

PART 5 • SECOND EDITION • SEPTEMBER 2026 (FIRST PUBLISHED JULY 2026)

Cooling & Thermal Management

Every watt that enters a chip must leave through a declared thermal path. The named rack profile, heat split, interfaces, climate and service objectives select the cooling envelope.

THE COOLING-SELECTION ENVELOPE

30–40 kW typical RDHx; >50 kW with active fans

There is no universal rack-kW cliff: the named equipment and facility envelope decide.

Check liquid heat fraction, airflow and inlet limits, water availability, climate, residual room heat, serviceability, redundancy, and the future density tail.

SEMIANALYSIS, DATACENTER ANATOMY PART 2 – COOLING SYSTEMS, 13 FEBRUARY 2025 · AS-OF 2025-02-13

THE REFRAME

You choose a complete cooling-service envelope, not a technology from a rack-kW lookup table.

Rack profile, heat split, airflow, water systems, climate, rejection, redundancy, serviceability, structure, and the roadmap all move together.

ACT I


The selection envelope

Seven years of accelerator roadmap took the rack from 15 kW toward a 600 kW announced design. Feasible cooling paths overlap, then narrow with the named rack profile.

THE TERRITORY **Cooling paths overlap; product heat split and facility limits decide**

Cooling selection is an overlap, not a cliff

Airflow, heat split, climate, plant, service, warranty and roadmap move every boundary.



Schematic · Source: SemiAnalysis (nVent) + NVIDIA · as of August 2026

Treat published kW bands as examples. Prove the selected path against today's rack and the two-generation tail before fixing water, structure, power, and rejection.

THE DRIVER

2.3 kW TDP support per Rubin GPU (Pegatron)

Vera Rubin GPU power (Pegatron): 2.3 kW TDP support per GPU — an OEM rating, not measured draw

Pegatron lists the rack at 188 kW Max Q / 228 kW Max P but no per-GPU Max Q; Rubin Ultra's package TDP is unpublished. Size the loop from the named profile you will run, never from a roadmap TDP.

PEGATRON RA4803-72N3 PRODUCT PAGE AND DATASHEET · AS-OF 2026-10

GENERATION BY GENERATION

Each accelerator generation narrows the feasible cooling set

GENERATION	RACK POWER (STATE THE BASIS)	NAMED-PROFILE DESIGN BASIS
A100 / HGX (2020–22)	Air-cooled HGX racks — no named product figure	Air: raised floor + containment
H100 / HGX (2023)	Air or RDHx: 30–40 kW typical RDHx capture, >50 kW with active door fans	Verify the actual server airflow/inlet envelope; air or RDHx may fit
NVIDIA GB200 NVL72 by HPE (2024–25)	132 kW nominal TDP	HPE profile: 115 kW liquid + 17 kW residual air
Vera Rubin NVL72 (2026)	188 kW Max Q / 228 kW Max P (Pegatron); 330 kW NVIDIA DSX facility design basis	DLC on the 800 VDC power path
Rubin Ultra Kyber (NVIDIA H2 2027; SemiAnalysis 2028)	~600 kW roadmap per-rack planning figure, boundary unpublished	DLC; in-chip microfluidics on the roadmap

THE PHYSICS

~3,500×


Volumetric heat capacity of water vs air — liquid isn't a better fan, it's a different regime

That factor is why a hose the width of your thumb replaces a hurricane of hot air.

THERMODYNAMIC REFERENCE • AS-OF 2025

AIR'S LAST STAND Containment in section: squeezing the last racks out of air**Containment keeps hot exhaust out of the cold supply**

Same contained hot aisle; a raised-floor CRAH (left) or a slab fan wall (right) delivers the supply.



Either way, fan energy and recirculation set the limit, not a fixed rack-kW wall.

Schematic · Source: Guide synthesis · as of October 2026

Air, RDHx/AALC, and DLC overlap. Qualify heat split/flux, airflow/inlet, water, residual room duty, climate, service, redundancy, and refresh; then compare delivered-watt economics.

THE PRICE OF GUESSING LOW

~\$5–6M/MW (20–40 kW/rack repositioning estimate)

Existing-hall repositioning to 20–40 kW racks: ~\$5–6M/MW (Enabled Energy, July 2025)

A liquid-cooled rack-scale conversion needs a separate facility scope and budget.

JOE STAIB, ENABLED ENERGY, SPONSORED DCD ARTICLE, 8 JULY 2025 · AS-OF 2025-07-08

ACT II

The liquid toolkit

A bridge for the building you have, a product-supported path for the equipment you select, and a cautionary tale about the technology that looked best on paper.

THE FIRST DECISION

Bridge the hall you have, or plumb the hall you need?

Rear-door / air-assisted (the bridge)

- Can extend a brownfield hall when the named rack and coil/airflow envelope close
- ~40–55% lower capex than full DLC in the bridge band
- Door fails silent: servers keep running while the room heats
- You hit the ceiling again in one GPU generation

Direct-to-chip liquid (the destination)

- Carries 120 kW today, 600 kW on the roadmap
- Needs CDUs, a reinforced slab and, unless a sidecar rejects to room air, facility water
- Loss of flow trips a 1 kW GPU in seconds — unmissable
- The design basis when the named equipment profile requires source capture


SO WHAT

Bridging is correct when the building already exists and the selected rack fits the verified RDHx or air-assisted envelope. In a greenfield, compare that lifecycle path with DLC against the declared roadmap.

THE BRIDGE, UP CLOSE A rear door makes the rack room-neutral — a liquid coil where all the air already goes

A rear door cools the exhaust where all the air already goes

No cold plates, no server changes: the rack becomes room-neutral at its own back door.



30–40 kW typical RDHx; not a universal limit, not a rack-kW lookup.

Keep the coil water above the room's dew point.

Below-dew-point water adds capacity but condenses; tempered water keeps the door dry at less capacity.

Schematic with sourced ratings · Source: SemiAnalysis + Motivair + Vertiv · as of September 2026

The door borrows the servers' own airflow, so the room never knows the rack is dense. The tax: coil water must stay above dew point, and passive doors run out of pressure long before active ones.

THE MAINSTREAM LIQUID PATH

>55%

Cold-plate cooling's forecast share of the liquid-cooling market (PMR; no single-/two-phase split)

PMR does not split the figure by phase. Two-phase is not thermally worse; it is better, but its fluid supply chain broke.

PERSISTENCE MARKET RESEARCH, DATA CENTER LIQUID COOLING MARKET • AS-OF 2026

THE STALL

31 Mar 2025

The last day to order 3M Novec — the fluid two-phase cooling was built on

3M announced its PFAS exit in Dec 2022 (completed end-2025), before its separate US public-water settlement (\$10.5–12.5B nominal, 2024–2036). Two-phase stalled as AI demand spiked.

FRASER TECHNOLOGIES SUPPLIER NOTICE • AS-OF 2025-03

THREE ARCHITECTURES


Why good-enough PUE beat best-in-class PUE

	DIRECT-TO-CHIP	1-PHASE IMMERSION	2-PHASE IMMERSION
Supplier-quoted PUE (boundaries unmatched)	~1.05–1.15 design illustration	1.02–1.03 (GRC graphic)	1.01–1.02 (GRC graphic)
Serviceability	Per-tray, hot-swap-ish	Full lift-out; repair time set by the procedure	Worst: lift-out + vapor
2026 status	Mainstream default	Viable niche	Stalled on PFAS

THE CAUTIONARY TALE Two tanks, two fates: single-phase circulates; two-phase boiled into a supply dead end

Immersion in section — single-phase circulates; two-phase boils

Same heat-transfer idea, different service model, fluid market, and regulatory exposure.



Both tanks share the unglamorous costs: floor point-loads, fluid logistics, and hoist-based service.

Immersion earns its costs through heat reuse, environmental protection or component heat-flux headroom.


Schematic · Source: 3M + Vertiv · as of October 2026

Thermal elegance didn't win. Single-phase survives where its serviceability tax pencils out; two-phase — the best PUE on paper — stalled when 3M walked away from PFAS.

INSIDE THE LOOP Cold plate to facility water: the budgets that make DLC work

Liquid and residual air each need their own heat path

Illustrative split heat path · fluid circuits remain separate



Both paths hold in a fault

Accept the cooling path

Cold plate passes; air fails

Add air duty or reduce load

Schematic · Source: Guide derivation · as of Sep 2026

Flow, delta-T, and pressure are one budget spent across the whole chain — and ~10–20% of each rack's heat still leaves on air.

THE DESIGN NUMBERS

GB200 NVL72 loop figures — two OEM profiles, one derived rule

~1.44 L/min/kW

The OCP flow rule per kW of chip heat — the number the silicon hands your pumps

GUIDE HEAT-BALANCE DERIVATION USING WATER PROPERTIES AT THE DECLARED DESIGN POINT · AS-OF 2026

45 °C maximum liquid inlet; 65 °C maximum liquid return (separate limits)

QCT GB200 NVL72 QoolRack temperature maxima — separate limits, not a setpoint

QCT, QOOLRACK STAND-ALONE ADVANCED LIQUID COOLING FOR NVIDIA GB200 NVL72 SYSTEMS · AS-OF 2025-04

132 kW nominal TDP: 115 kW liquid + 17 kW air


HPE GB200 NVL72 heat split at 132 kW nominal rack TDP — the air tail you still catch

HPE, NVIDIA GB200 NVL72 BY HPE DATA SHEET, OCTOBER 2026 EDITION. · AS-OF 2026-10

DRAWING THE MEMBRANE In-rack, row, or central plant — CDU placement sets the blast radius of a pump trip

Where the CDU sits sets how much one isolation takes out

Illustrative case · four CDU placements · one isolation in each



Head check: 240 > 213 kPa at design flow (candidate B)

Selecting the unit does not tell you which racks survive an isolation.

Schematic · Source: Guide derivation · as of Sep 2026

Every placement holds the same clean/dirty membrane; what changes is how much compute one failure takes down and how many units you service. Row-level with N+1 is the mainstream compromise.

ACT III

The facility-water design point

Supported facility-water temperature is one input to compressor hours, water use and heat-reuse grade alongside climate, approaches, rejection method, load and redundancy.

THE MASTER SETPOINT

Chilled water or warm water?

Chilled — W17/W27 (~17–27 °C)

- Generous thermal margin and CDU approach
- More compressor hours where ambient and approach temperatures cannot meet the supported supply point
- Minutes of ride-through stored in the cold-water volume
- Waste heat leaves low-grade — costly to sell

Warm — W40/W45 (~40–45 °C)

- Can expand economizer hours when the climate-bin and approach-temperature calculation closes
- Dry rejection can reduce site water, subject to design-day capacity and fan-energy tradeoffs
- Return temperature follows operating point and load ΔT ; reuse still needs a qualified offtake
- Thermal ride-through must be calculated from actual loop volume, controls and rack response


SO WHAT

ASHRAE W-classes name facility-water supply capability; OEM TCS limits and facility operating points are separate. Overlay climate, approach, load, rejection and redundancy before choosing the plant.


THE LOOP, END TO END **Separate TCS acceptance limits, FWS capability and the operating point**

The facility loop, and the W-class band that sets its economics

Warmer supply can buy economizer hours; climate, approach temperatures, and OEM limits decide.



ASHRAE W-classes — max supply °C and what each band unlocks; all share a 2 °C freeze-protection floor



← legacy chilled regime

more free-cooling hours → W+ above 45 °C

Why warmer water matters: more economizer hours and a more useful return temperature.
Compressor-free hours depend on climate, cooler/HX approaches, and the supply temperatures the OEM accepts.

Schematic · Source: ASHRAE + QCT (Quanta) · as of August 2026

A warmer supported supply can expand economizer hours, but the result follows the actual dry-/wet-bulb bins, heat-exchanger and cooler approaches, load, redundancy and control range.

WHAT A WARMER SUPPORTED POINT CAN BUY

Evaluate W45 capability in context

55–65 °C (illustrative FWS return)

Illustrative return of a 45 °C facility loop: the heat grade a warmer loop gives for reuse

GUIDE CALCULATION · AS-OF 2026-02

~0 L/kWh


Microsoft's design target for water use in closed-loop dry-cooled halls

MICROSOFT, SUSTAINABLE BY DESIGN · AS-OF 2024-12

CASHING OUT THE SETPOINT Heat rejection spends one of three currencies: power, water, or temperature

Heat rejection — what each option spends

Schematic: where each option sits on water and electricity. Your loop temperature decides which you can reach.



Free cooling at one Suwon lab building
Water-side economizer cut chiller, tower and pump kWh:
52.8% modeled (5 winter months), **42.0%** measured (42 days)

Site WUE, all cooling types
US average ~0.36 L/kWh (LBNL estimate)
Microsoft owned fleet ~0.27 L/kWh (FY25)

The fork is **burning kWh vs evaporating water** — and your loop temperature decides how hard the choice is.

Schematic positions · Source: LBNL + Microsoft + NOAA + Kim (2024) · as of October 2026

Chillers burn energy, towers evaporate water, dry coolers need the warm setpoint you did or didn't buy. The climate you sited in decides which currency is cheap — pick the spend to match.

CHIP TO SKY **Four delta-Ts stack between the junction and the weather**


The junction must keep its margin in a fault

Illustrative · a fault raises inlet temperature, resistance or both



$$T_{\text{junction}} = T_{\text{inlet}} + Q \times R_{\text{local}}$$

R_{local} : interface + cold plate to inlet



Schematic · Source: Guide derivation · as of Sep 2026

Every approach temperature you spend narrows the free-cooling window. The budget closes at the worst-case branch — or it doesn't close at all.

THE THERMAL BUDGET

Closing the budget at the worst branch

A Blackwell-class rack on a warm-water loop, from the weather to the silicon:

Outdoor design wet-bulb	~24 °C
+ heat-rejection approach	+4–8 °C
+ CDU heat-exchanger approach	+3–5 °C
+ worst-branch manifold penalty	+1–3 °C
+ coolant rise across the cold plate	+8–10 °C
+ junction-to-coolant resistance	+25–35 °C
Typical margin left before the GPU throttles (~85–90 °C)	~10–15 °C

A representative mid-case — the extremes of these ranges swing the margin from zero to ~25 °C. A 5 °C CDU approach instead of 3 °C costs margin you never get back.

THE UNFORGIVING PART

~99.999%

Fleet-wide CDU availability Google reports since 2020 — redundant pumps and heat exchangers on UPS-backed power

Liquid cooling has seconds of thermal inertia, so the floor is engineered: redundant pumps, N+1 rejection, UPS-backed CDUs — and commissioning must prove the failover, not just the capacity.


GOOGLE CLOUD · AS-OF 2025-04

THE HEAT YOU COULD SELL

Waste heat is a graded product — the upstream loop temperature fixed its grade

Heat reuse: what you have, what buyers need, and the heat-pump toll be

Schematic temperature ladder, warmer water higher up. Check your buyer's actual supply schedule.



Schematic · Source: EnEfG §11 + Columbia Climate Law Blog + IEA DHC + NVIDIA · as of October 2026

Warm-water return feeds modern district heat directly; chilled-water return needs a heat-pump lift that eats the margin. Reuse-readiness is becoming a permit condition, so grade is a siting decision.

THE REGULATORY FLOOR

10% / 15% / 20% by applicable start-date cohort

Germany's mandatory energy-reuse factor for new data centers, stepping up by commissioning year

Heat reuse is becoming a permit condition, not a sustainability slide — and the grade you can sell was fixed by the loop temperature you chose upstream.

GERMAN FEDERAL MINISTRY OF JUSTICE • AS-OF 2023-11-13

THE THROUGH-LINE

The vendor fixes the junction limit and the TDP. Everything you control is one budget — coolant temperature and the resistance stack — spent once, from the weather to the die.

TAKE IT HOME

What to do Monday

- 01 Size the substrate — floor, facility water, CDU space, pipe racks — for the rack two generations out, not the one you're buying.
- 02 Choose the lowest-lifecycle-cost path that closes the named rack, facility, service, and refresh envelope—not a kW band.
- 03 Default to single-phase direct-to-chip; pilot immersion only — and hold two-phase until a PFAS-free fluid supply exists.
- 04 Choose the facility-water temperature first — it sets compressor hours, water use, and heat-reuse grade before any plant buy.
- 05 Buy back ride-through: protected pump power where the study needs it, N+1 rejection, and a loss-of-flow test at commissioning.
- 06 Close the thermal budget at the worst-case branch — a budget that closes on average throttles GPUs in production.

|| THE DEFINITIVE GUIDE TO
AI Data Centers

This deck is the map. The guide is the territory — every number here is dated, sourced, and kept current in the full chapters.

- Ch 5.1 Thermal fundamentals and the density wall
- Ch 5.4 Direct-to-chip liquid cooling — the 2026 default
- Ch 5.7 Facility water loops and warm-water cooling
- Ch 5.11 Reliability, leak detection, and commissioning

READ THE FULL PART, WITH EVERY SOURCE AND DERIVATION, AT [AIDATACENTERGUIDE.COM](https://aidatacenterguide.com)

SOURCES

Every figure in this deck traces to the guide's dated register (1/4)

SemiAnalysis, Datacenter Anatomy Part 2 – Cooling Systems, 13 February 2025 (RDHx section) · as-of 2025-02-13 · <https://newsletter.semianalysis.com/p/datacenter-anatomy-part-2-cooling-systems>

Pegatron RA4803-72N3 product page and datasheet · as-of 2026-10 · <https://svr.pegatroncorp.com/products/pega-pod-rack-level/nvidia-rubin/ra4803-72n3>

HPE, NVIDIA GB200 NVL72 by HPE data sheet, October 2026 edition. · as-of 2026-10 · <https://www.hpe.com/psnow/generateDDS/NVIDIA%20GB200%20NVL72%20by%20HPE%20data%20sheet-PSN1014890104MYEN.pdf>

Pegatron RA4803-72N3 datasheet · as-of 2026-08 · https://svr.pegatroncorp.com/SpecFile/PEGATRON_RA4803-72N3_DS_2026v1.pdf

NVIDIA GTC (Jensen Huang); DCD, Tom's Hardware · as-of H2 2027 roadmap (GTC 2025); 2028 reported July 2026 · <https://www.tomshardware.com/pc-components/gpus/nvidia-shows-off-rubin-ultra-with-600-000-watt-kyber-racks-and-infrastructure-coming-in-2027>

Thermodynamic reference; ASHRAE TC 9.9 · as-of 2025 · https://handbook.ashrae.org/Handbooks/S20/SI/s20_ch13/s20_ch13_si.aspx

Joe Staib, Enabled Energy, sponsored DCD article, 8 July 2025; republished 26 September 2025 · as-of 2025-07-08 · <https://enabledenergy.net/news/ai-legacy-data-centers/>

SOURCES

Every figure in this deck traces to the guide's dated register (2/4)

Persistence Market Research, Data Center Liquid Cooling Market (PMRREP13330, March 2026); repeated by Schneider Electric, 5 May 2026 · as-of 2026 ·

<https://www.persistencemarketresearch.com/market-research/data-center-liquid-cooling-market.asp>

Fraser Technologies supplier notice (final orders ending 31 March 2025); 3M, Operations and innovation (exit completed end-2025) · as-of 2025-03 ·

<https://www.frasertech.co.uk/legislation/3m-alternative-solvents/>

CIBSE Journal, CPD Module 254, Liquid cooling in data centre applications (October 2025); SemiAnalysis, Datacenter Anatomy Part 2 (2025), secondary. A screening range quoted by practitioners, not an Uptime cooling-architecture statistic. · as-of 2025-10 · <https://www.cibsejournal.com/cpd/modules/2025-09-lcdca/>

GRC, Two-Phase vs. Single-Phase Immersion Cooling, PDF p.4 comparison graphic. Undated primary supplier brochure. · as-of undated brochure, verified 2026-10-19 ·

<https://www.grcooling.com/wp-content/uploads/2025/09/grc-blog-library-tech-comparison-%E2%80%94-two-vs-single-phase-immersion-cooling.pdf>

Guide heat-balance derivation using water properties at the declared design point · as-of 2026

QCT, QoolRack Stand-Alone Advanced Liquid Cooling for NVIDIA GB200 NVL72 Systems · as-of 2025-04 · https://blog.qct.io/wp-content/uploads/2025/04/QCT-Qoolrack-Stand-Alone_Advanced-Liquid-Cooling-for-NVIDIA-GB200-NVL72-Systems.pdf

SOURCES

Every figure in this deck traces to the guide's dated register (3/4)

Guide calculation; boundary definitions from ANSI/ASHRAE 127–2020 Addendum b, §§3.1 and 6.4–6.5 · as-of 2026-02 · https://www.ashrae.org/file%20library/technical%20resources/standards%20and%20guidelines/standards%20addenda/127_2020_b_20260227.pdf

Microsoft, Sustainable by design (2024-12-09); Microsoft, Mount Pleasant first facility completion (2026-06-23); Microsoft, Inside the world's most powerful AI datacenter (2025-09-18) · as-of 2024-12 · <https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/>

OCP, Project Deschutes: Data Center Facilities Specification 1.0.0 (Google), effective 2026-02-26 · as-of 2026-10 · <https://www.opencompute.org/documents/ocp-specification-deschutes-v1-0-pdf>

Open Compute Project, OAI System Liquid Cooling Guidelines (March 2023 update), section 1.3 Coolant Flow Rate, printed p. 7 · as-of 2023-03 · <https://www.opencompute.org/documents/oai-system-liquid-cooling-guidelines-in-ocp-template-mar-3-2023-update-pdf>

Google Cloud (Project Deschutes, OCP EMEA) · as-of 2025-04 · <https://cloud.google.com/blog/topics/systems/enabling-1-mw-it-racks-and-liquid-cooling-at-ocp-emea-summit>

SOURCES

Every figure in this deck traces to the guide's dated register (4/4)

German Federal Ministry of Justice / Federal Office of Justice
– EnEfG of November 13, 2023, BGBl. I No. 309; EnEfG §11(2),
German Energy Efficiency Act – controlling cohort percentages
and dates; German Energy Efficiency Act – statutory reuse
obligation; Columbia Climate Law Blog – Kristin Weiß, From EU
Framework to National Action: How Germany Regulates Data Center
Energy Use, October 24, 2025, commentary on heat reuse and
district-heating integration only. Cohort percentages and dates
follow the official §11(2) text. · as-of 2023-11-13 ·
https://www.gesetze-im-internet.de/enefg/__11.html

SOURCES

About these figures

Figures are resolved by reference from the guide's claim register at build time (Second edition · September 2026 (first published July 2026)) — a number appears here only as it appears, dated and sourced, at aidatacenterguide.com. Newer data may exist: check the live register.