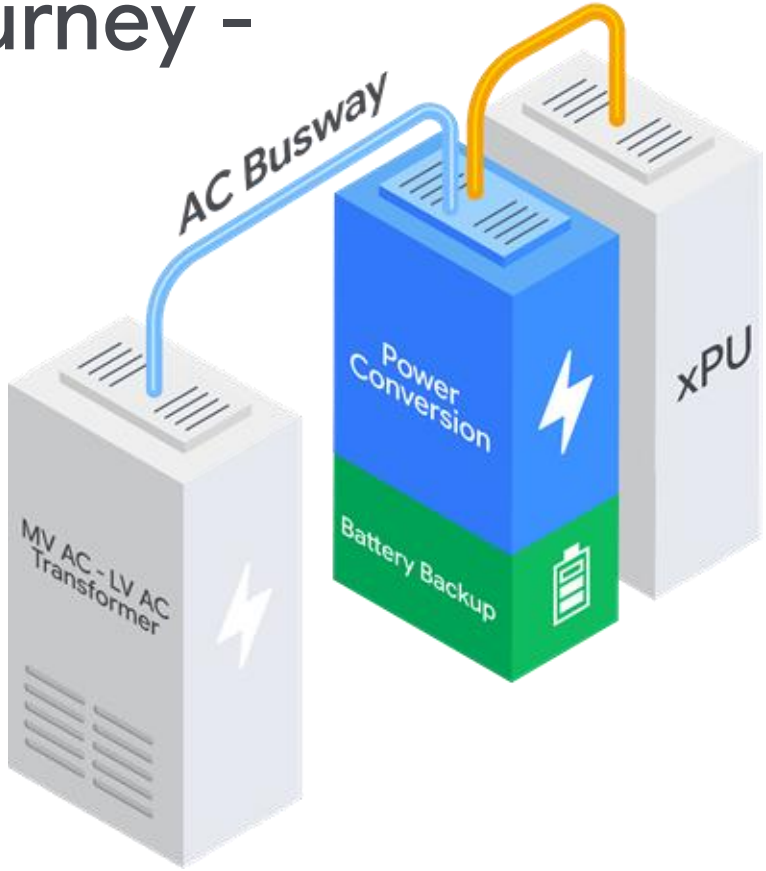


LVDC Sidecar

Standard AC Distribution

LVDC Journey - Today





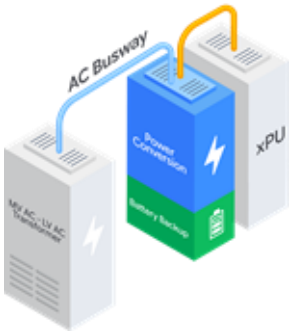
Today
Standard AC Distribution

LVDC Sidecar
Hybrid Transition

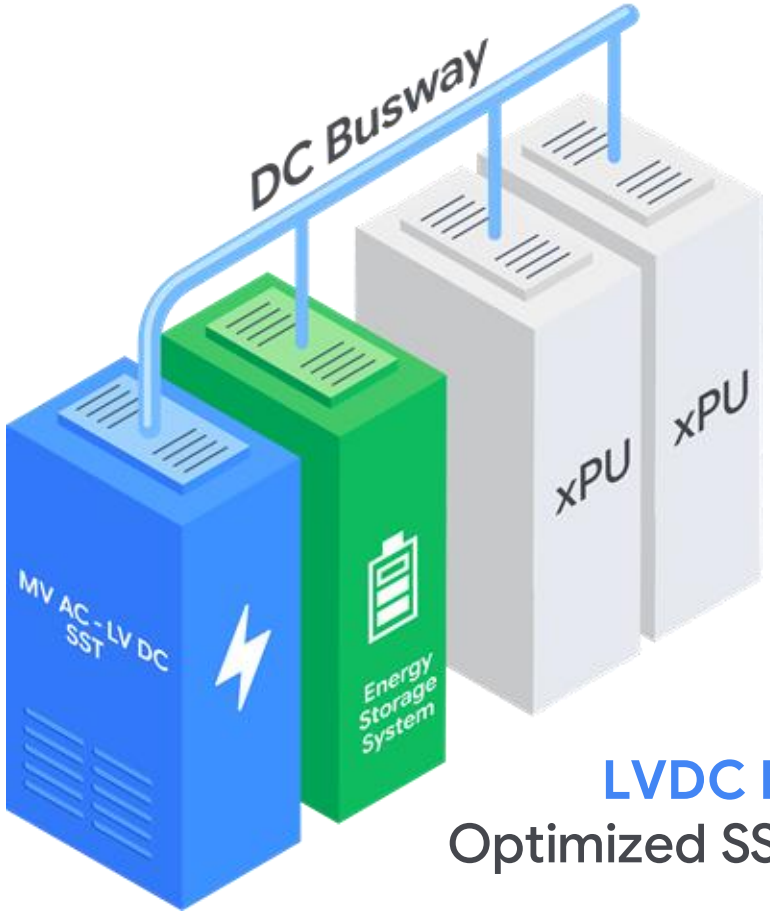
LVDC Journey - Native LVDC



Today
Standard AC Distribution



LVDC Sidecar
Hybrid Transition



LVDC Native
Optimized SST/TRU

LVDC Native Architecture Benefits

20–30%

Floor space densification

Sidecar footprint returns to the data hall

9,000

Install hours eliminated

Fewer field-assembled boxes and connections

>1%

Chip-to-grid efficiency

One conversion stage instead of three

Hundreds of rectifiers and batteries, distributed → Centralized SST and BESS

Consolidating conversion drives both density and velocity.

Figures are internal estimates compared to current AC baseline.


Open Standards
For The AI Era

Why OCP: Align in Months, Not Years

2024
~1,500 MW dropped offline at once

Feb 2026
First draft

Jun 2026
Critical items aligned

Ride through dips
90% load back in 2 s

Sep 2026
v1.0 in effect



Speed
Specs at product speed

Scale
One spec, many vendors

Trust
A neutral venue

800 VDC, From Grid to Chip



1

Mt Diablo

±400 VDC power rack

2

Native LVDC

800 VDC to the row

3

SST

MV AC to 800 VDC · v0.3

4

On-board power

New workgroup, Q4 2026



Liquid Cooling, Open Data Center

Blueprints the whole supply chain can build to

Project Deschutes 1.0

Open 5th-gen CDU spec

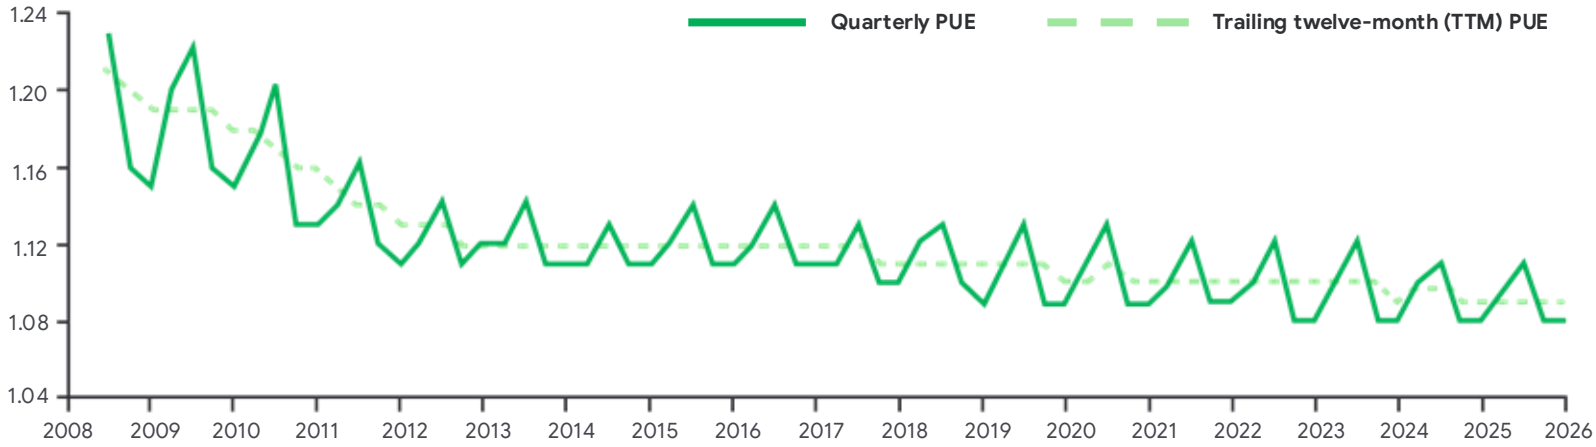
Project Brazos

Liquid Cooling for Air-Cooled DCs

Open Data Center 1.0

7 standard spaces, 12–24 MW blocks

PUE — Power Usage Effectiveness



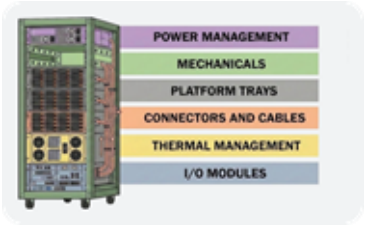
Google strives to build the world's most energy-efficient computing infrastructure, outfitting each Google data center with high-performance servers designed to use as little energy as possible compared to the amount of data they process.

Reducing Embodied Carbon



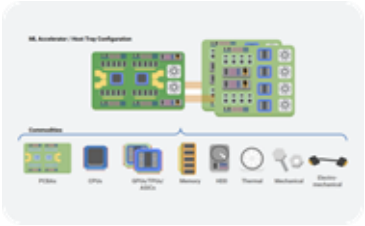
Green concrete

Goal: >50% lower concrete carbon



Carbon disclosure

0–10 data quality score



Product category rules

PCR v0 for ICT hardware



Clean backup power requirements

<10% of diesel CO_2
<10 s to full power

Shared methods let every builder reduce embodied carbon.

Responsible by Design

Energy



12 GW
new clean energy, 2025

35 GW
since 2010

1 GW
demand response

Water



120%
replenishment goal, 2030

7B+
gallons replenished

$500M+
for community water

Sources: datacenters.google/energy · Google 2026 Environmental Report

Go Deeper at OCP This Week

GRID & POWER

Grid Requirements & Mitigation
Wed 4:05 PM · 210CG

Hardware Mitigations: Power Fluctuation
Thu 1:00 PM · 210CG

800 VDC

On-Board Power
Wed 8:00 AM · 210CG

LVDC Ecosystem
Wed 9:10 AM · 210CG

800V Power Safety
Wed 12:55 PM · 210CG

SST Integration
Wed 3:15 PM · 210CG

COOLING & OPEN DC

CDU Sub-Project
Tue 2:25 PM · 220C

Single-Phase Cold Plate
Thu 10:45 AM · 220C

Open DC Architecture
Thu 12:30 PM · 210CG

SUSTAINABILITY

Sustainability Project Updates
Thu 8:00 AM · LL21DE

High-Temperature Fuel Cells
Thu 2:45 PM · 210CG

Open Enablers
Thu 12:50 PM · LL21DE

PCR for LCA
Thu 2:45 PM · LL21DE

Carbon Disclosure
Thu 3:35 PM · LL21DE

AT BOOTH C1 · ML power-fluctuation demo · Mt Diablo 800 VDC power sidecar · Ironwood rack



Collaborate With Us at Product Speed

Visit us at Booth C1

Next to the Central Park Lounge
Tue 11 AM–7 PM · Wed 9 AM–5 PM · Thu 8 AM–2 PM

IN THE BOOTH

- ML power-fluctuation mitigation demo
- Mt Diablo 800 VDC power sidecar
- TPU Ironwood rack
- Ask a Google engineer: office hours

Scan: Google @ OCP
sessions & office hours