



Intelligent power for AI

IQ SST: Distributed, solid-state transformers for AI data centers

Safe harbor

Use of Forward-Looking Statements

This presentation contains forward-looking statements made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995, including, but not limited to, Enphase Energy's business strategies, including anticipated trends and developments in markets in which it operates and in the markets in which it plans to expand; the capabilities and benefits of its technology and future products, including the features, architecture, functionality, performance, reliability, and serviceability; the suitability of its future products for AI data centers, utility-scale battery storage, utility-scale solar, and high-power DC fast charging; the availability of its new products and technologies and the market adoption of same; and its manufacturing operations in the United States. Any statements that are not of historical fact, may be forward-looking statements. Words used such as "anticipates," "believes," "could," "potential," "predicts," "continues," "designed," "estimates," "expects," "goal," "intends," "likely," "may," "ongoing," "plans," "projects," "pursuing," "seeks," "should," "will," "would" and similar expressions are intended to identify forward-looking statements, although not all forward-looking statements contain these words. All forward-looking statements are based on Enphase Energy's current assumptions, expectations and beliefs, and involve substantial risks and uncertainties that may cause results, performance or achievement to materially differ from those expressed or implied by these forward-looking statements. You are cautioned that these forward-looking statements are only predictions and may differ materially from actual future events or results. A detailed discussion of risk factors that affect Enphase Energy's business is included in the filings it makes with the Securities and Exchange Commission (SEC) from time to time, including its most recent reports on Form 10-K and Form 10-Q, particularly under the heading "Risk Factors." Copies of these filings are available on Enphase Energy's website through its investor page or on the SEC website. All forward-looking statements in this presentation are based on information currently available to Enphase Energy, and Enphase Energy assumes no obligation to update these forward-looking statements in light of new information or future events.

Industry Information

Information regarding market and industry statistics in this presentation is based on information available to Enphase Energy that it believes is accurate. Such information is generally based on publications that are not produced for purposes of economic analysis.

AI data centers create a new power problem

Three constraints emerging at the same time: capacity, dynamic response, and efficiency

Capacity gap

Demand is growing faster than active AI-dedicated capacity and available power infrastructure.

Dynamic loads

AI training loads can swing from idle to full power several times per second, stressing the grid.


Efficiency losses

Legacy architectures rely on too many conversion stages between the grid and compute racks.



AI demand is outpacing power infrastructure

U.S. data center power demand could reach ~80 GW by 2031, far ahead of expected AI-dedicated active capacity — and today's power architecture cannot deliver it efficiently.




Source: ABI Research for AI-dedicated active capacity; McKinsey for U.S. data center power demand through 2030; 2031 extrapolated from prior-year trend.

~ 80 GW


U.S. data center power demand by 2031 versus ~26 GW AI-dedicated active capacity

AI also has a dynamic load problem

AI training loads can swing from idle to full power several times a second. Power infrastructure must respond in milliseconds, not seconds.



Graph source: NVIDIA Technical Blog, "How New GB300 NVL72 Features Provide Steady Power for AI," Figure 1 (AI training workload power profile), July 28, 2025.



Picture every light in San Francisco going dark.


Then blazing back on.

Then dark again — several times a second.

That's a 1 GW AI data center.

The shift to 800 VDC architecture

800 VDC Distribution - MV to 800 VDC




Source: NVIDIA white paper "800 VDC Architecture for Next-Generation AI Infrastructure" by Jared Huntington & Mike Tu

Opportunity

- Convert MV AC to 800 VDC and deliver directly to GPU rack
- Reduce complexity by eliminating multi-stage power conversion
- Eliminate the sidecar and move energy storage near the SST

The IQ SST solution: supports two use cases



Enphase's differentiated solution: SST coupled with a BESS in the medium voltage section, and an SST for the compute load in the low voltage section

Enphase differentiated solution

- IQ SST supports two use cases:
 1. **Without LV battery:** Sub-millisecond response time enables us to reflect rapid compute load changes to the BESS SST
 2. **With LV battery:** A DC/DC block to charge and discharge a LV battery which is part of the IQ SST skid
- Reliability: Modular architecture enables very high reliability to achieve 99.999% uptime
- Manufacturing scale and cost: Soft switching and widely available components drive lower cost; supplied from U.S. manufacturing facilities

AI needs a new power architecture – IQ SST

IQ SST will be a distributed solid-state transformer that converts medium voltage AC directly to 800 VDC for next-generation AI racks.

	Today	Tomorrow
Architecture	UPS, PDU + sidecar (PSU, BBU, CBU)	Reduced sidecar and backup infrastructure. More space for compute
Voltage	48 V power delivery	800 VDC (± 400 VDC) power delivery to align with NVIDIA/OCP architecture
Reliability	Important	Distributed redundancy; five-nines availability
Supply chain	Fragile, offshore, custom builds	Supplied from U.S. manufacturing facilities
O&M	Specialized field service team	Modular hot-swap; no specialists

Why Enphase, and why now?

AI is setting a new standard for power delivery. Enphase is ready.

A 20-year head start

The core technology is what we've been shipping for two decades.

~87.8 million microinverters in the field. Nine microinverter generations. Five ASIC generations. Established and scalable manufacturing.

Distributed by design

No single point of failure. 99.999% availability.

Hundreds of power modules per unit with approximately 10% built-in redundancy.

Performance

Single-stage power conversion with sub-millisecond response time.

Fewer failure points. Higher efficiency. Simpler manufacturing. Lower service cost. Built on platforms we already ship at scale.

We changed solar. Now we're bringing distributed power to AI.

20

years of distributed
power electronics

9

Microinverter
generations

~87.8M

Microinverters
shipped

~5.2M

Systems
worldwide

~31.5 GW

of managed
solar systems

~2.5 GWh

of deployed
energy storage

165

Countries
served

~550

patents and pending
applications



Our distributed architecture

Enphase power module

- 4 kW building block; over 98.5% peak efficiency expected
- 250 – 500 kHz switching frequency; control ASIC: Kestrel (proprietary Enphase)
- BDS and UDS GaN for input and output power stage
- Weight: ~8 lbs; plastic enclosure with silicone potting
- Soft-switching design for low EMI

How power modules create an IQ SST rack

- Multiple power modules stacked depending on MV input and power
- Series connection on the input; parallel connection on the output (ISOP)
- Delta configuration; no neutral
- Input has triple redundant short-on-fail circuit (anti-fuse)
- Dual sided backplane for power module connection
- Sub-millisecond response time
- Output has short circuit current limit (fused)



Heart of the Enphase power module

Two breakthroughs, one single-stage converter.



Kestrel ASIC

- 5th-gen custom 22 nm control ASIC powering the single-stage converter
- Multi-core software with built-in communications, security, and safety
- Proprietary tandem function supports fiber-optic interface for SST application
- [Read the Enphase white paper](#)



GaN bi-directional switch

- Bi-directional GaN switching technology proven in IQ9 Microinverters
- Enables compact, high-efficiency power conversion
- Supports low EMI and fast switching in single-stage design
- [Read the Enphase white paper](#)

IQ SST 2.5 MW rack @ 34.5 kV

- Power: 2.5 MW, ambient-air cooled
- Architecture: supercluster of 342 power modules (Δ configuration)
- Will support 35 kV and 15 kV utility interconnection classes
- Designed for five-nines availability
- 10-year limited warranty
- Modular hot-swap design — service with a cobot
- No internal batteries required
- U.S.-based manufacturing
- FEOC-compliant supply chain
- Scales up to 5 MW per skid

Demo system available in Q4 2026



IQ SST roadmap

From functioning power modules to volume shipments

9+

months of active development


120+


Engineers dedicated to IQ SST

20+

Prospective customers & partner engagements

PRODUCT ROADMAP





Designed in Silicon Valley. Manufactured in Texas and South Carolina.

In the heart of the AI revolution, with current U.S. manufacturing and FEOC-compliant supply chain built to scale.

- High-volume automated manufacturing
- Leverages existing manufacturing processes
- Made from standard off-the-shelf components
- Potted, all-plastic enclosure
- FEOC-compliant supply chain



Enphase IQ SST

Intelligent power for AI

Distributed by design. Backed by 20 years of Enphase.




Learn more at enphase.com/iq-sst

Appendix

AI has an efficiency problem

Today's architecture has too many conversion stages. IQ SST will convert medium-voltage AC directly to 800 VDC, reducing sidecar infrastructure and power losses.



Source: Based on NVIDIA, "800 VDC Architecture for AI Data Centers"; NVIDIA Developer Blog, "NVIDIA 800 V HVDC Architecture Will Power the Next Generation of AI Factories," May 2025; Open Compute Project, "Power Architecture Evolution in Data Centers," Sept. 2025.

Rearchitecting AI power — the Enphase advantage



Reduce or eliminate sidecar infrastructure

Less power infrastructure. More space for GPUs. More tokens. More revenue.



Sub-millisecond response

Single-stage power conversion.
>250 kHz average switching frequency



Five-nines availability

Distributed architecture with 342 power modules. No single point of failure



Lower field service cost

Hot-swap without shutdown — no specialized technicians required



U.S. supply chain

Supplied from U.S. manufacturing facilities; standard high-volume components



Semiconductor-based, software-defined

Custom Kestrel ASIC and bi-directional GaN. 98.5% efficiency