

APPENDICES • SECOND EDITION • SEPTEMBER 2026 (FIRST PUBLISHED JULY 2026)

Appendices & Reference Data

The cheat-sheet deck: the envelopes, multipliers, crosswalks, and failure rates worth pinning to the wall — each one dated, sourced, and re-checkable in the live register.

WHAT THIS DECK IS

**The chapters teach the decisions.
The appendices put document
identities, OEM envelopes, scoped
calculations and failure modes
within reach at the bench.**

*Seven of them: crosswalks (A), reference designs (B), calculators (C),
the numbers register (D), glossary & timeline (E), FMEA catalog (F),
regional deltas (G).*

ACT I

Power and cooling envelopes

Rack power is one input. The named equipment heat split, airflow and inlet limits, water path, residual room heat, and roadmap case decide what closes.

PIN THESE FIRST

Freeze the named equipment profile before selecting the plant

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

The HPE GB200 NVL72 liquid/air heat split—not a universal rack-kW cliff—sets the basis

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

**135 kW nominal TDP /
up to 155 kW peak**

Lenovo GB300 NVL72: 135 kW nominal rack TDP, up to 155 kW peak – the 2026 liquid default

LENOVO PRESS LP2357 PRODUCT GUIDE, UPDATED 2026-08-30. · AS-OF 2026-08

~600 kW

Kyber NVL144 on 800 VDC (NVIDIA: H2 2027; SemiAnalysis: 2028) – per-rack planning figure

NVIDIA GTC · AS-OF H2 2027 ROADMAP (GTC 2025); 2028 REPORTED JULY 2026

THE SELECTION ENVELOPE

Cooling is not a rack-kW lookup

SELECTION INPUT	FREEZE THIS EVIDENCE	WHY EQUAL KW CAN DIVERGE	COMMITMENT
OEM thermal profile	Input, heat flux, supported architecture and inlet limits	Products split the same total kW differently	Modality + acceptance envelope
Liquid / residual air	Worst supported liquid capture and room-air remainder	High liquid capture can leave material room heat	CDU/FWS + room-air plant
Air path	Flow, pressure, containment, fan power, clearances	Hall and rack airflow margins differ	CRAH/in-row/RDHx + layout
Water + climate	TCS/FWS, fluid, design days, water and rejection	Site conditions reshape the same rack	Plant, water budget + controls
Reliability + roadmap	Maintenance/fault outcomes, isolation and future tail	Viable systems carry different outage/retrofit costs	Topology + reserved path

THE MULTIPLIER CHAIN

384 HPE GB200 racks: 50.688 MW nominal rack IT

Illustrative rack subtotal: 48 SUs × 8 HPE GB200 racks. Supporting network, storage and management loads must be added before total-IT admission.

Reference arrangement – 48 integer SUs	48 SUs
NVL72 racks – 48 × 8	384 racks
GPUs – 384 × 72	27,648 GPUs
GPU-rack subtotal – 384 × HPE 132 kW nominal	50.688 MW
Rack-only facility screen – assumed PUE 1.2	~61 MW
Rack-only POI screen – assumed 1.15 margin	~70 MW POI

Illustrative transformer allowance; qualify surviving capacity

2–3 × 75 MVA

~70 MW screens rack IT only. Add supporting IT; qualify surviving plant. Reserve demand for the named future rack/population or phase it separately: a new queue slot

ACT II

Thermal & water

Warm water rewrote the reference card. The numbers that matter now are supply-temperature classes, liters per minute per kilowatt, and megawatts per CDU.

READ THE CLASS NUMBER AS A TEMPERATURE

ASHRAE liquid classes: the W-number is a water temperature

CLASS	MAX FACILITY SUPPLY	HEAT REJECTION	AI-RELEVANCE
W17	≤ 17 °C	Chiller typically required	Conservative DLC; high chip margin, poor PUE/WUE
W27	≤ 27 °C	Economizer in cool climates, chiller assist	Common transitional DLC operating point
W32	≤ 32 °C	Largely chiller-less in many climates	Mainstream warm-water target
W40	≤ 40 °C	Dry coolers; broad free-cooling window	Strong PUE/WUE; district heat viable
W45	≤ 45 °C	Free cooling in nearly all climates	Forward roadmap – NVIDIA's Vera Rubin example runs 45 °C rack coolant

PLUMBING BY HEART

The liquid-loop multipliers

~1.44 L/min/kW

Secondary-loop flow rule of thumb – multiply by rack kW to size every manifold

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

1.3 MW

HPE NVL72 in-row CDU: 1.3 MW, up to eight racks – declare redundancy separately

HPE, NVIDIA GB200 NVL72 BY HPE QUICKSPECS, VERSION 6 · AS-OF 2026-09

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

QCT GB200 NVL72 QoolRack temperature maxima – separate limits, not a pair

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

THE EFFICIENCY ENVELOPES

PUE and WUE bands to quote from memory

1.05-1.15

Liquid-cooled PUE range practitioners quote (CIBSE, 2025): a screen, not a measured band

CIBSE JOURNAL, CPD MODULE 254, LIQUID COOLING IN DATA CENTRE APPLICATIONS · AS-OF 2025-10

~0.36 L/kWh

LBNL modeled US annual-average site WUE: ~0.36 L/kWh in 2023, across all cooling types

LBNL, 2024 UNITED STATES DATA CENTER ENERGY USAGE REPORT, FIGURE 4.7 · AS-OF DECEMBER 2024 REPORT; 2023 ESTIMATE

ACT III

Reliability & failure data

Measured events belong to named populations. Treat catalog couplings as hypotheses, validate the project topology, and assign shared events or beta only from justified evidence.

THE GROUND TRUTH

Measured references from named populations — not universal rates

1 per ~3 hr

Unexpected-interruption rate, Meta's Llama 3 405B run — 419 of 466 in 54 days

META · AS-OF 2024-07

35.3% / 17.2%

Share of interruptions from faulty GPUs / HBM3 — the silicon, not the facility

META AI, THE LLAMA 3 HERD OF MODELS, ARXIV V1, JULY 31, 2024, §3.3.4 · AS-OF 2024-07

~1,500 MW

Load lost in the July 2024 Virginia event — the case behind NERC's Level 3 alert

NERC, INCIDENT REVIEW: LARGE LOAD LOSS · AS-OF 2024-07

TRANSLATE INTENT, THEN VERIFY THE STANDARD TEXT

Resilience families: compare scope, not matching numerals

STANDARD FAMILY	CLASSIFICATION	PRIMARY SCOPE	DO NOT INFER
Uptime Institute	Tier I–IV	Topology and operational outcomes	An availability percentage or equivalence to another family's same numeral
ANSI/TIA-942-C	Rated-1–4	Telecom, architectural, electrical and mechanical facility requirements	Automatic Uptime or EN/ISO certification
EN 50600 / ISO/IEC 22237	Availability Class 1–4	Facility infrastructure, availability and management framework	Automatic Uptime or TIA certification
Project service objective	Named maintenance and fault cases	End-to-end path, transfer, loading, common modes and workload recovery	A guaranteed outcome from a topology name alone

THE REDUNDANCY ANTI-PATTERN

~25–40% facility

Whole-facility planning premium for Tier IV fault tolerance over Tier III concurrent maintainability

The bodies certify different scopes. Reprice the actual design at equal usable load, then select by named maintenance/fault outcomes, transfer, common modes and workload recovery.

SAVILLS RESEARCH, EUROPEAN DATA CENTRES, MAY 2024 · AS-OF 2024-05

ACT IV

Borders, money & provenance

The same design crosses a frequency boundary or a code family and stops being compliant. Annualized capital cost changes with earning life; a book-policy change affects reported expense.

THE DELTA MAP

Frequency and code basis are separate jurisdiction fields

JURISDICTION	FREQUENCY / COMMON SERVICE	INSTALLATION BASIS
United States	60 Hz · 208Y/120 or 480Y/277 V examples	NEC edition adopted by state/local AHJ
Canada	60 Hz · utility/province-specific service	CSA C22.1 adopted provincially/territorially
Mexico	60 Hz · NOM/utility-defined voltage	NOM-001-SEDE + national/local/utility rules
EU / United Kingdom	50 Hz · 230/400 V commonly	EU: national HD 60364 basis · UK: BS 7671 + duties/DNO
Japan / Gulf exceptions	Japan splits 50/60 Hz · Saudi 60 Hz · other GCC 50 Hz	Separate national codes, voltages and utility rules

DIFFERENT RULES, DIFFERENT BOUNDARIES

Regional requirements constrain the tested operating point

**≤ 1.25 large/very-large;
 ≤ 1.2 at national hub
nodes**

China: new/expanded large facilities versus national-hub projects

NDRC ET AL., DATA CENTER GREEN AND LOW-CARBON DEVELOPMENT SPECIAL ACTION PLAN, JULY 2024, §II(2). · AS-OF 2024-07

≥ 500 kW

EU EED threshold for mandatory PUE/WUE/ERF/REF reporting, annual by May 15

DIRECTIVE · AS-OF 2026-08-24

26 °C+

SS 697: 26 °C+ air; IMDA cites potential 2–5% cooling-energy saving per 1 °C increase

IMDA, TROPICAL DATA CENTRE STANDARD · AS-OF 2023

THE MACRO ANCHORS

Market vitals — the three biggest numbers in the guide

~\$6.7T

Global data-center capex required by 2030, midline of three scenarios

MCKINSEY, 'THE COST OF COMPUTE' · AS-OF 2025-04

~950 TWh

Global data-center electricity demand by 2030 — about 3% of global electricity

IEA, KEY QUESTIONS ON ENERGY AND AI · AS-OF 2026-04

>2,060 GW

Generation and storage waiting in US interconnection queues at end-2025

LBNL, QUEUED UP: 2026 EDITION · AS-OF END-2025

HANDLE WITH THE CAVEAT ATTACHED

Keep the service and date on each economic example

~\$1.90/M tok

Guide self-hosted Llama-70B scenario;
separate from Ramp's March-2025 customer
cohort

SPHERON · AS-OF 2026-04

~\$2.8/hr

Neocloud H100 rental index, Oct 2026 –
volatile; re-pull before committing

SILICON DATA H100 RENTAL PRICE INDEX ·
AS-OF 2026-10-19

\$3.0/GPU-hour

Guide assumption: idle hours remain paid; use
Chapter 1.8's matched cash-flow case

GUIDE ASSUMPTION, 2026-09-08. EPOCH
SUPPLIES THE COST METHOD, NOT THESE
PRICES. · AS-OF 2026-09-08

FIXED-INPUT SENSITIVITY

~70%

Three earning years cost about 70% more annually than seven

Epoch's May 2026 inputs give about \$12B versus \$7.3B/year; the exact five-year case is \$8.8B/year. Earning life is disputed; the arithmetic is reproducible. Book policy only moves reported expense.

EPOCH AI, AMELIA MICHAEL AND BEN COTTIER, MAY 2026 WORKBOOK • AS-OF 2026-05

THE PROVENANCE DISCIPLINE

A number with no date, no source, and no caveat is not a fact — it is a rumor with a decimal point.

THE INDEX CARD

Which appendix, when

- 01 A – match a Uptime Tier requirement, ASHRAE class or KPI acronym to its document, edition and applicability before contracting.
- 02 B – check 384 HPE GB200 racks and 101,376-GPU BOM: reconcile supporting IT, counted fabric, failure-state plant and quote scope.
- 03 C – the live calculators: TCO and \$/GPU-hr, \$/M-tokens, PUE/WUE, and the cooling-selection envelope.
- 04 D – re-pull any volatile figure before committing capital; every number is date-stamped and sourced.
- 05 E – bench definitions, a Chapter 2.1 schedule delta and links to discipline programs and primary evidence.
- 06 F – physical FMEA IDs, consequences and tests; Part 11 initiators link only through validated shared top events.
- 07 G – the deltas that break a home-market design abroad: frequency, codes, water law, climate.

|| THE DEFINITIVE GUIDE TO
AI Data Centers

Pin the envelopes. Check source vintage, actual verification and next review separately before capital depends on the number.

- Ch B The multipliers behind every sizing sanity-check
- Ch C Live calculators for TCO, tokens, and the cooling-selection envelope
- Ch D The date-stamped register – re-pull before you commit
- Ch F Physical FMEA IDs, qualified tests and separate security initiators

READ THE FULL PART, WITH EVERY SOURCE AND DERIVATION, AT AIDATACENTERGUIDE.COM

SOURCES

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

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>

Lenovo Press LP2357 product guide, updated 2026-08-30. · as-of 2026-08 · <https://lenovopress.lenovo.com/lp2357-lenovo-nvidia-gb300-nvl72-rack-scale-ai>

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>

NVIDIA DGX SuperPOD GB200 Reference Architecture · as-of 2025-11 · <https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-gb200/latest/index.html>

ASHRAE, Emergence and Expansion of Liquid Cooling in Mainstream Data Centers (2021), p. 4: fifth-edition class changes. · as-of 2021 ·

https://www.ashrae.org/file%20library/technical%20resources/bookstore/emergence-and-expansion-of-liquid-cooling-in-mainstream-data-centers_wp.pdf

NVIDIA Technical Blog, 'NVIDIA Vera Rubin POD: Seven Chips, Five Rack-Scale Systems, One AI Supercomputer' (2026-03-16) · as-of 2026-03-16 · <https://developer.nvidia.com/blog/nvidia-vera-rubin-pod-seven-chips-five-rack-scale-systems-one-ai-supercomputer/>

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

SOURCES

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

HPE, NVIDIA GB200 NVL72 by HPE QuickSpecs, Version 6 (September 8, 2026) · as-of 2026-09 · <https://www.hpe.com/us/en/collaterals/collateral.a50009224enw.html>

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

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/>

LBNL, 2024 United States Data Center Energy Usage Report, Figure 4.7 · as-of December 2024 report; 2023 estimate · https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf

Meta (Llama 3 Herd of Models) · as-of 2024-07 · <https://arxiv.org/abs/2407.21783>

Meta AI, The Llama 3 Herd of Models, arXiv v1, July 31, 2024, §3.3.4 / Table 5. · as-of 2024-07 · <https://arxiv.org/html/2407.21783v1>

NERC, Incident Review: Large Load Loss (2025-01-08), pp.7-8; NERC comments to FERC, Docket EL25-49 (2025-04-23), Attachment A (April 17 presentation), PDF p.17 · as-of 2024-07 · https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf

SOURCES

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

TIA, Revision C announcement, 9 May 2024; Rated terminology predates this revision. · as-of 2024-05-09 · <https://tiaonline.org/press-release/tia-unveils-revision-c-of-the-global-data-center-infrastructure-standard/>

Savills Research, European Data Centres, May 2024; publisher report. Quoted secondary estimate; not measured project evidence. · as-of 2024-05 · <https://pdf.euro.savills.co.uk/european/european-commercial-markets/spotlight-eu-data-centre---may-2024---final.pdf>

IEC 60038 / IEC 60364; ANSI C84.1 · as-of standard · <https://webstore.iec.ch/en/publication/153>

OCCTO / utility data · as-of standard · https://www.occto.or.jp/en/assets/en/information_disclosure/report_on_the_quality_of_electricity_supply_data_for_fy2024/2024_qualityofelectricity_260113.pdf

NDRC et al., Data Center Green and Low-Carbon Development Special Action Plan, July 2024, §II(2). · as-of 2024-07 · <https://www.ndrc.gov.cn/xxgk/zcfb/tz/202407/P020240723625582947550.pdf>

Directive (EU) 2023/1791 Art. 12; Delegated Regulation (EU) 2024/1364 Arts. 3 and 5 · as-of 2026-08-24 · https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj

IMDA, Tropical Data Centre Standard (SS 697:2023) public summary. Year precision encoded as January 1. · as-of 2023 · <https://www.imda.gov.sg/how-we-can-help/green-dc-roadmap/tropical-dc-standard>

SOURCES

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

McKinsey, 'The cost of compute' · as-of 2025-04 ·
<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>

IEA, Key Questions on Energy and AI (Apr 2026) · as-of 2026-04 ·
<https://www.iea.org/reports/key-questions-on-energy-and-ai>

LBNL, Queued Up: 2026 Edition (data through end-2025) · as-of end-2025 ·
https://eta-publications.lbl.gov/sites/default/files/2026-06/queued_up_2026_edition.pdf

Spheron (vendor blog, 2026-04-04) — assumptions stated; H100 median from GetDeploying (2026-09-29) · as-of 2026-04 ·
<https://www.spheron.network/blog/ai-inference-cost-economics-2026/>

Silicon Data H100 Rental Price Index (SDH100RT), 2026-10-19 reading (\$2.84, +7.2% over seven days); earlier reading \$2.81 on 2026-10-05 · as-of 2026-10-19 ·
<https://www.silicondata.com/products/silicon-index/h100>

Guide assumption, 2026-09-08. Epoch supplies the cost method, not these prices. · as-of 2026-09-08 ·
<https://epoch.ai/data-insights/ai-datacenter-cost-breakdown>

Epoch AI, Amelia Michael and Ben Cottier, May 2026 workbook; Guide fixed-life calculation. · as-of 2026-05 ·
<https://epoch.ai/data-insights/ai-datacenter-cost-breakdown>

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.