

skeleton+

**FREE
DOWNLOAD**

THE EXECUTIVE BLUEPRINT

MAXIMIZE THE FULL POWER OF GPUs IN AI DATA CENTERS

**REDUCE OPEX & CAPEX
WITH SUPERCAPACITORS**



Table Of Contents

The Insider's Guide: AI Data Centers

- "The AI revolution is not coming, it's already here."
- Key Operational Challenges for Data Centers
- Impacts of Data Center Power Challenges
- Why Choose Supercapacitor Technology?
- Energy Storage Technology Comparison

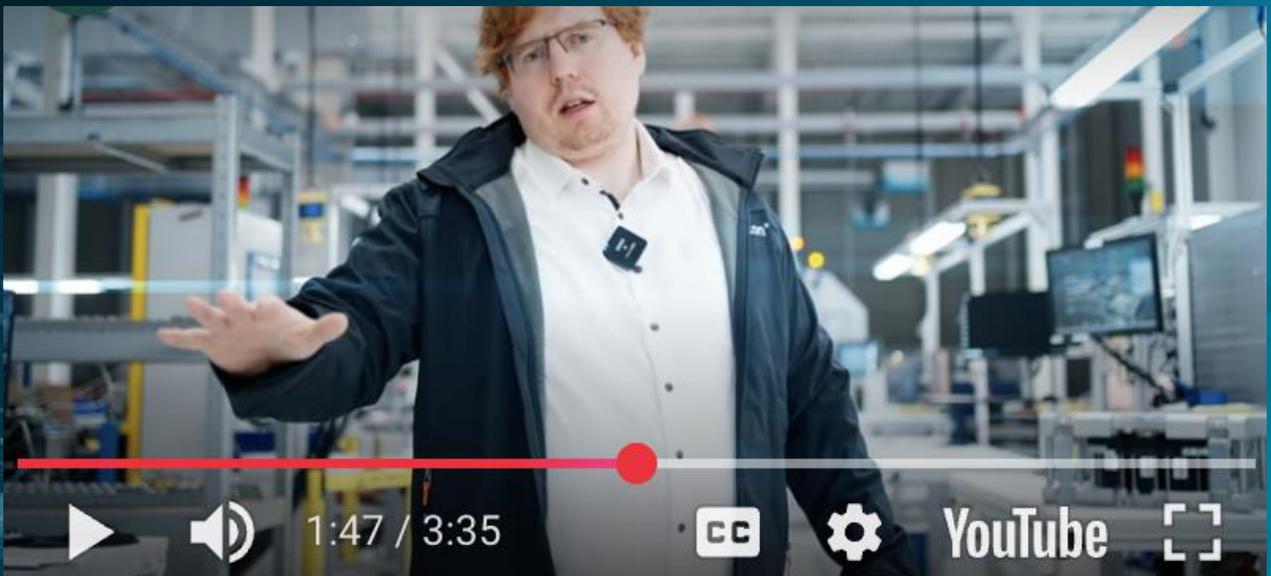
GrapheneGPU: For GPUs' Full Compute Power

- GrapheneGPU PSC50
- Skeleton's GrapheneGPU Solution for AI Data Centers
- Systems and Modules for AI Data Centers at All Voltage Levels

Peak Shaving: Get EVO's Ultra-Low ESR Advantage

- The Costly Risk of Ignoring Peak Shaving
- The Skeleton Advantage
- Maximize Your AI Data Center's Full Potential

"The AI revolution is not coming, it's already here."



Watch Taavi Madiberk, CEO & Co-Founder of Skeleton Technologies' speech now on our YouTube channel [@skeletontechnologies](https://www.youtube.com/@skeletontechnologies)

The Insider's Guide: AI Data Centers

The global AI data center market is experiencing rapid growth, with projections indicating a compound annual growth rate (CAGR) of 27%, potentially reaching a market value of €147.9 billion by 2034.

20%

Estimate of these data centers, utilized for AI workloads

€100BN

Data centers globally could cut by adopting peak shaving technologies by 2030

+20%

Peak power costs can be of total annual energy bills for data centers

11,000+
Data Centers
in Globally

Market.us. (n.d.). AI data center market report. Retrieved March 10, 2025, from <https://market.us/report/ai-data-center-market/>

Intelligent Data Centres. (2024, October 4). Energy-efficient data centres could save up to €25 billion by 2030. Intelligent Data Centres. Retrieved March 10, 2025, from <https://www.intelligentdatacentres.com/2024/10/04/energy-efficient-data-centres-could-save-up-to-e25-billion-by-2030/>

McKinsey & Company. (n.d.). How data centers and the energy sector can sate AI's hunger for power. McKinsey & Company. Retrieved March 10, 2025, from <https://www.mckinsey.com/industries/private-capital/our-insights/how-data-centers-and-the-energy-sector-can-sate-ais-hunger-for-power>

JLL. (2024, March 5). Global data center demand surges despite supply and power constraints. JLL. Retrieved March 10, 2025, from <https://www.us.jll.com/en/newsroom/global-data-center-demand-surges-despite-supply-and-power-constraints>

Key Operational Challenges for Data Centers



Power Demand Spikes

AI GPUs create fast, large power swings (up to hundreds of MW), forcing data centers to oversize grid connections — increasing costs, energy use, and grid stress.



Artificial compute loads

To smooth power fluctuations, many data centers run artificial compute loads between real compute cycles, using more energy, generating excess heat which introduces thermal de-rating of GPUs that reduces their compute capacity.



Grid Stability

Sudden power spikes from data centers can disrupt grid stability by causing voltage fluctuations or power quality issues.



Impacts of Data Center Power Challenges

1. Power Demand Spikes: GPU-Driven Load Volatility

Modern AI GPUs operate with highly dynamic power profiles, causing short, unpredictable spikes in demand. These transient events happen in microseconds, outpacing traditional power infrastructure. To accommodate these peaks, data centers must **oversize grid connections**, power distribution units (PDUs), and backup systems. This leads to inflated **capital expenditure (CAPEX)**, underutilized infrastructure, and additional **grid connection delays**.

2. Artificial Dummy Loads: Idle Energy Between Compute Cycles

To avoid sharp fluctuations in grid draw, many data centers deploy **artificial "dummy loads"** that run between actual GPU compute cycles. These loads keep energy demand constant — at the expense of efficiency.

3. Grid Stability: How One Data Center Can Disrupt an Entire Region

Large AI data centers introduce fast-ramping loads that cause voltage sags, harmonic distortion, and frequency deviations on local grids — all in real time. As the number of high-density compute centers grows, they become **active threats to grid stability**. Without fast-reacting grid support technologies (like supercapacitors), the risk of regional grid instability increases.

Why Choose Supercapacitor Technology?

- + 10× higher power density than li-ion.
- + Designed for millions of charge-discharge cycles
- + Extreme temperature resilience from -40°C to +65°C (-40°F to 149°F)
- + No thermal runaway. No fire risk. Inherently safer in mission-critical infrastructure such as data centers.
- + Sustainable by design — with no rare metals. A lower environmental impact, easier recyclability, and greater supply chain security.



SkelCap
supercapacitor

Energy Storage Technology Comparison

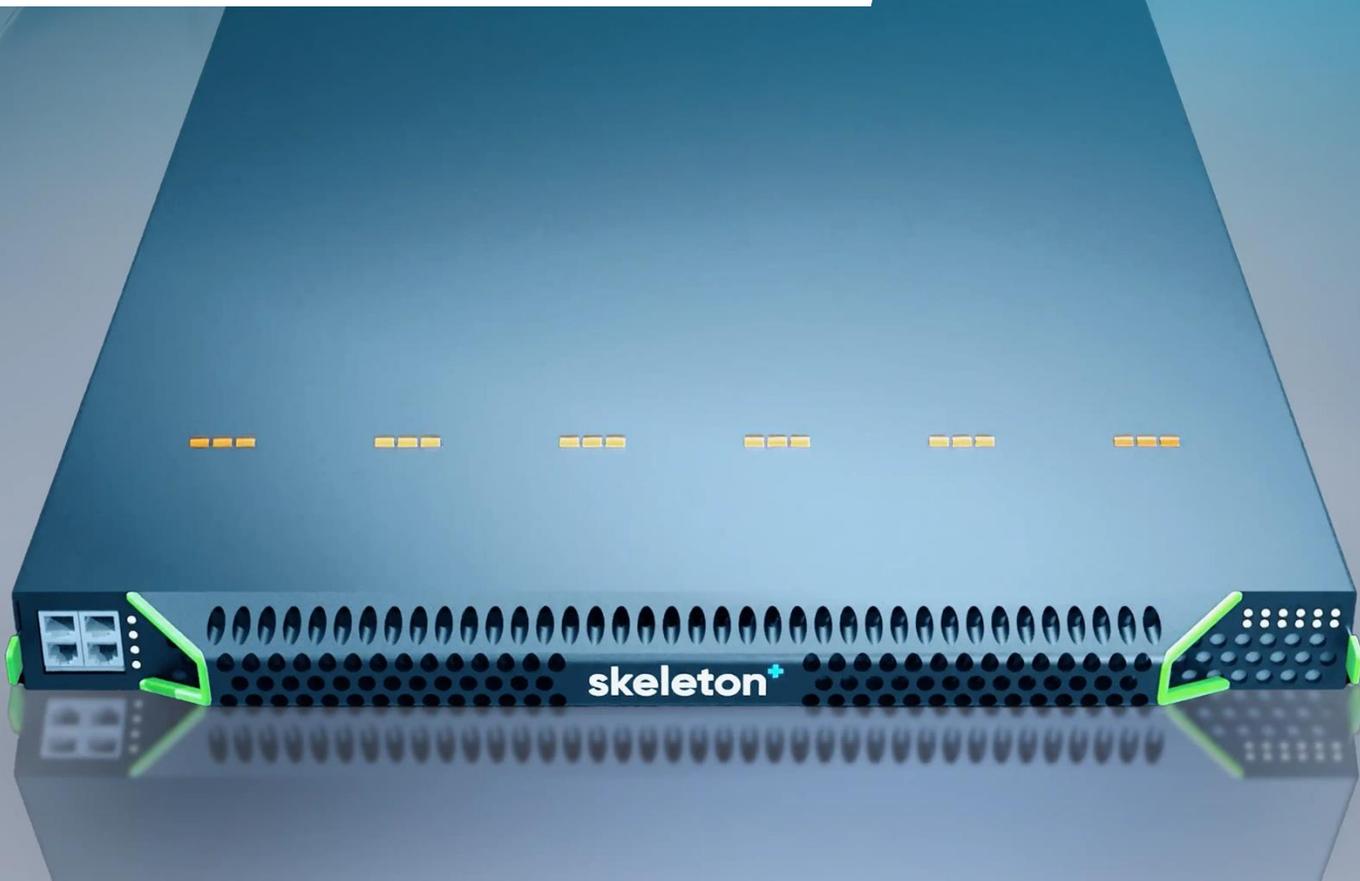
Feature / Metric	Skeleton Supercapacitor	Li-ion (NMC/other)	LFP (Lithium Iron Phosphate)	Li-Caps
Energy Density (Wh/kg)	Low (5–16)	High (150–250)	Medium (90–160)	Medium (30–100)
Power Density (W/kg)	Very High (10,000+)	Medium (250–1,000)	Medium (500–2,000)	High (2,000–5,000)
Cycle Life	Millions	1,000–5,000	3,000–7,000	10,000–100,000
Charge/Discharge Duration	seconds	30 min – 2 hours	15–45 minutes	Minutes
Reaction speed	Microseconds	Milliseconds	Milliseconds	Milliseconds
Temperature Tolerance	-40°C to +65°C	0°C to +45°C	-10°C to +60°C	0°C to +55°C
Thermal Runaway Risk	None	High	Medium	Low
Maintenance Requirement	Very Low	High	High	Medium
Cost per Cycle	Low	High	Medium	Medium
Environmental Safety	Non-toxic, recyclable materials	Toxic materials, complex recycling	Toxic materials, complex recycling	Mixed (depends on blend)
Scalability for Critical Infrastructure	Excellent (modular, fast-acting)	Limited (slow response time)	Limited (slow response time)	Good for certain hybrid systems

skeleton⁺

Supercharge AI Workloads by Eliminating Artificial Compute Loads and Unlocking Full GPU Power

GrapheneGPU

For GPUs' Full Compute Power



skeleton⁺

GrapheneGPU

PSC50



Core Advantages:

Voltage

48 V

Peak power

1250 A

Continuous power

830 A

+ 40%
More FLOPS

- 45%
More Energy
Efficient

- 44%
Smaller Grid
Connection

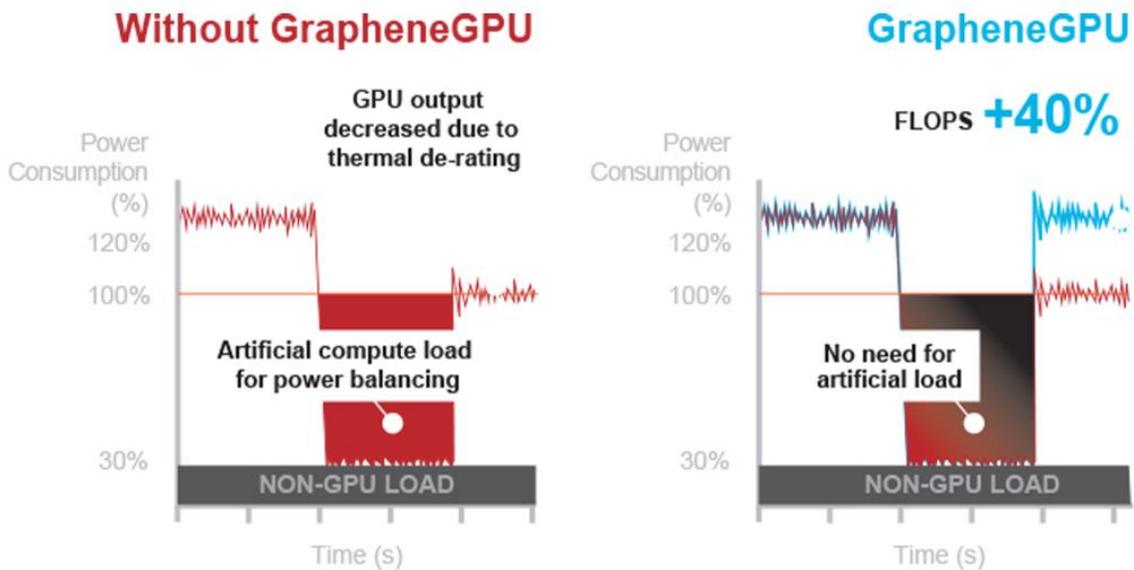
Order now

www.skeletontech.com

Skeleton's GrapheneGPU Solution for AI Data Centers

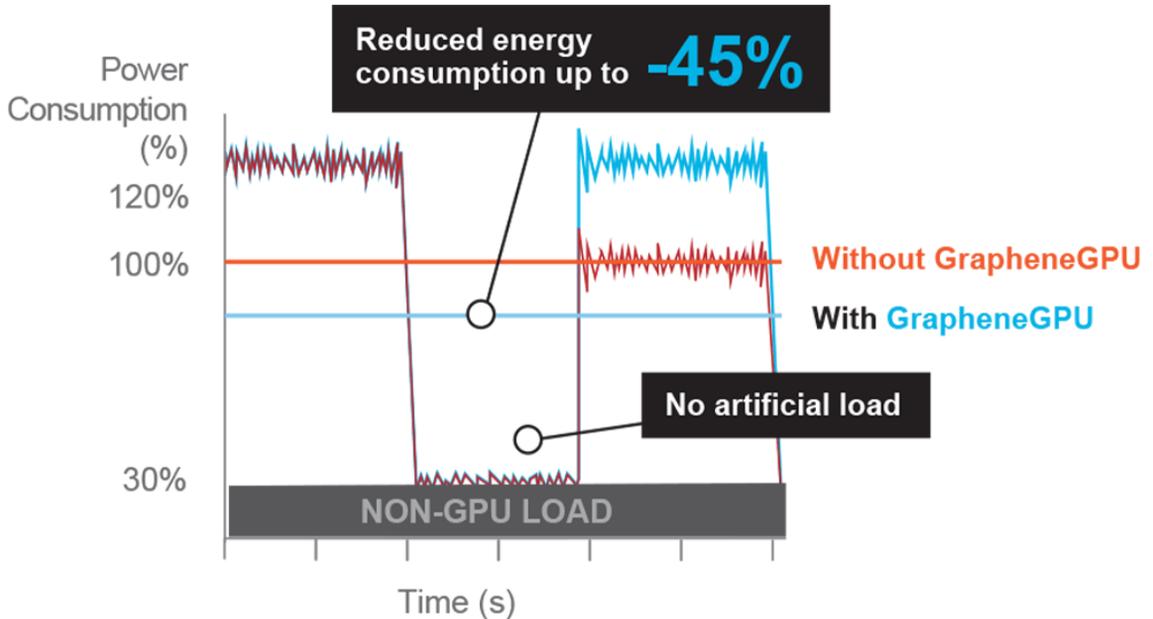
1. Maximize Compute Power (+40% More FLOPS)

GrapheneGPU optimizes GPU performance by charging during idle cycles and discharging during peak load. This flattens the power demand curve, eliminates the need for artificial compute loads, and prevents GPU thermal de-rating, ensuring consistent, full-performance output.



2. Optimize Energy Use

By removing the need for dummy loads and minimizing cooling requirements, GrapheneGPU reduces overall energy consumption by up to 45%, leading to significant savings.



3. Reduce Grid Connection Requirements (Lower CAPEX)

GrapheneGPU smooths power fluctuations by shaving peaks and filling valleys in demand, lowering the peak grid connection needs by up to 44%. This means less infrastructure, lower upfront investment, and faster scalability for new data centers.



Systems and Modules for AI Data Centers at All Voltage Levels



Graphene Sidecar (Voltage based on requirements)

High-performance cabinet rack designed to operate in next-generation AI data centers. Can be equipped with GrapheneGPU PCS 50 or GrapheneGPU PCS 400 modules.



GrapheneGPU PCS 50 (48V)

delivers 60 kW peak power when GPUs compute, and charges when GPUs are idle. Maintaining a smooth power demand from the grid and eliminating the need for dummy loads.



GrapheneGPU PCS 400 (400 V)

delivers 160 kW peak power when GPUs compute, and charges when GPUs are idle. Maintaining a smooth power demand from the grid and eliminating the need for dummy loads.

skeleton⁺

Peak Shaving: Get EVO's Ultra-Low ESR Advantage

In modern data centers, speed, efficiency, and reliability are non-negotiable. Skeleton's EVO cells with ultra-low Equivalent Series Resistance (ESR) deliver rapid power response, minimal energy loss, and reduced heat generation — enabling faster peak shaving, lower cooling costs, and longer component life. It's the key to powering AI workloads and scaling sustainably.

Two versions

SkelCap SCA0300 EVO

Weldable Terminals
PCB-Mountable



How Peak Shaving Helps Lower Operational Costs

Every second of downtime has a high cost.
Every energy spike eats into your margins.
As energy prices rise and demands intensify,
Skeleton supercapacitor-based solutions give data centers a crucial edge: instant power that lowers costs, protects uptime, and strengthens your infrastructure.

Peak Shaving 101

Risks and Advantages





The Costly Risk of Ignoring Peak Shaving

Inefficient CAPEX

Without peak shaving, many data centers overspend on oversized backup capacity to compensate, resulting in higher costs and inefficient space use.

Increased Downtime Risk

Lagging or overloaded backup systems may fail when demand peaks, jeopardizing uptime, SLAs, and trust.

Soaring Peak Energy Costs — and a Hidden Scaling Problem

You pay premium rates for short power spikes lasting just seconds, while also limiting your ability to scale AI infrastructure, underutilizing energy and compute potential.

Higher Maintenance Costs

Sudden surges stress cooling systems, power electronics, and backup generators, leading to faster degradation and higher maintenance costs.



The Skeleton Advantage

Skeleton helps your AI data center control what the grid can't.

Instant Power Delivery

Supercapacitors charge/discharge in microseconds, making them ideal for absorbing sharp, short energy spikes that traditional batteries can't respond to fast enough.

Lower Energy Costs

We don't just flatten the energy demand curve — we eliminate the need for artificial, empty compute loads used to smooth power demand. The result? Lower energy consumption and reduced operating costs.

Increased Reliability & Uptime

Peak loads can cause equipment stress or failure. Skeleton buffers those peaks, reducing downtime risk for critical infrastructure like data centers.

Reduce Overbuilt Infrastructure

Supercapacitors reduce the need to oversize backup generators, batteries, or transformers, translating to smarter CAPEX allocation and lower maintenance.

skeleton⁺

Maximize Your AI Data Center's Full Potential: Plug & Play Integrate GrapheneGPU Into Your Data Center – Contact Our Sales

Contact our energy storage experts today for step-by-step guidance on optimizing energy usage, eliminating artificial loads, and cutting costs with Skeleton's GrapheneGPU — the power behind AI's exponential growth.

www.skeletontech.com

sales@skeletontech.com

Follow us on

[in](#) [X](#) [f](#) [@](#)



**TIME Magazine's Top 30
GreenTech Companies of
2025**

skeleton+

© 2025 Skeleton Technologies. All rights reserved.

This document is the intellectual property of Skeleton Technologies and is provided for informational purposes only. It may not be copied, distributed, reproduced, publicly displayed, or sold in any form without prior written consent from Skeleton Technologies. The information contained herein is subject to change without notice. Skeleton Technologies makes no warranties, express or implied, regarding the accuracy or completeness of this content. Use of this material does not transfer any rights of ownership or copyright.