The Fedora Model

Advancing AI and Sustainable Grid Solutions

Prepared by:

• James Tanneberger

Date:

• June 14, 2025

Overview:

This document presents groundbreaking solutions to integrate next-generation AI technologies with sustainable grid practices. It highlights key benefits for stakeholders such as rural economies, AI enterprises, and grid operators.

The Fedora Model: Reshaping AI Infrastructure for a Resilient Energy Future

Whitepaper

Executive Summary

The exponential growth of artificial intelligence is straining the traditional infrastructure supporting it—both in terms of compute demand and energy consumption. Meanwhile, electric grids across the U.S. are under pressure from the transition to renewable energy, leading to increased volatility and reduced baseload capacity.

Fedora Strategies LLC introduces a new solution: the Fedora Model. By embedding AI data centers directly near electric cooperative substations and leveraging middle- and last-mile fiber networks, Fedora creates a decentralized, grid-responsive computing architecture that meets the rising demand for AI while helping stabilize the very grid that supports it.

This white paper outlines how this novel model transforms idle grid infrastructure into a national AI platform, generates new revenue streams for cooperatives, and supports clean energy adoption without the costly and disruptive footprint of hyperscale data centers.

1. Introduction: A Grid at the Crossroads, a Cloud Under Pressure

Two disruptive forces are reshaping the digital and energy landscapes at once:

- The AI revolution, which is driving unprecedented demand for computing resources
- The energy transition, which is rapidly increasing the penetration of intermittent renewable power

Modern AI workloads require massive compute capacity. Hyperscale data centers, while efficient at scale, demand gigawatts of energy and often clash with the limitations of local grids, especially in rural and renewable-heavy regions. At the same time, electric cooperatives—which serve 42 million Americans—struggle to monetize their infrastructure in this new energy economy. Many substations operate far below capacity, while their fiber investments remain underleveraged.

The Fedora Model bridges these gaps. It introduces a new paradigm for infrastructure: one that distributes AI processing across the grid rather than concentrating it in overburdened urban nodes. It replaces volatile, spiky load profiles with high-load-factor, dispatchable AI demand, turning liabilities into assets and uncertainty into control.

2. The Problem: Colliding Crises in Energy and AI

2.1 A Grid Under Strain

The U.S. electric grid is undergoing a seismic shift. Aging infrastructure, the retirement of fossil fuel plants, and the accelerated deployment of renewable energy have created a fragile balancing act. The increasing penetration of solar and wind introduces variability into a system designed for predictability, making grid stability harder to achieve.

Grid operators are scrambling to manage peak demand spikes, frequency disruptions, and declining reserve margins—all of which threaten reliability. Traditional solutions like peaker plants or transmission upgrades are expensive, slow to deploy, and often difficult politically.

Result: The grid needs to add new types of base-line load—flexible, dispatchable, and available where infrastructure already exists.

2.2 Hyperscale Headwinds

At the same time, artificial intelligence is catalyzing an explosion of computational demand. Generative models and high-volume inference workloads are driving enterprises to seek ever-larger compute resources—often located in hyperscale data centers that require hundreds of megawatts, access to transmission-grade power, and dedicated water and cooling infrastructure.

This creates three core challenges:

- Location inflexibility: Hyperscalers concentrate in regions with limited grid capacity (e.g., Northern Virginia, Phoenix, Dallas), compounding local stress.
- Grid imbalance: Large, concentrated loads reduce grid flexibility and leave no room for fast response to fluctuations in demand.
- Community resistance: Rural and suburban communities increasingly resist hyperscale citing due to land use, noise, and environmental concerns.

Result: Centralized compute is hitting structural, economic, and political resistance.

2.3 Underutilized Cooperative Infrastructure

Electric cooperatives operate over 8,000 substations across the United States, many of which are significantly underused relative to their transformer capacity. These substations already serve rural and exurban communities with growing fiber footprints and access to renewables—but they lack a scalable economic model to capitalize on their infrastructure.

In parallel, G&T cooperatives and generation partners are seeking new sources of steady demand to justify system investments and increase load factors—a key metric for asset utilization.

Result: The grid has the space and fiber to scale compute, just not the model.

Conclusion: A Rare Opportunity for Convergence

The grid is calling for help, and AI is calling for space. But their trajectories are out of sync. The Fedora Model offers a chance to align these forces—distributing AI while stabilizing the grid and unlocking economic opportunities for rural utilities along the way.

3. The Fedora Model: Grid-Aware Computing, Built for Scale

The Fedora Model is a strategic framework that combines distributed AI compute, gridaware load shaping, and fiber-based workload orchestration—all embedded within the power system's most underutilized asset: the distribution substation.

Rather than build new hyperscale facilities, Fedora Strategies partners with electric cooperatives to deploy modular, dispatchable AI data centers directly at substations, using existing transformer capacity and tapping into brand-new, cooperative-based, last-mile and middle-mile fiber networks. This model creates a decentralized AI infrastructure layer that behaves like a cohesive cloud platform—aggregated, efficient, and grid-integrated.

3.1 Core Principles of the Fedora Model

- Compute as Load-Shaping Infrastructure: Each Fedora site hosts a modular AI data center capable of scaling compute power up or down based on real-time grid conditions. When the grid has surplus power, Fedora's model absorbs it. When supply tightens, Fedora can reduce consumption.
- Substations as Digital Launchpads: Fedora deploys modular data centers on small, 1-acre to 5-acre tracts near existing cooperative substations. These sites are already fiber-connected and often require no substation upgrades, making them ideal for scalable AI deployments without new transmission or extensive land-use complications.
- Fiber Aggregation for Virtual Hyperscale: Using local, regional, and national fiber aggregation, Fedora links dozens—eventually hundreds—of substation compute nodes into a cohesive, low-latency AI network. This creates virtual hyperscale power without the geographic or infrastructure constraints of a centralized cloud.

3.2 What Makes Fedora Different

Feature	Hyperscale Data Centers	Fedora Model
Location	Urban/exurban, land- intensive	Small plots near existing substations
Load Profile	Flat, inflexible demand	Flexible, dispatchable, grid- responsive
Grid Impact	Stresses local infrastructure	Optimizes and stabilizes the grid
Deployment Cost & Speed	High capex, long timelines	Low-friction, modular, faster ROI
Community Acceptance	Resistance due to land, water, sewer, and energy use	Cooperative-driven – requires little land, no ongoing flow of water or sewer, and utilizes existing electric infrastructure.

3.3 Strategic Benefits

- For Co-ops: Unlocks new revenue from existing infrastructure, increases load factor, reduces rate pressure
- For Grid Operators: Provides a controllable load asset, enhances renewable absorption
- For AI Customers: Offers scalable, cost-effective, distributed compute to achieve scale, reduce single-site risk, and grow a more resilient and sustainable AI ecosystem across the US.

4. Benefits by Stakeholder: A Shared Value Framework

The strength of the Fedora Model lies in its ability to create aligned incentives across sectors. By turning idle energy assets into productive infrastructure, Fedora Strategies generates economic, operational, and strategic wins for utilities, AI enterprises, and grid operators alike.

4.1 Electric Cooperatives & G&T Utilities

- Unlocking New Revenue Streams: Substations and excess fiber strands, historically viewed as sunk costs now support value-generating compute hubs.
- Improved Load Management: Modular data centers have backup generators. The generators allow operators to remove the data center from the grid during grid emergencies. These dispatchable loads offer grid-stabilization, reducing AI's load on the grid as required to meet sudden reliability needs.
- Rural Economic Development: Lowers the electric cooperative's overall cost of power, reducing financial strain and allowing the cooperatives to better serve their communities. The Fedora Model brings next-gen technology infrastructure to regions often left out of digital innovation.

4.2 AI & Cloud Enterprises

- Scalable, Distributed Compute: Bypasses siting constraints of hyperscale data centers by leveraging preexisting substations and small plots of land.
- Lower Cost per Kilowatt: Improves overall electric efficiency (improves load factor), greatly reduces land acquisition costs, and has no lengthy permitting timelines.
- Green AI by Design: Aligns with sustainability goals through grid-sensitive dispatch allowing further enabling the use of renewables as base-line resources.

4.3 Grid Operators (ISOs & RTOs)

- A Dispatchable Load Tool: Fedora AI loads behave as flexible demand resources, capable of timely response to emergency grid conditions.
- Improved Resource Planning: Reliable, predictable AI demand enables better planning of generation, transmission, and storage resources.
- Accelerates Clean Energy Integration: With the ability to absorb more base-line generation and drop load at peak periods, Fedora supports the growth of renewables by helping make the grid more stable.

Each stakeholder gains, but the biggest benefit is systemic: The grid gets stronger, Al gets smarter, and rural America gets empowered.

5. Architecture & Technical Overview: Building the Fedora Compute Grid

The Fedora Model isn't just a strategic framework—it's a blueprint for real-world infrastructure that integrates digital computing with physical energy systems. This section outlines how Fedora's modular components combine to create a distributed, responsive, and fiber-orchestrated AI network rooted in substation infrastructure.

5.1 Modular Data Centers at the Edge

Each Fedora node is a containerized or modular microdata center, deployed directly at or near cooperative substations. These facilities include:

- GPU-accelerated servers optimized for AI training or inference
- Edge cooling systems (two-phase, liquid or high-efficiency air-cooled)
- Dedicated power distribution hardware, backed by transformer capacity
- Direct integration with substation telemetry for real-time grid coordination

These compute pods can scale from 2 MW to 20+ MW per site, depending on available capacity.

5.2 Dispatchable Load Control Layer

At the heart of Fedora's uniqueness the creation of load orchestration opportunities, which align computing operations with grid conditions. The model empowers:

- Real-time grid telemetry monitoring (frequency, voltage, demand)
- Al workload shaping algorithms, which prioritize or defer tasks based on conditions
- Integration with utility SCADA and DERMS platforms for dispatch signals
- Support for both autonomous response and grid-operator coordination

In its most simple form, the Fedora model allows grid operators to utilize the data center's backup generation resources to shed load from the grid at critical periods without interrupting AI compute cycles. As it grows in complexity, the Fedora model allows data centers to behave more like flexible generators in reverse—ramping power draw up or down on command, without sacrificing AI task completion.

5.3 Fiber-Linked Compute Clusters

Each local compute pod connects to a regional mesh of fiber routes, enabling workload distribution across multiple substations. Fedora leverages:

- Middle-mile and last-mile fiber operated by cooperatives, municipalities, or private providers
- Distributed workload orchestration using custom software or container-based schedulers (e.g., Kubernetes variants)
- Low-latency communication to route AI jobs dynamically across the grid-aligned network

This approach provides the benefits of hyperscale cloud—without centralized fragility.

5.4 Site Security, Redundancy & Monitoring

Every Fedora site is designed with:

- Physical security via fenced enclosures, surveillance, and remote access controls
- Operational redundancy, including UPS backup, generator failover, and faulttolerant networking
- Environmental & performance telemetry, tracked centrally for uptime and optimization

Security and reliability are built into the fabric, not bolted on after the fact.

6. Implementation Framework: From Pilot to National Platform

6.1 Stage 1 – Targeted Pilots with Electric Cooperatives

The Fedora Model begins with one or more strategic pilot deployments at electric cooperative substations in Indiana. These pilots serve as both technical proof points and relationship-building platforms with key stakeholders.

Key Activities:

- Partner with a forward-thinking cooperative (e.g., SCI REMC or similar)
- Install multiple small-scale AI data modules (2 MW) at multiple substations to achieve a 10 30 MW "cluster" of AI computing capacity



- Integrate with existing fiber routes and SCADA systems
- Use pilot workloads to demonstrate dispatchable load functionality and compute performance

Goals:

- Validate real-time grid alignment
- Collect operational data for optimization
- Establish economic case for cooperative partners
- Refine regulatory positioning and local engagement protocols

6.2 Stage 2 – Regional Expansion and Fiber Aggregation

Once pilot viability is confirmed, Fedora begins to scale within a defined service region. This phase leverages:

• Local fiber networks and middle-mile fiber consortiums like Accord Telecommunications which are made up of partnering electric cooperatives and telecom providers



• Clustering of multiple substation compute nodes into a regional cloud utilizing proximity to minimize latency



• Expansion to adjacent co-ops and utility districts through joint participation models creating larger aggregated clusters.



Indiana - 750 MW Cluster Example

Deliverables:

- Full workload orchestration between sites
- Cooperative revenue-sharing models
- Utility-side controls for load dispatch and fault response
- Optional integration with regional Independent System Operators (ISOs)

6.3 Stage 3 – Grid-Integrated AI at National Scale

With regional success, the Fedora model then scales nationally deploying a continental network of AI modules (a cluster of clusters) tuned to utility topology and ready to meet grid reliability needs.

Core Systems:

- A national operations center managing AI workload routing, grid integration, and maintenance analytics
- Real-time energy market integration via participation in demand response, frequency regulation, and capacity markets
- Distributed compute environments that respond to carbon signals, marginal pricing, or load curtailment directives



Tapestry – 10 Gigawatt Cluster of Clusters Example

Outcome:

- A federated, resilient, and grid-beneficial AI infrastructure spanning the U.S.
- The creation of a national "Compute Grid" that supports the electric grid—balancing Al productivity with electric grid reliability

7. Financial & Economic Rationale: Aligning Capital with Cooperative Intelligence

7.1 Capital Efficiency Through Infrastructure Reuse

The Fedora Model dramatically reduces timelines and capital requirements when compared to hyperscale development by leveraging existing cooperative substations and fiber infrastructure.

Element	Hyperscale Build	Fedora Deployment
Land acquisition	\$1M-\$10M+	Minimal
Transmission interconnect	Multi-million-dollar upgrades	\$0 (distribution-level connectivity)
Permitting & zoning	12–24 months	< 3 months with cooperative support
Total time to deploy	2–5 years	6–11 months

This efficiency allows for faster deployment of AI, faster ROI, lower upfront investment, and scalable replication across multiple cooperative service territories.

7.2 Financial Benefits for Cooperatives

The Fedora Model benefits electric cooperative partners through:

- Site lease payments or shared savings agreements
- Energy consumption revenue from increased substation utilization
- Lower average cost of power due to improved load factor
- Participation in demand response programs or load aggregation markets

8. Vision for the Future: A New Layer of National Infrastructure

8.1 From Infrastructure to Intelligence

For over a century, America's electric grid has been the backbone of productivity and growth. Today, intelligence infrastructure—AI, cloud, and data analytics—is emerging as the new utility layer. Yet without thoughtful integration, this new layer threatens to destabilize the very system that powers it.

The Fedora Model envisions a future where compute and energy are no longer in tension but in partnership. Where the nation's substations don't just deliver electrons, but also host intelligence, enable grid balancing, and become anchors for digital equity in every ZIP code.

8.2 Ethical, Equitable, and Efficient AI at Scale

As the global AI arms race accelerates, so does concern over its energy footprint, access inequality, and geopolitical concentration. The Fedora Model offers a path forward—one rooted in the cooperative ethos, regional diversity, and ecological mindfulness.

By aligning AI growth with rural infrastructure, Fedora:

- Democratizes access to compute beyond cloud incumbents
- Reduces emissions by design, not retrofit
- Creates economic dignity through partnerships, not displacement

This is not just sustainable compute. It's accountable intelligence.

8.3 A National Compute Grid, Inspired by the Electric Grid

The ultimate vision? A nationally coordinated, locally operated mesh of AI infrastructure, integrated with utility controls and optimized by carbon-aware routing. Just as the electric grid revolutionized the 20th century, a Compute Grid—anchored by Fedora—can empower the 21st.

In this future:

- Every watt counts.
- Every workload flexes.
- Every substation matters.

Conclusion:

Fedora Strategies LLC isn't just deploying servers. It's building alignment—between electrons and algorithms, between rural potential and global needs, between legacy systems and future intelligence.

The opportunity is here. The grid is ready. The workloads are waiting. It's time to reimagine what infrastructure can be.