SOLUTION BRIEF

800 VDC power rack

Accelerate deployment in next-generation AI data centers

KEY BENEFITS

- Lower cost per token with reduced electrical losses and greater power efficiency
- More AI output per rack and per square foot enabled by higher rack density
- Faster deployment with standardized, prefabricated power integration

Disaggregated architecture supports high-density IT racks

The 800 VDC power rack uses a disaggregated architecture that shifts power products out of the IT rack, maximizing space for compute and allowing more GPUs to be packed densely, communicate more efficiently, and deliver greater value from the IT hardware. Featuring Flex's power shelf for the NVIDIA Vera Rubin platform, the 800 VDC power rack also enables current data centers to support high power density, next-generation accelerated computing racks without costly retrofits.

More compute power. Greater resilience.

NVIDIA and partners are adopting 800 VDC architectures to support 400 kW to 1 MW racks, enabling end-to-end energy efficiency gains of 5 percent to 7 percent and reducing the amount of copper needed by up to 50 percent.[1] When deployed with a future GPU platform, this disaggregated architecture increases available compute power per rack from approximately 125 kW to up to 880 kW.

The Flex 800V power rack includes:

- **8 110kW power shelves** suitable for NVIDIA Vera Rubin platform
- **Up to 24 RU** space for CESS, BBUs and auxiliary equipment
- **Options to customize** as per your power and compute needs

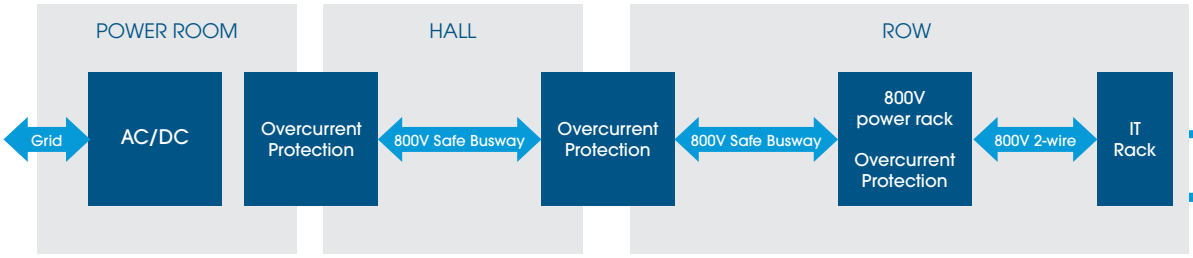

[1]NVIDIA, NVIDIA 800VDC Architecture will power the next generation of AI factories, M. Blake, M. Hsu, et. Al ., May 20, 2025

ABOUT FLEX

Flex provides advanced manufacturing capabilities, innovative power and cooling products, and services that solve for data center power, heat, and scale challenges in the AI era.

Accelerate next-generation data center deployment with Flex.

Increase tokens per dollar at scale

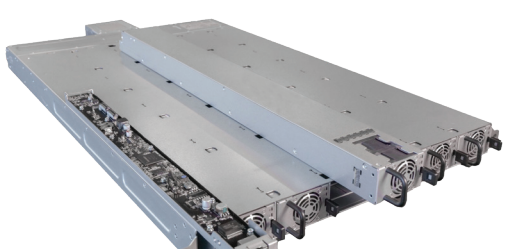

800 VDC to 50 VDC Power Shelf

Standardized, high-voltage, rack-adjacent power simplifies data center deployment and scaling, reducing capital and integration costs across large AI factories. As AI data centers evolve into large-scale inference and training factories, infrastructure economics are increasingly defined by cost per token rather than peak compute alone. Flex's 800 VDC rack-adjacent power architecture is designed to support this model by enabling dense, standardized, megawatt-class deployments that translate installed power capacity into higher effective token production with lower marginal cost.

- **Increases token output** per megawatt by enabling higher GPU density without proportional infrastructure growth
- **Lowers capital** and operational cost per token by reducing non-compute power and integration overhead
- **Improves realized token throughput** by supporting higher operational availability and simplified serviceability

> "Megawatt-scale AI workloads are redefining data center power infrastructure. The Flex 800 VDC power rack pairs advanced power products with global manufacturing scale to enable efficient, scalable power delivery and faster deployment of next-generation infrastructure."
>
> **Chris Butler**
> President, Embedded and Critical Power, Flex

Learn more

Unlock data center capacity

Flex AI Infrastructure Platform

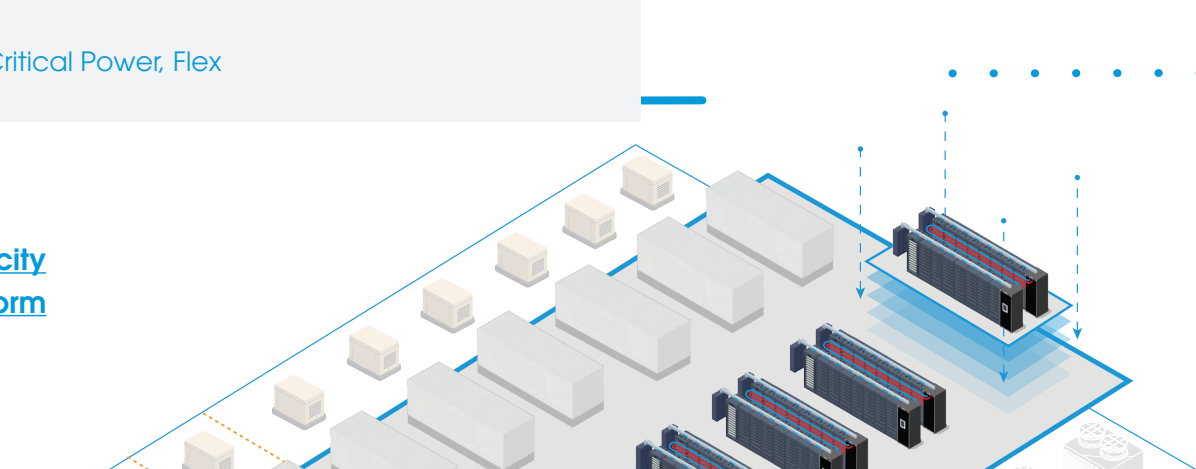

For more information, visit flex.com