



PF

Power Factor in Data Centers

Real power, reactive power, and why it still matters at hyperscale

Part 1: The basics (no EE background assumed) • Part 2: Data centers, from the server to the grid

By Joseph T. Wunderlich, PhD · 2026 · prepared with Claude (Anthropic)

Why should anyone care about power factor?

4.4%

of U.S. electricity used by data centers in 2023

6.7–12%

projected share by 2028,
driven largely by AI

100+ MW


a single modern campus;
some proposals exceed 1 GW

At that scale, every few percent of wasted current capacity means transformers, cables, and grid upgrades that cost millions. Power factor is one of the main measures of how efficiently a facility uses the electrical capacity it pays for.


Start here: AC voltage and current are waves

In alternating current (AC), voltage and current swing back and forth 60 times per second in the U.S. The key question is whether they swing together.

Heater or light bulb: in step




Motor: current lags voltage



In step: every bit of current does useful work. Out of step: some current just sloshes back and forth without doing work, but it still flows through the wires.

An analogy: the mug of root beer



Real power = the drink

What you actually wanted. It does the work: **computing, cooling, lighting.**



Reactive power = the foam

Takes up space in the mug and you paid for the mug.



Apparent power = the whole mug

What the **wires, transformers, and generators** must be sized to carry.

Power factor = how much of the mug is drink.

Three kinds of power, three units



Real power

watts (W, kW, MW)

Energy converted to useful work or heat. This is what the kWh meter bills.

$$P = V \cdot I \cdot \cos\theta$$



Reactive power

volt-amperes reactive (var, kvar)

Energy stored and returned each cycle by magnetic or electric fields. No net work, but real current.

$$Q = V \cdot I \cdot \sin\theta$$




Apparent power

volt-amperes (VA, kVA, MVA)

Simply voltage × current. Sets the size of every wire, breaker, transformer, UPS, and generator.

$$S = V \cdot I$$

The power triangle and the power factor



Pythagoras applies:

$$S^2 = P^2 + Q^2$$

Power factor:

$$PF = P / S = \cos \theta$$

PF runs from 0 to 1. A PF of 1.0 means all current does work.

Lagging: current behind voltage (motors, transformers).

Leading: current ahead of voltage (capacitors).

Worked example: same work, different current

A 100 kW load on a 480 V three-phase supply. Current: $I = S / (\sqrt{3} \times 480 \text{ V})$

	PF = 0.80 lagging	PF = 0.99 lagging
Real power P	100 kW	100 kW
Apparent power $S = P / \text{PF}$	125 kVA	101 kVA
Reactive power $Q = \sqrt{S^2 - P^2}$	75 kvar	14 kvar
Line current I	150 A	122 A
Relative wire losses ($\propto I^2$)	100%	65%

19%

less current for the same useful work

35%

less heating loss in the wires

Rule of thumb: wire losses scale with $1 / \text{PF}^2$. The meter bills the same 100 kW either way; the difference shows up as hotter cables and bigger equipment.

One more wrinkle: distorted current


Electronics (including every computer power supply) don't draw smooth sine-wave current. A simple rectifier (to convert from AC to DC for computers) draws short, sharp pulses near each voltage peak. Those pulses contain harmonics, extra currents at 3, 5, 7... times 60 Hz.

True PF = displacement PF × distortion factor

Displacement PF = $\cos \theta$ (the phase shift from before)

Distortion factor = $1 / \sqrt{1 + \text{THD}^2}$


THD measured relative to the fundamental (the 60 Hz component)



A load can be perfectly “in phase” and still have a poor power factor because its current is badly shaped. Measuring it takes a true-RMS (Root-Mean-Square) meter or power analyzer.

A building's PF is a vector sum, not an average


Real powers add. Reactive powers add with their signs. Leading and lagging loads partly cancel.



	kVA	PF
IT load	504	0.99 lead
Cooling	250	0.80 lag
Sum of kVA	754	
Actual total	706	0.99 lag

$|S_{\text{total}}| \neq |S_1| + |S_2|$. The transformer sees 706 kVA, not 754.

Factories: the classic case for capacitor banks



- Induction motors (Inductive load) **need magnetizing current**, so a plant full of motors runs lagging, commonly around 0.7–0.85
- A capacitor bank (Capacitive load) in parallel supplies that reactive power on site, so it no longer travels from the utility
- Power factor rises toward 0.95 or better; line current and kVA demand fall; utility **power factor charges go away**
- The motors' real power (kW) and the work they do are unchanged
- Example: a 2,500 kVA plant corrected from 0.85 to 0.95 lagging needs about 618 kvar of capacitors

Capacitors raise (improve) the power factor by reducing reactive power.

Correcting power factor has side effects

Capacitor banks are the classic fix for lagging PF (sized on the previous slide). They are not free.

1

Switching transients

Energizing a capacitor bank causes a decaying ringwave of roughly 200 Hz–5 kHz, reaching 140–200% of peak voltage. The ITI (CBEMA) curve exists partly because of this event.

2

Harmonic resonance

Capacitors plus system inductance form a resonant circuit. If it lands near the 5th or 7th harmonic from drives and rectifiers, it amplifies the very distortion you were trying to tolerate.

3

Overcorrection

A fixed bank sized for full load pushes the facility leading at light load, which is exactly the generator problem in Part 2.

Part 2

Power factor in data centers

Less important than it used to be at the server, but still important at the facility and the grid, often in unexpected ways.

From the field: clean power in practice



La Jolla, CA • 1985–86

J. Wunderlich

Director of Projects / Architectural Engineer

About 100,000 sf of hi-tech office and light-manufacturing space. Clean power was a design requirement for the computer and manufacturing areas; power factor was not addressed explicitly.



IBM S/390 Lab, Poughkeepsie • 1996–98

J. Wunderlich

Researcher & Hardware Development Engineer

Mainframe-supercomputers developed with IBM Germany, from about \$2M entry-level to about \$43M fully deployed (today's dollars). Extensive time inside large data centers held to strict clean-power and power-consumption limits. The move from water-cooled bipolar to air-cooled CMOS cut power sharply; AI racks are **now** liquid-cooled again.

At the server: largely a solved problem

1990s – early 2000s


- Capacitor-input supplies drew pulsed current
- True PF often around 0.6–0.7
- Third-harmonic currents add up in the shared neutral wire instead of cancelling
- Required oversized neutrals and special “K-rated” transformers

Today

- Active PFC circuits shape current into a clean sine wave
- PF typically 0.95–0.99 at moderate to full load
- Efficiency standards (80 PLUS tiers) effectively require it
- Much of the old “power factor folklore” comes from the earlier era

Active power factor correction (PFC): a small boost converter inside the supply that forces input current to follow the voltage waveform. Its regulated DC bus capacitor also carries the server through short dropouts (the ITI curve’s 20 ms).


The modern problem: leading PF at light load



- Every supply has an input filter capacitor that draws a roughly fixed leading current
- At full load it is tiny compared with the real current
- At idle, or in a new hall at 20–30% of design load, it dominates
- Multiply by thousands of servers: the whole facility can go noticeably leading

Curve assumes capacitive current = 8% of full-load current; real equipment varies.

Why leading PF worries the facility engineer



Why: A generator controls its voltage by adjusting the current in its rotor field winding. Lagging loads are easy to handle. A leading load magnetizes the machine on its own, so the regulator must keep turning its field down until it runs out of range.

Result: UPS-generator compatibility under leading PF is a real commissioning test item, and it has caused transfer failures.

Power factor drives sizing and cost

≈ 4 MVA

extra upstream capacity for a 100 MW campus at PF 0.95 instead of 0.99 (105.3 vs. 101.0 MVA)

- Equipment is limited by current, so it is rated in kVA; IT loads consume kW.
- Losses scale as $1/\text{PF}^2$, so PF affects operating cost, not just capital cost.
- Utility tariffs often charge on kVA demand or penalize PF below about 0.90–0.95. Well-designed sites rarely pay, which is why PF is monitored.

How UPS ratings evolved

1

0.8 PF

Older UPS: 100 kVA unit delivers 80 kW

2

0.9 PF

Transition as server supplies improved

3

1.0 PF

Common today: kW = kVA

Field story: kVA is not amps



Westlake Oaks office building

Westlake Oaks Executive Park, Austin TX (13 buildings), 1984–85

J. Wunderlich

Project Manager / Architectural Engineer

One building was leased entirely to a computer company for its data center. The tenant's electrical requirement was stated in kVA, but it was read as an amperage. The building's service could not deliver it, and the developer had to trench a new feed from the utility to the main disconnect: roughly \$200,000 extra in today's dollars.




A tenant computer center

Same number, different meaning	208Y/120 V	480Y/277 V
A 1,000 A service delivers	360 kVA	831 kVA
A 1,000 kVA load needs	2,776 A	1,203 A

Amps mean nothing without voltage, and kVA says nothing about useful work without power factor. Always state voltage, phases, and kVA together.


What computers tolerate: the ITI curve



- Steady state: $\pm 10\%$ indefinitely
- Swell to 120% for up to 0.5 s
- Sag to 80% for 10 s, or to 70% for 0.5 s
- Dropout: zero volts for up to 20 ms (about 1.2 cycles)
- Between the lines, equipment keeps running. Below, it may stop but should not be damaged. Above, damage is possible.

Self-protecting loads can shake the grid

Sudden data center load transfers, Dominion zone (MW)



- A normally cleared 230 kV line fault caused a brief voltage dip
- Data center UPS controls saw it and transferred to backup generation, all at once
- With load suddenly gone, grid voltage and frequency rose, triggering a second wave of disconnections in 2026
- Grid operators say the facilities disconnect too early; the same problem was solved for solar and wind inverters with ride-through rules
- FERC ordered NERC in July 2026 to write enforceable standards for these “computational loads”; ride-through requirements are expected in 2027

At the grid: power plants in reverse

1

Grid-participant requirements

Grid operators increasingly expect large loads to provide reactive-power capability, voltage ride-through, and power-quality limits, not just a minimum PF.

2

AI training creates synchronized swings

Thousands of GPUs alternate in lockstep between computing and communicating, producing large, fast swings in real power that can excite grid and generator dynamics.

3

Co-located loads

Data centers seeking power directly from a generating plant, as debated around the Crane (Three Mile Island) restart in PJM, couple load behavior tightly to one generator.

Mitigations for AI load swings

- On-site battery storage to buffer swings
- Firmware-level GPU power smoothing
- Deliberately burning “filler” power during idle phases


Note: the AI oscillation problem is about real power fluctuations, not power factor itself, but it is a closely related power-quality issue.

Where it is heading: DC inside the data center

Today: many conversions



Emerging: 800 V DC distribution



Power factor doesn't disappear; it moves. Thousands of small PFC stages are replaced by a few large rectifiers that become the facility's single grid-facing device for PF, harmonics, and ride-through.

Caveat: where the problems come back



Legacy equipment

Older or cheaper supplies without active PFC (power factor correction) still mix into many facilities.



Lightly loaded halls

Leading PF and poorer harmonic performance at low load, as seen earlier.



Cooling plants

Variable-frequency drives on chillers, pumps, and fans draw harmonic current and are frequently the worse offender now.



GPU nodes off their sweet spot

High-power supplies running far from their design point lose some of their PF and efficiency advantage.

Takeaways

1

Power factor = real power $\div$ apparent power. It measures how much of the current you carry actually does work.

2

Low PF means more current, bigger equipment, and more heat for the same useful work. kW, kVA, and amps are not interchangeable.

3

At the server, active PFC has largely solved the problem.

4

At the facility, PF is a vector sum; watch leading PF at light load, generator compatibility, and capacitor side effects.

5

At the grid, data centers must learn to ride through disturbances; 2026 brought a 3,800 MW load loss and new FERC action.

Power factor moved from a server design problem to a systems-integration and grid-interconnection problem.