From Trader-Centered Workflows to Trading Systems: A Manifesto for the Next Era in Power Markets

Anguel Grigorov

July 15, 2025

Contents

1 Executive Summary **2**

2 Introduction: The Bottleneck in Virtual Power Trading **2**

3 Industry Responses—and Why They Fall Short **3**

4 The Core Structural Problem **4**

5 The Theory of Inventive Problem Solving (TRIZ) **4**
5.1 Three Core Concepts of TRIZ 4
5.2 Applying TRIZ to Virtual Power Trading 5
5.3 Summary of Impact 6

6 Trading Strategy: The New Unit of Decision-Making **6**
6.1 What Is a Trading Strategy? 6
6.2 Why Strategies? 7
6.3 Hyper-Parameters Define a Strategy 7
6.4 Evaluating a Strategy: Return vs. Ruin 7

7 From Backtest to Probability Theory **8**
7.1 Optimizing a Strategy: Historical-Forward Simulation 8

8 VPDaaS: Strategy-Centered Trading as a Service **8**
8.1 The Emergence of Market Meta-Knowledge 9

9 Key Takeaways **10**

10 Contacts **11**

1 Executive Summary

The virtual power market is outgrowing the capabilities of the human mind. With over 300,000 hub/hour positions available daily, a probabilistic structure, and a landscape of spatially entangled regime shifts, today's trader-centered workflow has become a bottleneck. Burnout, blind spots, and missed opportunities aren't anomalies—they're structural inevitabilities.

This white paper outlines a systemic shift: **from gut-based micro-decisions to engineered, probabilistically-optimized trading systems**. Drawing on TRIZ (the Teory of Inventive Problem Solving), probabilistic modeling, and forward-world simulation, we lay the intellectual and architectural groundwork for what comes next.

Enter VPDaaS — **Virtual Power Desk as a Service**: A new layer of infrastructure for power markets. Think **"the Bloomberg Terminal for probabilistic power trading"—but strategy-native, simulation-driven, and AI-assisted from the ground up**.

VPDaaS enables market participants to:

- **Design and deploy trading strategies** grounded in probabilistic forecasts,
- **Calibrate risk per position**, per day, per strategy,
- **Simulate possible worlds** to optimize returns while bounding ruin risk,
- **Scale trading insight** far beyond the cognitive limits of human decision-makers.

This is not about automating traders out of the loop. It's about **elevating the trader**—from being a fatigued decision machine to a designer of capital-efficient, risk-aware trading systems.

The transition is coming—with or without you. VPDaaS is building the rails for it.

2 Introduction: The Bottleneck in Virtual Power Trading

The modern virtual power trading desk faces a stark contradiction.

Markets are more complex than ever—over **300,000 daily trading positions** span **12,500+ nodes**—yet the workflow remains trapped in **trader-centered micro-decisions**. Every day, human traders are asked to decide:

- Which hub/hour positions to trade,
- In which direction (buy or sell), and
- How many MWh to allocate—hundreds of times, under pressure, and with limited visibility.

This paper focuses exclusively on the virtual power trading space—where speculative trades are placed in day-ahead markets. While similar principles may eventually apply to physical power or Financial Transmission Rights (FTR), the observations and solutions discussed here are grounded in the realities of virtual markets.

This model is outdated.

First, it misallocates scarce human capital. Humans are brilliant at **conceptual strategy design**, not endless, granular execution. The result? **Burnout and high turnover**.

Second, cognitive bottlenecks mean that even the most talented traders can engage with only a **thin sliver** of market opportunities. Vast swaths of the market remain *terra incognita*.

Third, power markets are structured by probabilistic, interacting patterns—not linear, deterministic behaviors. One particularly punishing pattern is what we call a **Flip-Over**: a sudden, systemic regime shift in the directional logic of market zones.

Flip-Overs are the power market's equivalent of equity market "corrections": disruptive, often sharp reversals that reset the rules mid-game. But unlike stock corrections, Flip-Overs are **spatially nonuniform, temporally asymmetric**, and far harder to model—making them even more dangerous to unprepared traders.

They tend to strike every 30—120 days, persist for several days, and routinely catch human traders off guard—often triggering **compounded losses** in the first 24—48 hours.

3 Industry Responses—and Why They Fall Short

The industry has attempted several workarounds—none sufficient:

1. **The Myth of the Super-Trader** Some shops chase the elusive "market whisperer" who *just knows*. Even if found (rarely), these individuals:
   - Can only track a few dozen positions,
   - Demand unsustainably high compensation, and
   - Often prefer to trade independently, not work for firms.
2. **More Data $\neq$ Better Decisions** Firms are awash in data—PDF reports, weather forecasts, price histories, system topology updates. But **data alone is not a decision-making framework**. It lacks:
   - Risk-measurement systems,
   - Risk-management frameworks,
   - Translation mechanisms from forecasts to actionable sizing.
3. **Quant Teams as Partial Solutions** Large players deploy internal quant teams. This helps—but doesn't solve the whole problem.
   - They're **expensive**.
   - They often focus **heavily on forecasting**—not risk modeling, capital sizing, or strategy engineering.
   - They often lack integration into **systematic, full-stack optimization workflows**.

- Without a unifying framework, quant teams can fall into a dangerous pattern: torturing the data until it "confesses"—producing spurious results to satisfy internal pressure. The result? Overfitting masquerading as insight.

4 The Core Structural Problem

Most desks rely on **Value-at-Risk (VaR)**—a 1990s-era tool invented by J.P. Morgan—for risk management. This framework is:

- **Too high-level**,
- **Statistically smoothed**, and
- **Inadequate for per-position risk modeling in a probabilistic** market.

Traders today operate without:

- A framework for **per-position risk profiling**,
- A way to simulate and **test strategies in alternate possible worlds**, and
- A system to align **capital allocation with probabilistic outcomes** and the trader's individual *ruin threshold*.

5 The Theory of Inventive Problem Solving (TRIZ)

In the 1940s, while imprisoned in a Stalin-era labor camp, Soviet engineer Genrih Al'tshuller developed the **Theory of Inventive Problem Solving**, known by its Russian acronym **TRIZ**. His radical idea: invention isn't magic—it follows **discernible, repeatable patterns**.

Drawing from thousands of patents, Al'tshuller distilled universal principles behind breakthrough innovations. He even formalized them into **ARIZ** (Algorithm for Inventive Problem Solving), a method widely adopted by Soviet engineers during the Cold War.

5.1 Three Core Concepts of TRIZ

1. **Ideal Final Result (IFR)** The perfect solution **appears only when needed—and vanishes when not**.

   Example: An **ice cream cone** eliminates the need for packaging. You eat the container.

   The IFR is rarely attainable—but it sets the **directional ideal** for system evolution.

2. **Technical Contradiction (TC)** A TC arises when improving one feature worsens another.

   Example: Making a turbine blade thinner boosts efficiency—but reduces durability.

   Al'tshuller's breakthrough: **innovation is the resolution of contradiction**, not compromise.

3. **Laws of System Evolution** Al'tshuller proposed eight laws governing how all technical systems evolve. Three are especially relevant here:

   - **Law of Increasing Ideality**: Systems evolve to deliver maximum function with minimum cost—eventually becoming invisible.
   - **Law of Transition to Super-System**: When a system hits performance limits, innovation often emerges from a **higher-level system** that includes the original.
   - **Law of Micro-Level Refinement: Major progress often comes by re-engineering low-level components**—the small, overlooked levers of change.

5.2 Applying TRIZ to Virtual Power Trading

Let's apply TRIZ to the core contradiction in today's virtual power trading landscape.

The Ideal Final Result (IFR):

All trader functions—forecasting, sizing, risk assessment—are performed automatically, precisely, and probabilistically, without requiring constant human intervention.

This isn't science fiction. **Algorithmic trading** and **AI-driven strategy design** already move us closer to that ideal.

The Technical Contradiction (TC):

Virtual power trading demands that humans make granular, high-stakes decisions across massive complexity.

But the human brain can't scale to the task. The result?

Burnout, blind spots, and misaligned risk.

TRIZ shows a path forward:

1. **Abandon the trader-centered workflow**, which merely manages—but does not resolve—the contradiction. It's an evolutionary dead end.

2. **Elevate to a higher-order system**: A **strategy-centered framework** where humans define and refine logic, rules, and risk tolerances—but no longer micromanage execution.

   - Humans don't toggle 300,000 levers—they define strategy families and capital thresholds.
   - They focus on **risk/return trade-offs**, not discrete buy/sell decisions.

3. **Re-engineer microstructures**: This includes rethinking how we:

    - Move from VaR to **per-position probabilistic risk metrics**.
    - Link forecasts and risk values directly to **position sizing logic**.
    - Simulate entire portfolios across **alternative futures**, not just historical backtests.

5.3 Summary of Impact

TRIZ frames the problem not as a **talent shortage** or **data challenge**, but as a **systemic contradiction**.

The solution isn't "more dashboards". It's **a new system architecture**.

The trader-centered era is over.

The next one is **systemic, probabilistic, and strategy-driven**.

6 Trading Strategy: The New Unit of Decision-Making

The implication is clear: sooner rather than later, virtual power trading will move beyond its inefficient, trader-centered workflow and adopt a streamlined, super-system-based, strategy-centered one. Human cognition is far better suited to designing and overseeing high-level systems than to executing hundreds of granular, high-stakes decisions under pressure.

This evolution centers on the **trading strategy**—a new unit of decision-making designed to replace human micromanagement with automated, rule-based execution.

6.1 What Is a Trading Strategy?

In the context of virtual power trading, a **trading strategy** is a highly structured, optimized, and nearly automated framework for deciding:

- Which hub/hour positions to trade,
- In which direction (buy/sell), and
- How much MWh to allocate.

In practice, this involves extensive use of AI/ML forecasting, per-position risk modeling, and parameterized logic. The "nearly" in "nearly automated" reflects the ongoing need for human oversight—to intervene in edge cases, debug issues, or apply judgment during *force majeure* events.

This is standard in algorithmic equity or futures trading—and it is high time virtual power markets caught up.

6.2 Why Strategies?

Moving from gut-driven decisions to strategy-centered design unlocks **major advantages**:

- **Cognitive leverage:** Human designers focus on what they do best—engineering risk-return tradeoffs, not chasing hourly volatility.
- **Scalability:** One person can design and oversee multiple strategies simultaneously—across hundreds of thousands of hub/hour positions.
- **Consistency:** Each strategy is tested, benchmarked, and governed by formal criteria.
- **Forward-looking optimization:** Strategies are tuned using historical-forward simulations, not retrofitted with hindsight-based backtests.
- **Accountability:** Performance becomes measurable, repeatable, and explainable.

A strategy is like a symphony played from sheet music—precise, disciplined, and replicable.

Gut-based trading is more like jazz: sometimes brilliant, often erratic, and hard to scale.

6.3 Hyper-Parameters Define a Strategy

Each trading strategy is parameterized by a structured list of **hyper-parameters**, such as:

- **Definition of ruin:** e.g. a 25% drawdown threshold.
- **Per-position risk model:** how much money is at risk if a forecast fails.
- **Hub/hour selection logic.**
- **Directional logic:** what triggers a buy or sell signal.
- **MWh sizing logic** per position.
- **Portfolio heat:** the level of capital risked relative to available funds.
- **Adaptive scaling rules:** how position size responds to P&L trends.
- **Forecasting engine(s)** used for signal generation.
- And more.

6.4 Evaluating a Strategy: Return vs. Ruin

Each strategy is characterized by two primary outcome metrics:

- **Expected Annualized Return (EAR)** — e.g., +15%/year.
- **Estimated Probability of Ruin (PoR)** — e.g., 0.7%.

7 From Backtest to Probability Theory

Backtests are single-shot simulations: one historical path, one outcome. They're prone to overfitting and false confidence.

Historical-forward simulation, by contrast, builds an ensemble of alternative market histories that *could have happened but didn't*.

This transforms strategy evaluation from anecdote into probabilistic science.

7.1 Optimizing a Strategy: Historical-Forward Simulation

Trading is about exposing capital to risk in pursuit of return. The ideal strategy:

- Assumes just enough risk to target a strong return,
- But not so much that the chance of ruin becomes unacceptable.

To find that sweet spot, we simulate performance across **many possible worlds**:

1. Historical price data is used to infer probability distributions for forecasted returns.
2. These are used to generate thousands (10K—1M+) of alternate market scenarios.
3. The strategy is run over each simulated world to assess its behavior.
4. The ensemble of outcomes gives us:
    - A robust estimate of **Expected Annualized Return**, and
    - A robust estimate of **Probability of Ruin**.

We then **optimize hyper-parameters** to target the most desirable return—risk profile—*all while respecting the time-bound nature of forecasts* (i.e., using only what would've been known at each historical point).

8 VPDaaS: Strategy-Centered Trading as a Service

The culmination of this shift—from trader-centered micro-decisions to system-centered probabilistic design—is **VPDaaS: Virtual Power Desk as a Service**.

VPDaaS is not a dashboard. It's not a "data platform." It's a **full-stack strategic engine** built for the unique demands of virtual power trading. Its goal is to democratize what only elite quant teams and algorithmic hedge funds could previously afford: a modular, scalable, AI-native system that designs, tests, and monitors trading strategies—at the individual trader level.

At its core, **VPDaaS builds and maintains trading strategies as probabilistic objects**, not gut-feeling heuristics. Each strategy is defined by:

- A clear set of **hyperparameters** (e.g., risk allocation rules, signal filters, trade sizing logic),

- An embedded **forecasting engine** (e.g., ML models, statistical ensembles),
- A **per-position risk model** rooted in the trader's ruin threshold,
- And a **simulation framework** that stress-tests performance over thousands of alternative historical-forward scenarios.

VPDaaS delivers for each strategy:

- An **Expected Annualized Return**,
- A **Probability of Ruin** estimate,
- And daily **position-level execution guidance** consistent with those strategic bounds.

It also offers:

- Live tracking of strategy performance vs expectations,
- Scalable oversight across portfolios of strategies,
- And rapid re-optimization tools when market conditions evolve.

In short: VPDaaS lets traders focus on what humans do best—defining the *goals*, the *boundaries*, and the *trade-offs*. It handles the rest: *the simulation, the sizing, the execution logic, the monitoring*. And because it's delivered "as a service", it's accessible—no in-house quant team required.

Just as CRMs redefined how businesses manage relationships, and ERPs transformed operations, **VPDaaS will redefine how power traders design and execute decisions**. It is not a tool for the trader. It is the *system* the trader operates from.

8.1 The Emergence of Market Meta-Knowledge

Every trading strategy evaluated in VPDaaS becomes a data point—not just in isolation, but as part of an evolving cloud of knowledge. Each historical-forward simulation (like our current 45,000+ $\times$ 10,000 matrix of tests across MISO hubs) generates a statistical fingerprint of strategy behavior across probabilistic futures.

These fingerprints aren't just optimization inputs—they're the foundation for meta-learning: the science of learning about strategies, not just within them.

As this corpus grows, VPDaaS will become more than a strategy design tool. It will evolve into a keeper of meta-knowledge for the virtual power market:

- Clustering strategies by behavioral archetypes,
- Mapping regions of high robustness vs fragility in strategy-space,
- Recommending strategy templates for new participants based on past analogues.

This is a long-term vision: to codify the epistemology of power trading itself.

In time, VPDaaS won't just help traders design strategies. It will help them ask better questions—based on everything that has ever been tested.

Conclusion: The Future of Power Trading Is Strategic

The contradiction at the heart of virtual power trading is no longer ignorable.

Human-centered micro-decision workflows—no matter how seasoned the traders, no matter how rich the data—are structurally misaligned with the scale, complexity, and probabilistic nature of today's power markets.

What's needed is not more dashboards, more PDFs, or more intuition. What's needed is a shift in the *unit of reasoning* itself:

From **trader actions** to **trading strategies**.

This paper has outlined the theoretical rationale behind that shift. From the lens of TRIZ, we see the trader-centered workflow as a system at the end of its evolutionary arc.

The natural resolution is a strategy-based, probabilistically informed super-system—capable of simulating risk, aligning capital, and letting human minds operate at their cognitive best.

Virtual Power Desk as a Service (VPDaaS) is our answer to that resolution.

VPDaaS embodies the next era in power trading: scalable, auditable, risk-aware, and strategy-centric. By offering modular services—forecasting, per-position risk modeling, strategy simulation, and real-time monitoring—we aim to democratize quantitative power trading and make its benefits accessible to far more players than before.

The goal is not to replace the human trader.

It's to **free the human trader**—to think strategically, to allocate capital intelligently, and to operate with clarity in the face of uncertainty.

The shift is coming. The super-system is forming.

This white paper is its early blueprint.

9 Key Takeaways

- **The Trader-Centered Workflow Is Obsolete**

  Virtual power markets have outgrown human-scale decision-making. With over 300,000 positions daily and probabilistic regime shifts, manual workflows are a structural bottleneck.

- **Strategy Is the New Unit of Reasoning**

  Probabilistic, parameterized trading strategies—optimized through historical-forward simulation across possible worlds—offer a scalable, risk-aware, and cognitively efficient alternative.

- **Flip-Overs Redefine Risk**

  Spatially asymmetric regime shifts ("Flip-Overs") regularly destabilize trading logic. VPDaaS embeds them into its modeling core—so they don't remain unmodeled surprises.

- **VPDaaS Builds the Infrastructure of the Next Era**

  Think "Bloomberg Terminal for probabilistic trading"—a platform to design, optimize, and monitor virtual power strategies with per-position risk, expected return, and ruin probabilities built in.

- **No Quant Team? No Problem.**

  VPDaaS delivers hedge-fund-grade infrastructure as a service—making advanced strategy design accessible to desks that lack in-house quant capabilities.

Quant Power, Democratized.

Further Reading

For readers interested in exploring the foundational concepts behind this paper, the following resources are highly recommended:

- **Theory of Inventive Problem Solving (TRIZ):**
  Altshuller, G. (1999). *The Innovation Algorithm: TRIZ, Systematic Innovation and Technical Creativity*.
  Amazon Link

- **Position Sizing and Risk Management:**
  Tharp, V. K. (2008). *The Definitive Guide to Position Sizing*.
  Van K. Tharp International

- **Probabilistic Forecasting:**
  Gneiting, T., & Katzfuss, M. (2014). Probabilistic Forecasting. *Annual Review of Statistics and Its Application, 1*, 125-151.
  DOI Link

- **Risk Metrics Beyond VaR:**
  McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools*.
  Princeton University Press

10 Contacts

Feel free to reach out for inquiries or collaborations.

Author: Anguel Grigorov, Founder @ VPDaaS

Email: anguel@vpdaas.com

Email: anguel@energymarketai.com

LinkedIn: My LinkedIn page

Blog URL: My Blog page

Phone: 936-256-0830

X: @vpdaas

Calendly: To schedule a call

Copyright and Licensing

Email: anguel@vpdaas.com