PART 1 · SECOND EDITION · SEPTEMBER 2026 (FIRST PUBLISHED JULY 2026)

1

# Strategy, Workload Archetypes & Economics

*Before the slab, the substation, or the silicon, there is one decision: what machine are you actually building? Part 1 is the strategy layer — archetypes, procurement, and the four numbers that decide the return.*

THE DEFINITIVE GUIDE TO
AI Data *Centers*

THE CHECK A GIGAWATT WRITES

# ~$38B

## The all-in program a single gigawatt campus commits you to — capex, plus what it burns every year

*One upstream choice — the workload you intend to run — decides whether that check earns its cost of capital or strands on the balance sheet.*

EPOCH AI, AMELIA MICHAEL AND BEN COTTIER, MAY 2026 MODEL AND LINKED WORKBOOK. MAY 19 IS THE ARTICLE'S UP-FRONT CAPEX CSV REVISION DATE · AS-OF 2026-05

THE REFRAME

# Workload leads the requirements brief, but it does not determine the facility by itself.

*Equipment profile, service objectives, site constraints, regulation, schedule and commercial terms join workload in the design basis. Part 1 is about closing that vector before steel is cut.*

ACT I

# Name the machine you're building

*Five workload archetypes and the wider requirements vector that constrains each design. Validate costly-to-change commitments — siting, structure, utility and distribution — against the full basis.*

THE TERRITORY

# One requirements vector, seven coupled decisions

## The workload sets the cascade; the hardware sizes each step

Walk it forward from the workload; a step that fails sends you back up the chain.

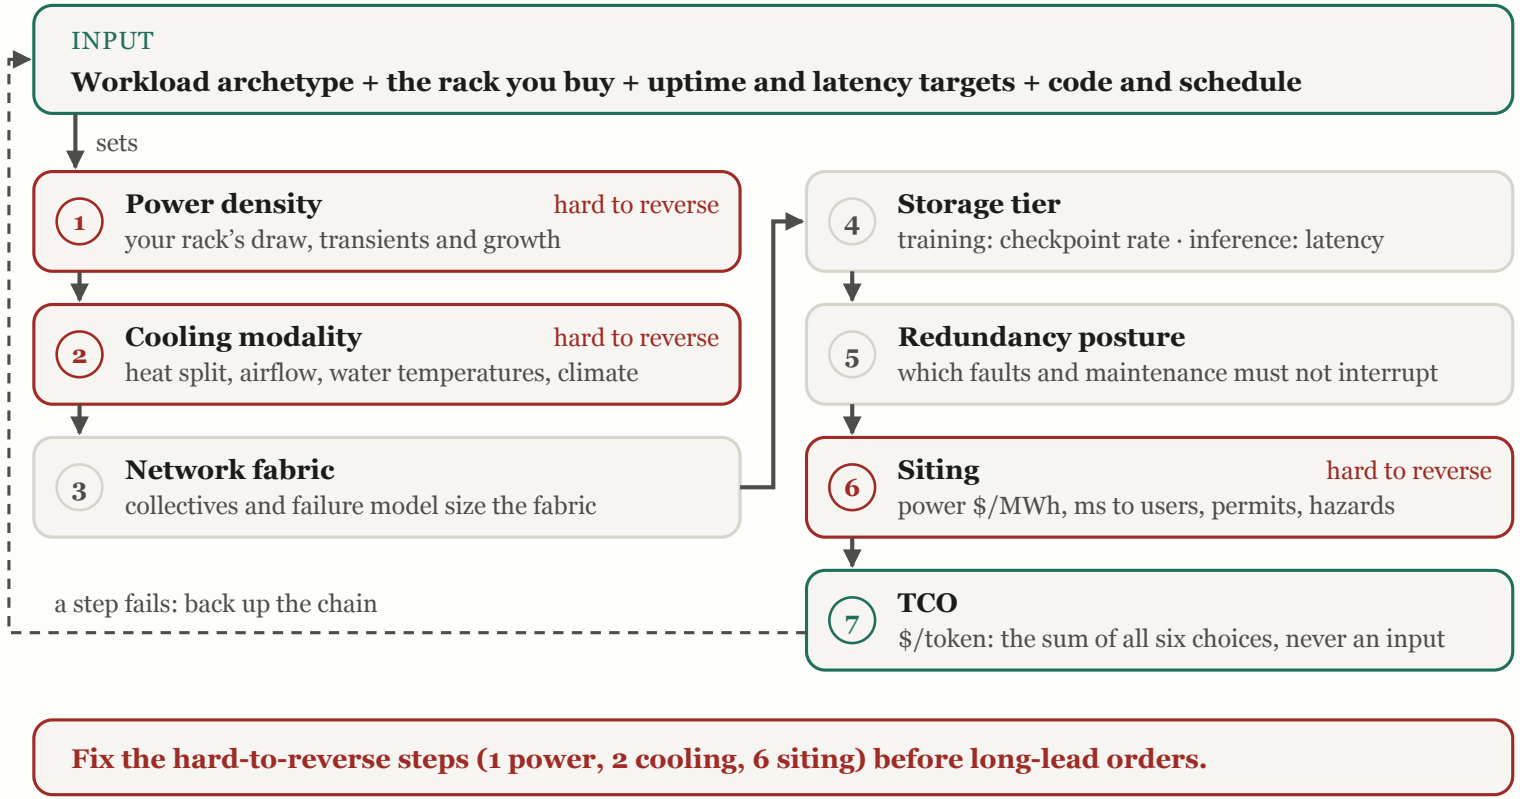


**Fix the hard-to-reverse steps (1 power, 2 cooling, 6 siting) before long-lead orders.**

Schematic · Source: Guide synthesis (Ch 1.1) · as of August 2026

*Workload leads, but equipment, service, site, code, schedule and commercial constraints can change every downstream choice. Iterate the full vector before fixing costly-to-change decisions.*

THE SELECTION DISCIPLINE

# Select from measured requirements, not workload labels

## Measure the workload

— Collective traffic and placement sensitivity

— Rack power, heat split, and transient shape

— Latency SLO and failure consequence

— Checkpoint, retry, and degradation behavior

— Growth envelope and hardware roadmap

## Close the design basis

— Choose cooling from heat and the OEM envelope

— Choose fabric from traffic, topology, and SLO

— Choose redundancy from consequence and recovery

— Choose site from power, latency, code, and schedule

— Declare the operating point, margin, and triggers

SO WHAT

Training and inference labels frame test cases; neither fixes the facility. Contract the requirements vector and declared operating points before costly commitments.

THE MEASUREMENT PLAN

# Five workload signals, one requirements-led selection process

| WORKLOAD SIGNAL | MEASURE | DECISION UNLOCKED | REVALIDATE AGAINST |
|---|---|---|---|
| **Synchronous collectives** | Traffic matrix; placement sensitivity | Fabric topology and rail count | Job size, scheduler, link-failure SLO |
| **Mixed RL pipeline** | Rollout/trainer concurrency and dataflow | Network and storage balance | Dominant phase; staleness tolerance |
| **Latency-bound serving** | p95/p99 latency, locality, request bursts | Placement and capacity headroom | Model size, batching, cache policy |
| **Queued batch serving** | Arrival curve, deadline, retry tolerance | Oversubscription and curtailment | Queue SLO, power contract, recovery |
| **Edge or regional serving** | Latency budget, data gravity, outage cost | Site and fault-domain placement | Jurisdiction, reach, fleet failover |

THE DENSITY RAMP

# The wall, the rack, and the roadmap

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

One practitioner air-cooling reference — not a universal product-independent cliff

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

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

NVIDIA GB200 NVL72 by HPE nominal rack TDP and liquid/air heat split

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

## ~600 kW

Kyber-class rack (NVIDIA: H2 2027; SemiAnalysis: 2028) — per-rack planning figure

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

COSTLY TO RETROFIT

# The density-ramp trap: a hall scoped for today's air-cooled racks cannot absorb the next liquid generation — not the floor, not the power chain, not the cooling plant.

*Don't build for the frontier on day one. Make the substrate that is costly to retrofit — floor loading, water, electrical headroom — carry two generations; match the IT fit-out to now.*

TRAINING IS ONE SYNCHRONOUS MACHINE

# Reliability is the training scoreboard

## 1 every ~3 hr (419 in 54 days)

Llama 3 405B: 419 unplanned interruptions in a 54-day window on up to 16,384 H100s

META · AS-OF 2024-07

## ~7 days (512-H100 cluster, typical)

Typical MTBF SemiAnalysis reported for a 512-H100 cluster — not a fleet scaling rule

SEMIANALYSIS, AI NEOCLOUD PLAYBOOK AND ANATOMY · AS-OF 2024-10

## 90% vs 96% scenario

Illustrative 90%-versus-96% training-goodput sensitivity — not an industry baseline

GUIDE ANALYSIS · AS-OF 2026-09-08

DECISION AND CONSEQUENCE

# A market forecast cannot allocate your fleet.

*Use measured training, prefill, decode and host demand to set the service and capacity boundary.*

THE HEADROOM TRAINING CAN'T TOUCH

# ~21% vs ~3%

## Power-oversubscription headroom, inference vs training — uncorrelated peaks vs synchronized ones

*Independent requests don't peak together; a synchronous training step does. The headroom sells more compute per megawatt — capture it in the design basis or strand it.*

PATEL ET AL., POLCA · AS-OF 2024

THE HYBRID MIDDLE

# 70–80%

## Rollout generation takes 70–80% of training time in AgentMath's synchronous Qwen3-8B baseline (March 2026)

*An RL cluster is a fleet of inference engines feeding a small trainer. Scope it as pure training and you pay for trainer-grade fabric on GPUs that spend their lives decoding.*

AGENTMATH, ARXIV 2512.20745V3 · AS-OF 2026-03-02

THE EDGE TEST

# ~30 / 50 / 100 ms

## The perceptibility thresholds — the only latency budgets that ever justify the edge

*If a centralized region already clears your SLO, the edge is a more expensive way to do what the region does better. Build micro-sites for physics-bound budgets, not preferences.*

QUALCOMM · AS-OF 2018-10

ACT II

# How you buy it

*Build, retrofit, lease, or rent. The right answer is almost never the one that minimizes unit cost — the expensive mistakes are made in time and optionality, not per-GPU-hour.*

THE FOUR LEVERS

# Four procurement archetypes, one bet on your own forecast

| ARCHETYPE | TIME-TO-POWER | CAPITAL | BEST-FIT DEMAND |
|---|---|---|---|
| **Greenfield self-build** | 24–36 mo construction once power is secured; grid waits can overlap (JLL, January 2026) | Highest — full capex | Durable, large, well-forecast, multi-year |
| **Brownfield retrofit** | Months — bounded by the existing shell | Moderate — ~$5–6M/MW for 20–40 kW racks (Enabled Energy, July 2025) | Bridge capacity; modest-density inference |
| **Colocation (wholesale/retail)** | 6–12 mo to a live cluster | Capex-light — lease plus your IT | Medium-term, scaling, demand-uncertain |
| **Neocloud / GPU rental** | Days to weeks | Opex only | Spiky, short, experimental, burst overflow |

THE GATE BEHIND EVERY BUILD

# The line you wait in for power

## >2,060 GW

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

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

## 18–24 mo

Greenfield hyperscale-campus estimate — building is fast; power is the gate

MASTT, TYPES OF DATA CENTERS · AS-OF 2026-04

THE MASTER PROCUREMENT FORK

# Is your demand durable, or is it a bet?

## Durable — price ownership

— First-party load, signed take-or-pay, multi-year

— Self-build or long build-to-suit lease

— Choose ownership when matched NPV pays

— Accept the capex and the interconnection queue

## Uncertain — price the exit

— Spiky, short, experimental load

— Colo anchor plus neocloud overflow

— Pay only for an enforceable exit right

— Check cancellation and paid standby

SO WHAT

Price a hybrid against the same service envelope: committed base load, qualified overflow, paid reserve and actual exit rights.

THE RENTAL MARKET, MID-2026

# ~+40%

## Rise in the 1-year H100 rental contract index, Oct 2025 to Mar 2026 — on-demand largely sold out

*Optionality is getting more expensive as capacity tightens. If renting is your hedge against the queue, lock the capacity before the wait year, not during it.*

SEMIANALYSIS H100 RENTAL INDEX · AS-OF MARCH 2026

ACT III

# Will it pencil?

*Four numbers decide the return: capex per watt, the economic earning life you underwrite, the utilization you actually achieve, and the price the market leaves you after deflation.*

EARNING-LIFE SENSITIVITY **Which earning life supports the commitment?**

# Shorter earning life raises annual cost

Annual cost, $B per year · same 1 GW fleet, fixed Epoch inputs

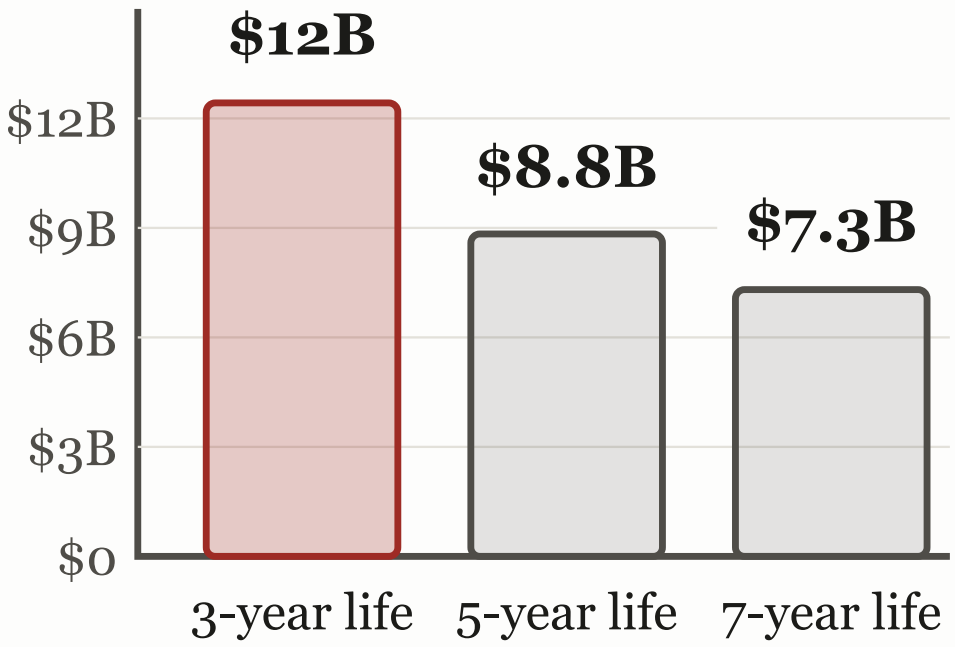


**≈70% more per year**

3-year vs 7-year life

Underwrite earning life before committing capital.

Book expense is a separate schedule: stretching it cannot create demand or disposal cash.

Guide scenario · Source: Epoch AI + Guide fixed-life calculation · as of May :

*With Epoch's May 2026 inputs fixed, three earning years cost about 70% more annually than seven. The exact five-year case is $8.8B/year; book depreciation is a separate expense schedule.*

THE SWING, QUANTIFIED

# $12B / $8.8B / $7.3B

## Annual TCO on the same 1 GW fleet at a 3-, 5-, or 7-year assumed IT life

*The 3-, 5-, and 7-year lives are guide-model assumptions. Goldman Sachs estimates 4–6 years; 2–3 years is a contested obsolescence bear case, not an observed universal life.*

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

MATCHED RENTAL CASH COMMITMENT

# $3.0 /GPU-hour

## Idle reserved hours still incur the calendar rental rate

*Guide assumption: idle and reserve hours remain paid through the annual commitment. Chapter 1.8 compares matched cash flows; Chapter 2.5 tests receipts against debt service.*

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

THE PRICE THE MARKET LEAVES YOU

# ~9–900x/yr by threshold

## LLMflation — the yearly fall in inference cost at fixed quality; recent measurements run faster

*The most dangerous pro-forma assumption is that today's token price holds. It doesn't. Margin survives only if your cost falls faster than the market's price.*

EPOCH AI, LLM INFERENCE PRICE TRENDS · AS-OF 2024–2025

THE PRICE LADDER

# Rental channels carry different service obligations

*Compare dated, workload-specific quotes with matched capacity, cancellation, support and locality.*

| | |
|---|---|
| Neocloud quote (matched commitment and SLA) | Dated, workload-specific quote |
| Hyperscaler on-demand (managed, enterprise SLA) | Matched service quote |
| **Investment handoff** | **Matched service → 1.8** |

*The dated neocloud comparison is a rental example. Chapter 1.8 prices complete owned and rented cash flows after technical qualification.*

THE THROUGH-LINE

**The scoreboard is not megawatts built. It is whether this asset, financed this way, earns its cost of capital over the economic life the workload actually obeys.**

TAKE IT HOME

# What to do Monday

01 Name the archetype in writing — workload profile sheet, ramp curve, design basis — before any long-lead order.

02 Sort decisions by cost to reverse; over-build only what is costly to change (slab, water, voltage path) and choose the site once.

03 Plumb the substrate for two density generations; fit out the IT for the current one.

04 Match redundancy to explicit maintenance, fault and recovery outcomes; workload labels do not select N, N+1 or 2N.

05 Pick procurement by demand durability, not unit cost — anchor the certain base, rent the volatile tail.

06 Underwrite to the economic life and utilization you can contract; stress the pro-forma against token-price deflation.

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 1.1 The archetype framework and the requirements cascade

Ch 1.6 Build vs buy vs rent, and the hybrid patterns

Ch 1.7 The requirements-and-consequences lookup table

Ch 1.8 Depreciation, breakeven, and the ROI stress tests


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 (1/3)

Epoch AI, Amelia Michael and Ben Cottier, May 2026 model and linked workbook. May 19 is the article’s up-front CapEx CSV revision date; annualized-cost CSV is marked May 14. This record’s asof_date identifies the CapEx data revision, not article publication or verification. · as-of 2026-05 · https://epoch.ai/data-insights/ai-datacenter-cost-breakdown

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

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

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

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

SemiAnalysis, AI Neocloud Playbook and Anatomy · as-of 2024-10 · https://newsletter.semianalysis.com/p/ai-neocloud-playbook-and-anatomy

Guide analysis — stipulated sensitivity scenario; no claim of an industry measurement. · as-of 2026-09-08

SOURCES

# Every figure in this deck traces to the guide's dated register (2/3)

Patel et al., POLCA (Microsoft Research), ASPLOS 2024 · as-of 2024 · https://par.nsf.gov/servlets/purl/10505206

AgentMath, arXiv 2512.20745v3 · as-of 2026-03-02 · https://arxiv.org/pdf/2512.20745v3

Qualcomm · as-of 2018-10 · https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/presentation_-_vr_and_ar_are_pushing_connectivity_limits_-web_0.pdf

JLL, 2026 Global Data Center Outlook · as-of 2026-01 (JLL outlook); bands compiled 2026-08 · https://www.jll.com/content/dam/jllcom/en/global/documents/reports/research-reports/26-research-global-data-center-outlook-new.pdf

Enabled Energy, originally published as DCD opinion, 2025-07-08 · as-of 2025-07 · https://enabledenergy.net/news/ai-legacy-data-centers/

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

Mastt, types of data centers · as-of 2026-04 · https://www.mastt.com/blogs/types-of-data-centers

SemiAnalysis H100 Rental Index · as-of March 2026 · https://newsletter.semianalysis.com/p/the-great-gpu-shortage-rental-capacity

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

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

SOURCES

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

Epoch AI, LLM inference price trends · as-of 2024–2025 · https://epoch.ai/data-insights/llm-inference-price-trends

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.*