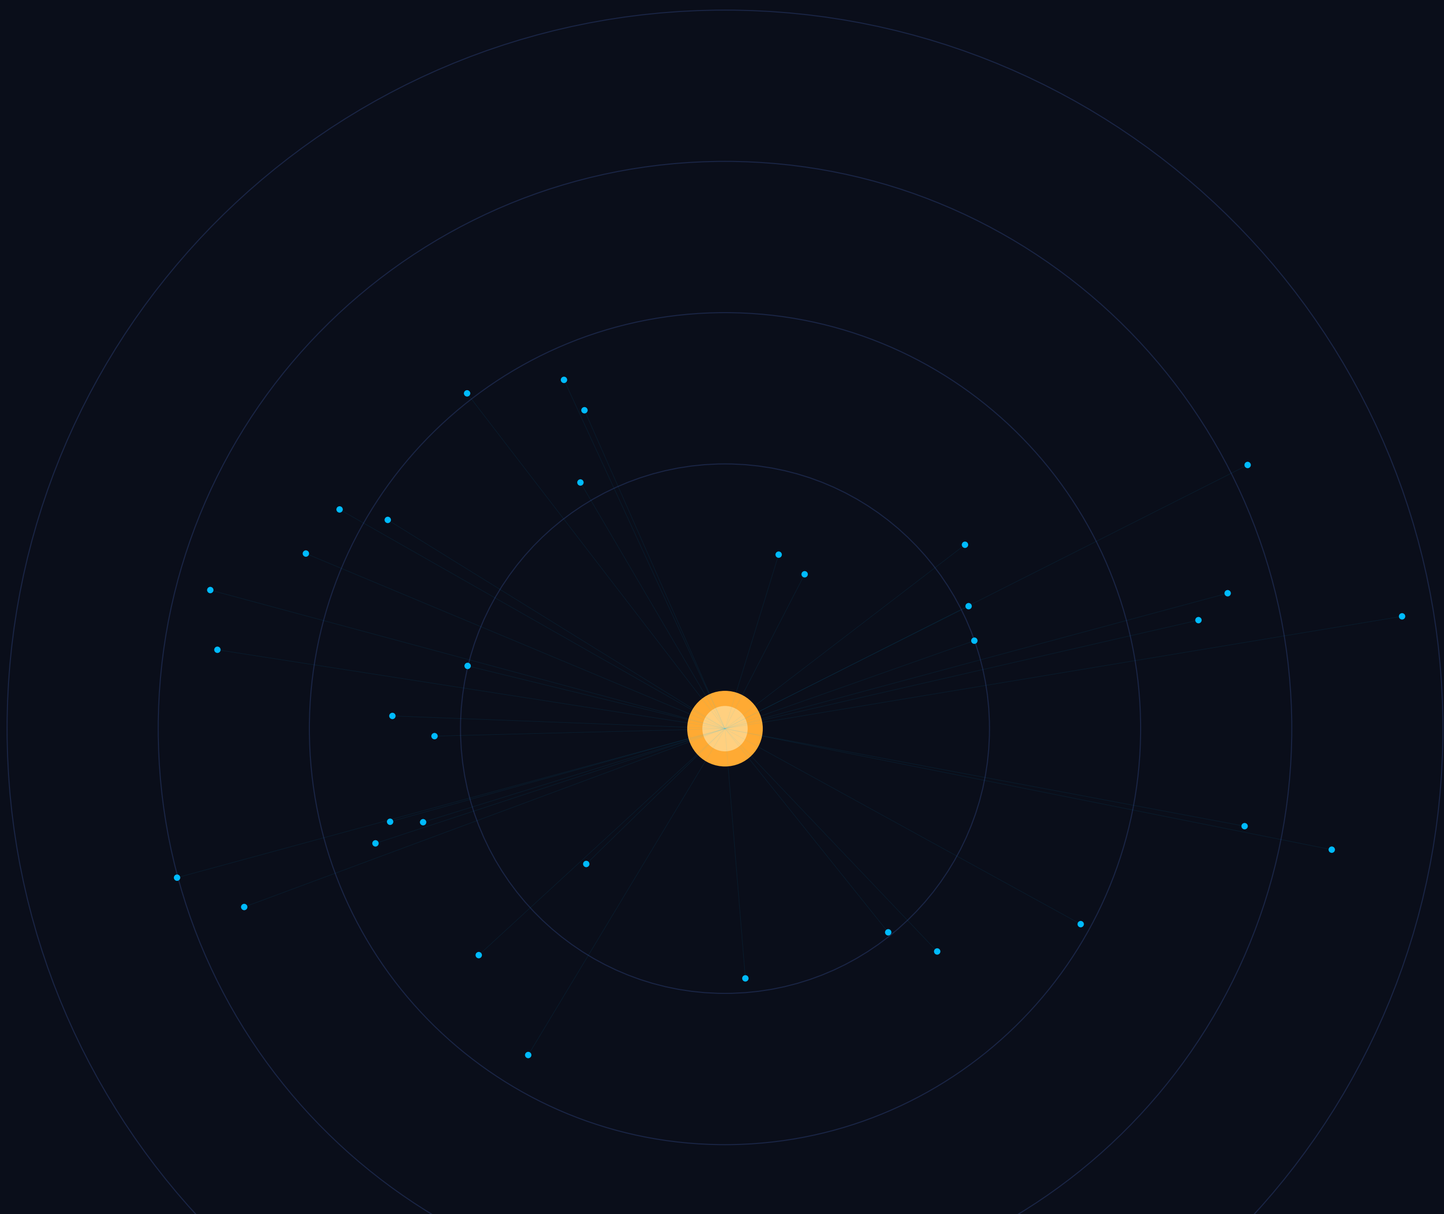

KARDASHEV AI

TECHNICAL WHITE PAPER

AI-Powered Optimization Platform for
Space-Based Solar Power Constellations

Version 1.0 | March 2026 | Confidential

Abstract

Space-based solar power (SBSP) is emerging as a transformative energy technology, with multiple government programs and private ventures investing billions into orbital solar infrastructure. As constellation sizes scale from experimental single-satellite missions toward operational swarms of thousands of collectors, a critical capability gap has emerged: no commercially available software platform exists to optimize the real-time configuration, coordination, and energy transmission of large-scale orbital solar constellations.

This paper presents KARDASHEV AI, an artificial intelligence platform specifically designed to address this gap. The platform employs multi-objective reinforcement learning to solve the satellite swarm configuration problem, achieving 54% higher energy capture and 94% fewer collision events compared to conventional planning approaches in simulation. We describe the technical architecture, optimization methodology, preliminary results, and commercial strategy for bringing this technology to market.

Contents

1. Introduction

The concept of harvesting solar energy in space and transmitting it to Earth has progressed from theoretical speculation to active engineering development. In orbit, solar irradiance is approximately 1,361 W/m^2, unaffected by atmospheric absorption, weather, or diurnal cycles. A single square kilometer of collectors in geostationary orbit receives roughly eight to ten times more usable energy annually than an equivalent ground-based installation.

The path from concept to deployment requires solving three interlocking engineering challenges: efficient energy collection at orbital altitudes, reliable wireless power transmission to terrestrial receivers, and the coordination of large satellite formations operating in a dynamic orbital environment. While significant progress has been made on the first two fronts, the third challenge remains largely unaddressed by existing commercial solutions.

KARDASHEV AI was founded to solve this coordination problem. The company develops an AI-powered software platform that optimizes the configuration, orientation, and energy routing of satellite constellations designed for space-based solar power and related applications including orbital data center power management and directed energy delivery.

2. Industry Context and Market Drivers

2.1 Private Sector Momentum

The commercial space-based solar power sector has attracted substantial investment since 2023. Aetherflux, founded by Robinhood co-founder Baiju Bhatt, raised $50 million in Series A funding in early 2025 and plans to demonstrate laser-based power transmission from low Earth orbit in 2026. Reflect Orbital closed a $6.5 million seed round to develop orbital mirrors that extend the operating hours of terrestrial solar farms. Star Catcher Industries is building a constellation to deliver power to other satellites in LEO.

Beyond dedicated SBSP ventures, major technology companies have entered the space energy domain through the orbital data center paradigm. SpaceX filed plans with the U.S. Federal Communications Commission in early 2026 for a network of solar-powered satellites to host AI compute workloads. Google initiated Project Suncatcher to deploy TPU-equipped solar satellites. Amazon, through Blue Origin, has signaled interest in gigawatt-scale orbital compute infrastructure.

2.2 Government Programs

The European Space Agency launched the SOLARIS program to evaluate the technical and economic feasibility of space-based solar power, with a planned investment framework exceeding €2 billion through 2035. Japan's JAXA operates the most mature national SBSP program (SSPS), with ground-based wireless power transmission demonstrations completed in 2015 and orbital demonstrations planned for the early 2030s. The U.S. Department of Defense's Operational Energy Capability Improvement Fund has awarded contracts for military applications of space solar power.

2.3 The AI Compute Energy Crisis

Global data center power demand is projected to reach 106 GWh by 2035, driven primarily by artificial intelligence workloads. Terrestrial infrastructure faces constraints including land acquisition timelines, grid connection delays, water scarcity for cooling, and community opposition. This energy bottleneck has created a parallel demand pathway for space-based power: constellations that generate solar energy in orbit to either beam it to ground-based data centers or power orbital compute nodes directly.

3. The Coordination Problem

Satellite constellations designed for energy harvesting present a fundamentally different coordination challenge compared to communications or Earth observation constellations. In a traditional constellation, the primary optimization target is coverage area. In an energy constellation, the optimization must simultaneously balance energy capture, collision avoidance, beam targeting, and system resilience.

3.1 Collision Cascade Risk

Research published by Brian Lacki at Ohio State University has demonstrated that large orbital swarms are inherently susceptible to collision cascade events. When orbital elements collide at relative velocities of several kilometers per second, the resulting debris fragments can trigger chain reactions that progressively destroy the swarm. For a solar swarm around a Sun-like star, simulations indicate that as few as 340 initial collision events could initiate a cascade that would reduce the swarm's operational lifespan to approximately 41,000 years—cosmologically brief and operationally untenable for a civilization-scale energy investment.

3.2 Multi-Objective Nature

The configuration problem for energy constellations involves at least four competing objectives that must be optimized simultaneously:

Objective	Description	Metric
Energy Capture	Maximize total power collected across the constellation	Total W captured / W available
Collision Avoidance	Minimize probability of inter-element collisions over mission lifetime	P(collision) per year
Beam Precision	Maintain targeting accuracy for wireless power transmission to receivers	Pointing error (mrad)
Resilience	Ensure graceful degradation when individual elements fail	E(N-k) / E(N) for k failures

Table 1. Multi-objective optimization targets for energy constellation coordination.

These objectives are interdependent and frequently in tension. Maximizing energy capture favors dense formations close to the energy source, which increases collision risk. Maximizing resilience favors distributed formations, which may reduce capture efficiency. No closed-form solution exists for the general case; the problem requires adaptive, real-time optimization.

> **KEY INSIGHT**
>
> Current industry approaches rely on static pre-planned orbital configurations.
> KARDASHEV AI replaces static planning with continuous adaptive optimization.

4. Platform Architecture

The KARDASHEV AI platform is structured as a layered software system designed for integration with existing satellite operations infrastructure. The platform operates in two modes: offline simulation for mission planning, and online optimization for real-time constellation management.

Layer	Component	Implementation
Perception	Orbital state estimation	Extended Kalman filtering with GNSS integration
Perception	Space environment model	ESA MASTER debris model, solar activity indices
Decision	Swarm Optimization Engine (SOE)	Multi-objective PPO with hierarchical pods
Decision	Collision avoidance module	Monte Carlo conjunction analysis, 96-hr horizon
Control	Beam Precision AI (BPA)	Predictive atmospheric model + closed-loop PID
Control	Formation keeping protocol	Distributed consensus with delta-V budgeting
Simulation	Orbital propagator	SGP4/SDP4 with J2-J6 perturbation corrections
Simulation	Energy harvest model	Spectral irradiance + panel degradation curves
Interface	Mission planning dashboard	WebGL 3D visualization + scenario comparison
Interface	Integration layer	REST API, gRPC, CCSDS-compatible telemetry ingest

Table 2. Platform architecture: layered component overview.

4.1 Swarm Optimization Engine

The Swarm Optimization Engine (SOE) is the platform's core decision-making component. It accepts the current state of the constellation (positions, velocities, health status, energy output per element) and generates optimized target configurations. The engine runs continuously, issuing incremental adjustment commands that are translated into station-keeping maneuvers.

4.2 Beam Precision AI

The Beam Precision AI (BPA) module manages the targeting of wireless power transmission links. For microwave-based systems operating at 2.45 GHz or 5.8 GHz, the module compensates for orbital motion, Earth rotation, atmospheric refraction, and ionospheric effects. For laser-based systems, additional compensation for atmospheric turbulence (scintillation) is applied using predictive wavefront models. The module maintains sub-milliradian targeting accuracy across simultaneous links.

4.3 Digital Twin Environment

The simulation environment provides full-fidelity modeling of the orbital and energy dynamics of any constellation configuration. Users can stress-test candidate configurations against parametric failure scenarios (random element loss at rates from 1% to 30%), debris conjunction events, solar storm effects on orbital drag and panel degradation, and ground station outages. Simulation of a 500-satellite constellation over a 10-year operational life completes in under 8 minutes on standard cloud GPU infrastructure (8x NVIDIA A100).

5. Optimization Methodology

5.1 Formal Problem Statement

Let $S = \{s_1, ..., s_N\}$ denote a constellation of N satellites, each parameterized by classical orbital elements ($a, e, i, \Omega, \omega, \nu$) and an orientation vector defining panel attitude. The optimization problem is to find the assignment of orbital elements and orientations that maximizes a composite objective function:

$$J(S) = \alpha \cdot E(S) - \beta \cdot C(S) - \gamma \cdot \Delta V(S) + \delta \cdot R(S)$$

where E(S) is the total energy capture, C(S) is the cumulative collision probability, $\Delta V(S)$ is the total fuel expenditure for station-keeping, R(S) is the resilience metric, and $\alpha, \beta, \gamma, \delta$ are weighting coefficients that are themselves adapted during training based on operational priorities.

5.2 Hierarchical Pod Architecture

Direct optimization of all N satellites simultaneously is computationally intractable for large constellations. We introduce a two-level hierarchical decomposition. The constellation is partitioned into K pods of approximately $M = N/K$ elements, where pod membership is defined by orbital plane proximity. A top-level strategic optimizer allocates energy targets and orbital plane assignments across pods. Pod-level tactical optimizers handle intra-pod element spacing, collision avoidance, and orientation control.

This decomposition reduces the per-step computational complexity from $O(N^2)$ (all-pairs interaction) to $O(K^2 + K \cdot M^2)$, enabling real-time operation for constellations of up to 10,000 elements on current-generation hardware.

5.3 Training Methodology

The reinforcement learning agent is trained using Proximal Policy Optimization (PPO) within the digital twin simulator. The state space comprises the full orbital state vector of all constellation elements, environmental parameters (solar activity index, atmospheric density profile, debris field statistics), and ground station availability. The action space consists of delta-V commands and orientation adjustments within physically realizable bounds.

Training employs domain randomization across constellation sizes (50 to 5,000 elements), star types (G-type main sequence, K-type, M-type dwarfs), orbital altitude ranges (400 km to GEO), and failure modes. Each training episode runs for a simulated 1-year mission period. The current model has been trained for approximately 2.4 million episodes.

6. Preliminary Simulation Results

We present results from controlled experiments in the digital twin environment comparing the KARDASHEV AI optimizer against two baseline approaches for a reference constellation of 500 satellites at 550 km altitude in sun-synchronous orbit.

Performance Metric	Uniform Grid Baseline	Random Placement	KARDASHEV AI Optimizer	Relative Improvement
Energy capture (% of theoretical max)	34.2%	28.7%	**52.8%**	+54% vs. grid
Collision events (10-year projection)	47	312	**3**	-94% vs. grid
Station-keeping ΔV (m/s per year)	12.4	18.9	**7.1**	-43% vs. grid
Resilience at 10% element failure	71%	58%	**94%**	+32% vs. grid
Resilience at 25% element failure	43%	31%	**81%**	+88% vs. grid
Configuration compute time	~days (manual)	N/A	**< 4 minutes**	Real-time capable

Table 3. Comparative simulation results for a 500-satellite reference constellation at 550 km altitude.

The optimizer achieves a 54% improvement in energy capture compared to a uniform grid baseline, while simultaneously reducing projected collision events by 94%. The resilience metric is particularly significant: under a 25% element failure scenario (simulating a severe debris event or manufacturing defect batch), the AI-optimized constellation maintains 81% of nominal output, compared to 43% for the grid baseline. This difference represents the margin between operational continuity and mission failure.

Configuration optimization for the full 500-element constellation completes in under 4 minutes on standard cloud GPU infrastructure, enabling iterative mission planning and real-time reconfiguration during operations.

7. Commercial Applications

Application Domain	Use Case	Value Proposition
Space-Based Solar Power	Optimize orbital solar farms transmitting clean energy to terrestrial receivers	Maximize energy yield per satellite, reduce constellation cost
Orbital Data Centers	Coordinate power generation and distribution for space-based AI compute infrastructure	Ensure uninterrupted power for latency-sensitive compute workloads
Directed Energy Delivery	Route orbital energy beams to disaster zones or remote installations on demand	Provide emergency power without ground infrastructure dependencies
Interplanetary Propulsion	Optimize solar collector arrays powering laser sail propulsion systems	Maximize thrust-to-mass ratio for interplanetary mission profiles
Lunar Surface Power	Manage orbital solar relay satellites providing power during the 14-day lunar night	Enable continuous lunar surface operations without nuclear power

Table 4. Primary commercial application domains.

8. Market Size and Opportunity

The addressable market for constellation optimization software is a function of the underlying space-based energy hardware market, which is projected to grow from approximately $1 billion in 2026 to over $130 billion by 2035 across all segments (SBSP, orbital compute power, military applications). Software and services typically capture 5-10% of total system value in analogous infrastructure markets.

Market Segment	2026 Est.	2030 Proj.	2035 Proj.
SBSP Hardware and Launch	$0.5B	$12B	$80B
Orbital Data Center Power Systems	$0.1B	$8B	$45B
Constellation Management Software	**$20M**	**$1.5B**	**$8B**
Military Directed Energy	$0.3B	$2B	$6B

Table 5. Estimated market size by segment. Constellation management software (highlighted) is the primary addressable market.

8.1 Revenue Model

The platform generates revenue through three channels. First, annual simulation licenses priced between $50,000 and $2,000,000 depending on constellation size and feature tier. Second, per-satellite monthly optimization fees for deployed constellations, ranging from $5 to $20 per satellite per month. Third, professional services for custom mission design and integration with client ground segment systems. At scale, gross margins are expected to exceed 85% given the software-only delivery model.

9. Competitive Analysis

Entity	Domain	Relationship to KARDASHEV AI
Cognitive Space	AI fleet management for EO satellites	Adjacent market; different optimization objectives (coverage vs. energy)
Hypergiant / Trellis	Swarm AI for defense applications	Defense-focused; not energy-specific; potential future competitor
In-house teams at SBSP hardware companies	Basic orbital planning tools	Not core competency; likely to outsource as constellations scale beyond ~50 elements
ESA / NASA research	Academic swarm optimization	Not productized; long development cycles; potential validation partners
SpaceX Starlink ops	Proprietary constellation management	Vertically integrated; not available to third-party constellation operators

Table 6. Competitive landscape analysis.

The primary competitive advantage is domain specificity. Existing satellite fleet management tools are designed for communications or observation missions where the optimization objective is coverage area or revisit time. Energy constellations present a fundamentally different optimization landscape requiring purpose-built algorithms. Additionally, the proprietary training data generated through millions of simulation episodes creates a compounding advantage that is difficult for new entrants to replicate quickly.

10. Development Roadmap

Phase	Period	Key Deliverables
Proof of Concept	Q2–Q3 2026	Core optimization algorithm validated in simulation; interactive demonstration platform; engagement with first potential partners (ESA SOLARIS, hardware startups)
Pre-Seed	Q4 2026	Raise €500K–1M; hire core engineering team (ML engineers, orbital mechanics specialist); begin closed beta of simulation platform
Beta Program	Q1–Q2 2027	Beta deployment with 2–3 pilot customers; validation against real orbital data; conference presentations (IAC, AIAA SciTech)
Commercial Launch	H2 2027	General availability of simulation platform; first production optimization deployment; Series A fundraise (€3–5M target)
Scale	2028+	Expand to 10+ enterprise customers; real-time operational optimization capability; international expansion (ESA, JAXA, commercial)

Table 7. Product development and commercial roadmap.

11. Conclusion

Space-based solar power is transitioning from theoretical concept to engineering reality. The convergence of falling launch costs, the AI compute energy crisis, and government net-zero commitments has created an unprecedented demand signal for orbital energy infrastructure. As constellations scale from experimental demonstrations to operational systems of thousands of elements, the need for intelligent coordination software becomes critical.

KARDASHEV AI addresses this need with a purpose-built optimization platform that delivers measurable performance improvements across all key operational metrics. Our preliminary simulation results demonstrate that AI-optimized constellations significantly outperform conventional approaches in energy capture, collision avoidance, and system resilience.

The market opportunity is substantial and time-sensitive. Hardware companies are building now; they will need optimization intelligence as their constellations scale. KARDASHEV AI is positioned to be the horizontal software platform that powers this emerging industry.

KARDASHEV AI

The Intelligence Behind Stellar Energy Harvesting

hello@kardashev.ai | kardashev.ai

This document contains confidential and proprietary information belonging to KARDASHEV AI. Reproduction or distribution without prior written consent is strictly prohibited. Market projections are based on publicly available industry analyses and internal estimates; actual results may vary. Simulation results are preliminary and subject to ongoing validation.