



AI & Data Center Energy Efficiency

Practical software and operational strategies to maximize computational output while minimizing grid demands, without renegotiating heavy infrastructure contracts.

OTC:ELEK

| The Massive Scaling Challenge

10x
INFERENCE VS SEARCH

The Computational Energy Crisis

AI model generation commands up to an order of magnitude more energy per query than standard database queries. As hardware expansions face severe electrical grid backlogs, software-level efficiency has transformed from a strategic benefit to an operational necessity.


- ❗ Heavy compute scaling causes severe thermal bottlenecks.
- ❗ Long hardware procurement lead times stall development.
- ❗ Operational expenditures rise linearly with parameter count.

Optimized Code & Algorithms

Maximizing Execution Efficiency

Substantial energy savings are unlocked by reducing redundant instruction cycles directly inside compiler configurations and training workflows.

-  **Algorithmic Pruning:** Eliminates silent or non-contributive model weights during active inference phases.
-  **Vectorized Compilation:** Speeds processor throughput with lower hardware power consumption.
-  **Green Design Standards:** Installs strict development metrics directly inside CI/CD pipelines.



Autonomous Energy Agents



Dynamic GPU Throttling

Intelligent agent overlays adjust individual clock speeds on inactive or low-priority inference pipelines instantly.



Predictive Dispatch

Balances ongoing server request arrays based directly on local grid cost metrics and thermal thresholds.



Automated Sleep

Transitions hardware banks to ultra-low idle sleep patterns during systemic low-traffic evening cycles.

Shift to Smaller Models (SLMs)

Targeted Parameter Efficacy

Scaling down parameters yields exponential power advantages.
Standardized 7B to 8B parameter models achieve up to 90% accuracy on specific tasks at a tiny fraction of the electrical cost.

$$E_{\text{inference}} \propto P^{1.5}$$

Standard LLM (175B Parameters)

100% Energy

Domain SLM (8B Parameters)

Using small language models (SLMs) delivers over 90% electrical usage reduction with minimal degradation in practical domain-specific tasks.

Joint R&D with Next Realm AI



Pioneering Green Inference

Collaborating closely with **Next Realm AI**, Elektros is building future-proof operational standards to dramatically streamline next-generation neural architecture deployment.

- ✓ **Neuromorphic Modeling:** Accelerating analog network emulation to replace standard digital overhead.
- ✓ **Edge Deployment:** Localizing inference onto smaller chips, completely bypassing continuous data-center sync.

Resource & Contact Hub

Resource Domain	Access Information	Department
Investor Portal	https://elektros.energy/investors.html	Investor Relations
Investor Newsletter	Newsletter Signup	Business Development
IR Agent	Visit IR Agent	Investor Support
Corporate Headquarters	16950 North Bay Road, Suite 1803 Sunny Isles Beach, FL 33160	Administrative Operations

Safe Harbor Statement

Forward-Looking Statements

This material contains forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995. These statements involve risks and uncertainties that could cause actual results to differ materially from those expressed or implied, including regulatory developments, execution risks, market conditions, and technological factors.

CONTACT:

Elektros, Inc.

Investor Relations & Media

Email: elektrosinc@gmail.com

Phone: 786-477-9003