Compute × Power Trading

Thirteen-tab synthesis dashboard for the compute-driven power complex.

Section-by-section feature guide.

Daniel Kaufman · June 2026 · Research and education — not investment advice.

Contents

01 — Market State
02 — Compute Factors (conversion stack, spark spread, regional pipeline, lease curve)
03 — Peak / Off-Peak / 24×7 (shape decision)
04 — Transmission & FTR
05 — Options Lab (Black-76, bullet vs monthlies, Asian, vega-neutral sizer)
06 — Trade Builder (A1, A2, P1, P2, P3, B3 + scenario tester)
07 — Margin (ICE / Nodal)
08 — NG Seasonality (UNG vs NG Jan27 options)
09 — Index Dispersion (SD vs OCPI factor decomposition)
10 — Compute Heat Map (physical compute × power)
11 — Peak/Off-Peak Convergence
12 — Take-or-Pay vs Spot
13 — Data & Method

TAB 01

Market State

What it does

Top-of-dashboard reference. Six KPIs anchor the conversation: PJM wholesale price Q1-26 ($136.53/MWh, +76% y/y, 63% data-center attributed), PJM 26/27 capacity ($329.17/MW-d at the FERC cap, vs $28.92 in 24/25), Henry Hub spot, the SDB200RT Mar-26 close (5.48), illustrative PJM-W RT Peak forwards for Jul/Aug-26, and the NoVa pipeline-to-inventory tightness ratio (16.8 GW / 4.0 GW = 4.2×). Two cards: a Silicon Data Q1-26 index time-series chart and a 'Why power is the compute trade today' table.

How to use it

- Use as the daily landing tab — these KPIs are the universe-state summary.
- Distinguish PRINT (dated, sourced) from INDICATIVE (yellow input, re-mark required) at a glance via the pill colors.
- The Silicon Data chart anchors the SDB200RT path — note the March squeeze from 4.40 to 6.11 and the partial mean-reversion to 5.48.
- The 'Why power' table is the elevator pitch for the rest of the dashboard.

The underlying math / data

- Published prints: PJM wholesale Q1-26 from PJM Monitoring Analytics; capacity from the FERC-capped 26/27 BRA result; HH from EIA; Silicon Data indices from SiliconNavigator + SiliconIndex.
- Illustrative forwards: PJM-W peak Jul-26 / Aug-26 shown at $95 / $108 — re-mark against actual ICE/Nodal quotes before any trading decision.

Reading the signal

When the PRINT KPIs move materially day-over-day, the INDICATIVE forwards likely need re-marking too — they're set against a stale snapshot. Treat the PRINT row as ground truth and the INDICATIVE row as a planning starting point.

TAB 02

Compute Factors

What it does

The conversion stack — chips, racks, PUE, utilization → grid MW — plus the spark-spread calculator and the regional pipeline → power demand visualization. Also includes the C1 lease-rate term-structure solver. Two columns: left has the chip-to-MW + spark spread; right has the regional pipeline chart, the basis-grid table, and the lease-curve solver.

How to use it

- Start in the conversion stack — pick chip class, GPU count, PUE, utilization. The output shows IT load, post-PUE grid MW, annual GWh.
- Move to the spark-spread calculator immediately below — set GPU rental price and host power price, see the energy share of revenue.
- Use the regional pipeline chart to see which zones absorb the biggest builds (NoVa, ERCOT N).
- The lease-curve C1 solver derives implied depreciation from spot/1-yr/3-yr lease stack — flag mispricing vs realistic obsolescence (~30%+/yr in $/FLOP terms).

The underlying math / data

- Conversion: IT load = chip TDP × GPU count; facility = IT × PUE; grid MW = facility / 1000.
- Spark spread = GPU rental – (kWh per GPU-hr / 1000 × power $/MWh).
- Chip TDPs (all-in rack basis): H100 1.4 kW, H200 1.6, B200 2.0, GB200 NVL72 1.67 (120 kW / 72 GPUs).
- Lease-curve solver: implied annual depreciation = (spot – 1yr_rate) / spot.

Reading the signal

The spark spread is fat (~4-6% energy at typical Blackwell rentals). Variance comes from the compute leg, not the energy leg. The lease curve's implied depreciation should track realistic obsolescence — when it collapses below 20%/yr, the term market is paying holders to lock supply (a scarcity signal).

TAB 03

Peak / Off-Peak / 24×7

What it does

Shape decision tool. Inputs: a fixed 24×7 block price, peak price, off-peak price, and a peak-hour fraction. The dashboard computes the implied 24×7 cost via the peak/off-peak weighted blend and compares to the fixed 24×7 offer. Output: which shape is cheaper, by how much. Also includes a 'Load Shape: DC vs System' chart that visualizes the divergence between flat DC load and system peak/off-peak demand.

How to use it

- Set your zone's peak and off-peak forwards.
- Set the fixed 24×7 quote from your counterparty.
- Adjust peak fraction to the actual peak-hour count for the contract month (typically 16/24 × 5/7 ≈ 47.6% for NERC peak).
- Read the recommendation — buy the cheaper shape, sell the richer if you're the supplier side.

The underlying math / data

- Implied 24×7 = peak × peakFrac + offPeak × (1 – peakFrac).
- Savings = fixed 24×7 quote – implied 24×7. Positive = blend is cheaper; negative = fixed is cheaper.
- Peak hour count varies by month — adjust for federal holidays.

Reading the signal

A DC operator with flat 24×7 load wants the cheapest 24×7 shape. When the fixed 24×7 quote is rich vs the peak/off-peak blend, replicate via two contracts instead. The chart underneath shows why the dispersion exists — DC load is flat, system load has a daily curve, and the off-peak hours undercut the average.

TAB 04

Transmission & FTR

What it does

The transmission constraints that price the compute corridors. A table of $/MWh basis costs by corridor (DOM ↔ West, ComEd ↔ West, AEP ↔ West, etc.). Below it, a Tightness Index chart and an FTR/CRR position calculator. Also includes the A1 queue-to-basis signal model — how to read the public PJM/ERCOT queue and convert it to basis exposure.

How to use it

- Start with the corridor table — identify which corridors are paying the highest basis.
- Use the FTR calculator to size a position based on MW and target P&L; per MW.
- The A1 queue-to-basis signal: high pipeline-to-installed ratio + low expected curve lift = basis underpriced.
- Cross-reference against the heat map (Tab 10) for the geographic context.

The underlying math / data

- Basis = downstream zone LMP – reference hub LMP.
- FTR settlement = MW × (Δbasis – auction price), accrued hourly.
- Queue signal: A1 score = pipeline MW / installed capacity, weighted by lead time to COD.

Reading the signal

DOM-West, AEP-West, and ERCOT North-Houston are the cleanest compute-driven basis trades. When the corridor cost in the table exceeds the FTR auction-clearing price by more than the calculator's break-even, that's an actionable long. Use the queue signal to confirm the build is real and dated.

TAB 05

Options Lab

What it does

Four interactive option calculators in one tab. (1) Black-76 pricer for monthly options on power futures — the building block. (2) Bullet vs monthlies dispersion (Trade P1): implied correlation calculator for buying monthlies and selling the bullet. (3) Daily vs monthly Asian convexity (Trade P2): the Jensen wedge between daily strikes and the monthly average. (4) Vega-neutral package sizer — automates the dollar-vega match across two legs.

How to use it

- Use the Black-76 pricer to anchor any individual leg before assembling the package.
- For P1: enter monthly vols + summer bullet vol, see the implied cross-month correlation. ρ above 0.85 is rich (sell bullet, buy monthlies); below 0.55 is cheap (other direction).
- For P2: enter daily vol + monthly Asian vol, see the Jensen convexity gap. Sell daily call spreads / buy monthly straddles when the gap exceeds 2× fair value.
- Vega-neutral sizer: balances the long-vol and short-vol legs to net zero dollar vega — pure dispersion exposure.

The underlying math / data

- Black-76: $C = e^{(-rT)} [F \times N(d1) - K \times N(d2)]$, with $d1 = (\ln(F/K) + 0.5\sigma^2 T) / (\sigma\sqrt{T})$.
- Implied ρ from bullet vol: $\sigma_{bullet}^2 = \Sigma w_i^2 \times \sigma_i^2 + \Sigma w_i \times w_j \times \rho \times \sigma_i \times \sigma_j$.
- Asian convexity: $\sigma_{asian} \approx \sigma_{daily} / \sqrt{(N/3)}$ for daily samples averaged to monthly.
- Vega-neutral: size_2 = size_1 × (vega_1 / vega_2).

Reading the signal

The cleanest single trade in power options is the bullet-vs-monthlies dispersion (P1) — almost always available, structurally short cross-month correlation. The Asian convexity (P2) is more episodic but pays out fast when realized vol diverges from the implied. Both isolate vol-of-vol from level vol — pure realized-correlation trades.

TAB 06

Trade Builder

What it does

Six trade structure cards (A1, A2, P1, P2, P3, B3) with each leg, sizing, and economics laid out, plus a scenario tester that lets you stress the portfolio under a chosen market path. The structures span power basis (A1, A2), vol dispersion (P1, P2, P3), and the compute spark spread (B3, at-listing).

How to use it

- Use the six cards as a structured trade catalog — copy the leg specs into a paper trade or live execution.
- The scenario tester runs the chosen structure under a market path you set: forward shift, vol shift, basis shock.
- A1 (long DOM-zone basis) is the highest-conviction Today trade — clean expression of compute supply curve via power forwards.
- P1 / P2 / P3 are the vol-dispersion trades — applicable as soon as the relevant options markets show width.
- B3 is parked until compute futures list; it's the at-listing version of the spark spread.

The underlying math / data

- Each card lists: legs (long/short), sizing (MW × $/MW × tenor), expected payoff, primary risk.
- Scenario tester: ΔP&L; = Σ leg payoff under the scenario path. Vol shock and basis shock applied independently.

Reading the signal

The Trade Builder is the synthesis layer — it takes the analytics from Tabs 1-5 and gives you executable structures. Read the structure cards top-to-bottom in conviction order; A1 is the modal trade; P-trades are the vol overlay; B3 is the destination.

TAB 07

Margin (ICE / Nodal)

What it does

Indicative initial-margin reference table for the contracts in scope (PJM-W RT Peak future, ERCOT North future, monthly options, daily option strips, FTRs) by venue, plus a portfolio-offset slider panel and a strategy margin calculator that prices the gross-to-net journey on a selected structure.

How to use it

- Use the reference table to anchor expectations — $5-9k IM for futures, ~1.3-1.6× that for short options.
- Adjust the portfolio offset sliders to your venue's actual credits (Nodal expected-shortfall, ICE inter-month spreads, cross-hub haircuts).
- The strategy margin calculator runs the gross-to-net flow on a chosen structure — long straddle vs short bullet, daily call spreads vs monthly, etc.
- The cross-CCP line is the warning: margining the same package at Nodal AND ICE doubles your IM with zero offset benefit.

The underlying math / data

- Gross IM = Σ short-leg IM + futures IM. (Long options consume premium, not IM.)
- Net IM = Gross × (1 – portfolio offset).
- Nodal expected-shortfall offsets are scenario-based; ICE uses spread credits per pair.
- Cross-CCP offset = 0 always.

Reading the signal

Consolidate venues where the dealer pricing is competitive — the margin savings from a single CCP often exceed the spread you'd pay for the better-priced execution at the second venue. Hold cash ≥ 2× net IM: vol-spike margin calls arrive the same morning as the P&L; event.

TAB 08

NG Seasonality (UNG vs NG Jan27 options)

What it does

The cross-product trade: UNG (front-month NG ETF) options vs NG Jan27 futures options. Multiple strike-mapping conventions (spot-ATM, seasonal-equivalent, time-weighted, hybrid), two variants (long convexity via calls, short tail via puts), and a terminal-payoff chart plus a P&L-through-the-life; chart. The point: capture the seasonal risk premium between front and winter via the options curve.

How to use it

- Pick a strike-mapping convention — start with 'hybrid' for the best wedge on puts.
- Choose variant A (long calls UNG / short calls NG Jan27) for long convexity, or B (short puts UNG / long puts NG Jan27) for short tail + carry collection.
- Adjust the terminal-front-price slider to see the payoff distribution; set curve mode to 'collapse' for the worst-case UNG drag scenario.
- Read the P&L-through-life; chart — find where the trough is and size to survive it, not the endpoint.

The underlying math / data

- UNG path simulated through 6 monthly rolls with 1.17% expense drag.
- Entry curve shape: 3.26 / 3.30 / 3.33 / 3.40 / 3.65 / 3.95 / 4.10 (indicative — re-mark).
- Strike mapping options encode the wedge between front-ATM and gas-equivalent strikes.
- Margin: UNG legs at OCC, NG legs at CME SPAN — no cross-margin.

Reading the signal

Roll drag is NOT in UNG option prices (risk-neutral drift = r). This is a position on the curve's seasonal risk premium, not an arbitrage. Variant A's trough is weeks 8-16 (theta paid before the winter window); Variant B's trough is early December. Size to survive the trough, not the endpoint.

TAB 09

Index Dispersion (SD vs OCPI)

What it does

Factor decomposition for the SD – OCPI spread, the compute-index basis. Six factor sliders (A: quotes vs prints; B: hardware/SKU mix; D: tier/SLA mix; E: physical/regional weighting; F+G: forward-vs-spot + Asian convexity; A': latency/timing). Each contributes a basis-point amount to the modeled spread; observed spread minus modeled = your residual edge.

How to use it

- Enter the observed SD–OCPI gap (live).
- Adjust the six sliders to your view on each factor's contribution.
- Residual = observed – modeled. Large residual = either a missing factor or your edge.
- Watch the waterfall chart on the right for the visual breakdown.

The underlying math / data

- SD = quote-based assessment; OCPI = transaction VWAP, regionally weighted.
- Each factor contribution in % of reference index level.
- Factor D (tier/SLA): hyperscaler index runs ~3.0× neocloud — the largest single structural component.
- Factor B (SKU): OCPI samples RTX + A100; SD does not. Material in markets where prosumer rentals dominate the print.

Reading the signal

The two structural drivers are D (tier) and B (SKU). Both become knowable the moment CME publishes the SD methodology in detail. Until then, the residual term is where the edge sits — model the band before launch and you quote the spread; everyone else discovers it.

TAB 10

Compute Heat Map

What it does

Geographic visualization of the basis grid. Eight DC clusters (NoVa, DFW, Columbus, Atlanta, Chicago, Phoenix, Santa Clara, plus Houston as a power node) plotted by approximate geography, with bubble size = operational MW and bubble color = power state (red largest shock, orange scarcity, amber cheap-but-congesting, green liquid, grey bilateral). Blue diamonds overlay the actual power trading nodes that price each cluster.

How to use it

- Hover any bubble for the underlying MW and pipeline figures.
- Cross-reference bubble color (load state) with diamond proximity (trading node) — the diamond is your execution venue for the cluster's basis.
- Use as the geographic context for any trade idea — every position lives in a specific bubble + diamond pair.
- Red bubbles (NoVa, DFW) are the modal long-basis trades; grey (Atlanta) are bilateral and harder to access.

The underlying math / data

- Bubble radius proportional to √MW (so visual area scales with MW).
- MW figures from JLL/CBRE/Cushman & Wakefield Q1-26 reports.
- Power state classification: based on Q1-26 wholesale + capacity print signatures.
- Schematic geography — not a survey map; positions approximate.

Reading the signal

A power shock in a heavily-weighted cluster (NoVa, DFW) moves OCPI (regionally weighted) while SD (composite) lags, and the same shock prices that cluster's blue-diamond power node. One event → OCPI basis + power-node basis + operator spark spread, all correlated. The heat map is the picture of that correlation structure.

TAB 11

Peak/Off-Peak Convergence

What it does

Models what happens to the on-peak/off-peak price differential as flat data-center load fills the demand stack. Inputs: system peak GW, system off-peak GW, flat DC block GW, stack knee, slope above/below, scarcity cap. Output: implied peak and off-peak prices after the DC block is added, and the implied convergence ratio.

How to use it

- Set the system peak and off-peak demand for your zone.
- Add a flat DC block size — typically 4-10 GW for a NoVa-style zone.
- Adjust the stack knee (where the supply curve steepens) and slopes above/below.
- The convergence ratio output tells you how much the peak/off-peak differential compresses.

The underlying math / data

- Off-peak price = max(slope_below × off-peak demand, 0).
- Peak price = min(scarcity cap, slope_above × (peak demand – knee) + base).
- DC block adds flat MW across both hours, pushing off-peak demand into the slope-above region first.
- Convergence ratio = (peak – off-peak) / peak; compresses as DC load lifts the off-peak baseline.

Reading the signal

Flat DC load mathematically compresses the peak/off-peak ratio because it shifts the off-peak demand into the steeper part of the supply stack first. Zones with the highest expected DC load growth (NoVa, DFW) will see the largest peak/off-peak ratio compression — that's a tradeable view on the shape (Tab 3) before it materializes.

TAB 12

Take-or-Pay vs Spot

What it does

Capacity decision optimizer for DC operators choosing between a multi-year take-or-pay power contract and rolling spot purchases. Inputs: capacity MW, term length, utilization, ToP fixed price, spot mean, spot vol, embedded cancellation option value. Output: total ToP cost vs spot expected cost ± 1σ band, recommendation, plus a term-cost-path chart showing the cumulative cost trajectory under both options.

How to use it

- Set capacity and term to match your build.
- Use the published ToP rate as your reference (typically 7-15% premium to current spot).
- Set spot mean to your zone's forward curve and vol to your view (25% is a calibrated default for compute-heavy zones).
- Critical: assign a value to the cancellation option. Even $2-5 / MWh changes the recommendation.
- Read the recommendation: ToP strongly preferred when ToP total is more than 1σ below spot expected.

The underlying math / data

- Annual MWh = capacity × utilization × 8760.
- ToP total = (annual MWh × ToP price – annual MWh × option value) × term.
- Spot expected total = annual MWh × spot mean × term.
- Spot term 1σ = annual std × √term (sqrt-of-time scaling).

Reading the signal

Cancellation clauses are puts on power rates; expansion rights are calls; take-or-pay floors with usage flexibility are swing options (gas-trading technology, directly portable). These embedded options are systematically given away free in current capacity negotiations because neither side prices them. A desk with the option toolkit extracts real value structuring around counterparty optionality.

TAB 13

Data & Method

What it does

Provenance ledger and methodology disclosure. Lists every PRINT (dated, sourced) and INDICATIVE (proxy) input across the dashboard, with its source, scrape date, and re-mark cadence. Plus the methodology notes for the trickier calculations: Asian convexity approximation, implied correlation from bullet vol, queue-to-basis signal model, factor decomposition for SD-OCPI.

How to use it

- Reference this tab whenever you're about to act on a number — verify the print is current.
- Methodology notes are the audit trail. If a calculator's output surprises you, the math is here.
- Re-mark cadences are guidance — scrape sources more frequently than indicated when markets move.
- Use as the spec doc when porting to a production environment.

The underlying math / data

- Power references: PJM Monitoring Analytics, ERCOT MIS, ICE / Nodal settlement files.
- Compute references: SiliconNavigator, SiliconIndex, Ornn (OCPI), Bloomberg.
- Pipeline references: JLL year-end 2025, Cushman & Wakefield H2-25, Avison Young Q1-26.
- Margin references: Nodal margin calculator, ICE margin tool, FCM-supplied IM runs.

Reading the signal

Treat this tab as the contract between the dashboard and the user. Every other tab depends on these inputs being accurate; the disclosure is what lets you trust the outputs. When in doubt about a number, start here.

END NOTES

Compliance and operating notes

The dashboard is a modeling tool, not market data. Every yellow input field and every 'indicative' level is a placeholder or a clearly-labeled proxy. Scraped provider levels (OCPI, Silicon Data) are dated and will go stale. Have the disclaimer language reviewed by counsel for your jurisdiction before going live as public-facing content. No analytics or tracking are included in the dashboard; the only browser storage is for last-viewed tab and input choices via localStorage — nothing is sent anywhere.

Going live with real data

Power side: an ISO feed toolkit replaces the basis, shape, and convergence proxies with PJM/ERCOT LMPs, load, and queues. Compute side: licensing Silicon Data and Ornn APIs replaces the index levels, the SD–OCPI spread, and the regional-weight proxy. Wire either into the relevant tab's input defaults (or fetch on load) and the signals go live.

Daniel Kaufman · Kinetic Alpha
kineticalpha.com/projects/compute-power-trading
dkaufmanrisk@gmail.com

Research and education only. The dashboard's indicative inputs are placeholders; production deployment requires live feed integration and counsel review. Not investment advice.