Satellite Inference

A staged architecture for processing space-originated data near its source

Compute where space data begins.

Satellite Inference proposes a staged orbital computing architecture for processing data that already originates in space. The program begins with a hosted compute pathfinder, advances to a proposed company-controlled node with 1 kW of continuous compute and 10 kW of BOL solar generation, and scales only after customer, power, thermal and flight-evidence gates are passed.

Author: Bekzad Dzhanpolatov

Organization: A program of RFID INC

Version: Public Whitepaper v0.1

Date: 1 September 2026

Evidence cut-off: 1 September 2026

Public concept summary. Proposed reference missions and preliminary analytical assumptions, not flight-release data or commercial commitments.

Contents

How to Read This Document
. Executive Summary
. Problem and Initial Market
. Proposed Product Architecture
. Proposed Reference Mission Ladder
. Orbit Selection and Propagation
. Rev A Power and Thermal Model
. Compute and Data Architecture
. Communications Strategy
. Proposed Mission Success Criteria
. Development Gates
. Key Risks
. Regulatory and Sustainability Framework
. Public Roadmap
. Public Evidence Ledger
. Explicit Unknowns
. Conclusion

How to Read This Document

This public whitepaper describes a proposed architecture and a sequence of reference missions. No spacecraft described here has been represented as an operational product or contracted launch.

The document uses four classes of statements:

- **Public fact:** supported by a linked primary or government source.
- **Rev A calculation:** produced by a transparent screening model for comparison and planning.
- **Working assumption:** selected to make the reference concept testable, but subject to analysis and evidence.
- **Business hypothesis:** requires customer discovery, a defined workload and a measurable acceptance criterion.

Compute power is stated as electrical input power. It is not a claim about FLOPS, TOPS, tokens per second or model throughput. Those metrics require selected hardware and a reproducible workload benchmark.

1. Executive Summary

Satellite Inference is developing orbital computing infrastructure for data that is created in space. The initial goal is not to reproduce a general terrestrial cloud in orbit. The proposed initial service processes optical, hyperspectral, synthetic-aperture radar, civil RF and spectrum analytics, telemetry, and scientific data near the sensor that creates it.

The commercial hypothesis is simple. When useful information can be extracted in orbit, an operator may avoid downlinking every raw byte before acting. Onboard processing can reject unusable observations, detect an event, prioritize a region, compress data, create a higher-level product and transmit a smaller evidence package.

The proposed program uses five reference stages:

1. **Hosted Pathfinder (SI-HP):** a 0.2-1.0 kW compute payload allocation on a third-party spacecraft.
2. **Node 1 kW (SI-N1):** the first proposed company-controlled node, with 1 kW continuous compute and 10 kW BOL solar generation.
3. **Node 10 kW (SI-N10):** a separate proposed node with 10 kW continuous compute and approximately 30 kW BOL solar generation.
4. **Node 100 kW (SI-N100):** a proposed modular industrial-scale node.
5. **Grid 1 MW (SI-G1MW):** a proposed aggregate cluster, not a monolithic spacecraft.

The number in every Node or Grid name refers to continuous compute power. Solar generation is a separate specification.

The v0.1 baseline is low Earth orbit, or LEO, at 500-600 km. GEO and Lagrange points are comparison cases, not baseline deployment commitments.

2. Problem and Initial Market

2.1 Data Already Originates in Space

Modern sensors can create more data than is valuable to transmit in raw form. The opportunity is greatest when the useful result is much smaller than the input and when earlier decisions matter.

Proposed initial workloads include:

- cloud, quality and empty-frame screening for Earth observation;
- event detection and priority selection;
- calibration, compression and higher-level product generation;
- civil RF and spectrum feature extraction;
- scientific instrument preprocessing;
- autonomous health and anomaly analytics for space infrastructure.

2.2 What a Customer Would Receive

A useful orbital data product should include more than a model output. The proposed evidence package includes:

- an input-data identifier;
- model and runtime version;
- timestamp and orbital context;
- confidence and quality flags;
- result hash;
- measured energy and execution time;
- temperature history;
- recovery or corrected-error events.

The intended outcome is a reproducible result that a customer can compare with a ground-processing baseline.

2.3 Initial Customer Hypotheses

Two public-safe customer categories guide the first experiments:

1. A sensor operator that wants faster event detection, reduced downlink volume or a more autonomous workflow.
2. A compute or electronics developer that wants on-orbit characterization and flight evidence for a defined workload.

These remain hypotheses until a customer provides representative data, an interface definition and an acceptance metric.

---

3. Proposed Product Architecture

Satellite Inference proposes an integrated architecture with five functional layers:

- a modular compute payload with memory, storage and telemetry;
- a radiation-aware runtime with monitoring, checkpointing and controlled recovery;
- a power and thermal scheduler that respects energy, temperature and communications limits;
- an interface for ingesting space-originated data and delivering verified results;
- an orbital platform with generation, storage, heat rejection, communications and mission operations.

The first deliverable does not require an owned spacecraft. Hosted Pathfinder can characterize compute behavior on a partner platform. SI-N1 is intended to validate the coupled power, battery, thermal and compute architecture on a company-controlled free-flyer.

3.1 What v0.1 Does Not Claim

Version 0.1 does not claim:

- a terrestrial solar business;
- wireless delivery of power from space to Earth;
- a general-purpose orbital cloud as the initial market;
- GEO or Lagrange operations as the baseline mission;
- a selected processor or promised inference throughput;
- a committed launch, manufacturing configuration or production schedule;
- a partnership with any launch, spacecraft or component provider.

4. Proposed Reference Mission Ladder

Reference stage	Primary objective	Continuous compute	BOL solar generation	Public status
Hosted Pathfinder (SI-HP)	Characterize compute, recovery and a real workload on orbit	0.2-1.0 kW in allocated windows	Host supplied	Proposed first flight step
Node 1 kW (SI-N1)	Validate an integrated power, thermal, battery and compute node	1 kW	10 kW	Proposed first company-controlled node
Node 10 kW (SI-N10)	Provide sustained multi-workload compute capacity	10 kW	approximately 30 kW	Proposed follow-on node
Node 100 kW (SI-N100)	Demonstrate modular industrial scaling	100 kW	approximately 288 kW in Rev A	Long-range reference architecture
Grid 1 MW (SI-G1MW)	Aggregate multiple modules into a cluster	1 MW aggregate	approximately 2.88 MW aggregate in Rev A	Long-range reference architecture

Each stage is a separate decision gate. Later stages are not launch commitments or completed designs.

4.1 Hosted Pathfinder (SI-HP)

A hosted payload is carried and supported by another operator's spacecraft. Under a future interface agreement, the host would provide defined mechanical, power, thermal, data, communications and operations services. Satellite Inference would provide the compute payload, software and experiment definition.

Proposed objectives are to:

- execute a signed representative workload;
- measure performance per watt;
- observe memory, storage and accelerator behavior;
- demonstrate checkpoint and controlled recovery;
- record board and interface temperature;
- process a real space-originated dataset;
- produce an auditable flight-evidence report.

The 0.2-1.0 kW range describes a potential electrical allocation, which may be scheduled rather than continuous.

4.2 Node 1 kW (SI-N1)

SI-N1 is the first proposed company-controlled spacecraft. Its reference objective is to combine:

- 10 kW BOL deployable solar generation;
- 1 kW continuous compute through sunlight and eclipse;
- battery-supported operation during Earth shadow;
- dedicated heat rejection;
- independent TT&C and a mission-data path;
- controlled end-of-mission disposal.

The installed 10 kW solar capability intentionally exceeds the minimum analytical requirement for 1 kW continuous compute. The margin is intended to support deployment validation, degradation, battery recharge, platform loads and controlled experiments. It is not a claim that 10 kW of compute can run continuously on SI-N1.

4.3 Node 10 kW (SI-N10)

SI-N10 is a separate proposed mission, not an operating mode of SI-N1. In the Rev A screening model, 10 kW continuous compute requires approximately:

- 30.0 kW BOL solar generation;
- 26.9 kWh target battery nameplate;
- 33.3 m^2 effective radiator area.

These values are analytical scaling outputs, not released subsystem specifications.

4.4 Node 100 kW and Grid 1 MW

SI-N100 and SI-G1MW are reference architectures for modular scaling. The preferred public hypothesis is repeatable modules rather than one monolithic spacecraft. Capacity would be added only after workload demand, operational evidence, communications and launch interfaces are demonstrated.

5. Orbit Selection and Propagation

5.1 Why LEO Is the Baseline

The v0.1 reference range is 500-600 km LEO, with 550 km used for screening calculations. LEO offers short propagation distance, access to rideshare services, a mature ground-station ecosystem and proximity to many Earth-observation workloads.

Using a circular two-body approximation with Earth equatorial radius $R_E=6,378.137$ km and gravitational parameter $\mu=398,600.4418\ \mathrm{km^3/s^2}$:

$$r=R_E+h$$

$$T=2\pi\sqrt{\frac{r^3}{\mu}}$$

At 550 km, the model produces an orbital period of approximately 95.65 minutes.

5.2 Eclipse

For a spherical Earth and a cylindrical shadow, the maximum eclipse at beta angle zero is:

$$t_{e,max}=\frac{T}{\pi}\arcsin\left(\frac{R_E}{r}\right)$$

At 550 km, the screening result is approximately 35.61 minutes of eclipse and 60.04 minutes of sunlight per orbit. Actual eclipse duration depends on orbit plane, date and beta angle.

5.3 Propagation Floor

The vacuum one-way propagation floor is:

$$t_{one-way}=\frac{\rho}{c}$$

where ρ is link distance and c is the speed of light.

Comparison case	Reference altitude or distance	One-way propagation floor	Interpretation
LEO, overhead at 550 km	550 km	approximately 1.835 ms	Baseline reference mission
GEO, overhead	35,786 km	approximately 119.37 ms	Comparison case, not baseline
Sun-Earth L1 or L2	approximately 1.5 million km from Earth	approximately 5 s	Separate deep-space architecture

These figures exclude processing, routing, queues, gateway transport, protocols and ground-station availability.

GEO should not be described as continuously illuminated. GEO spacecraft encounter seasonal eclipse periods near the equinoxes, with real operational eclipses reaching roughly 72 minutes in some cases. GEO also requires a different communications, stationkeeping, radiation and regulatory architecture.

5.4 Why Not L1 in v0.1

Sun-Earth L1 and L2 add seconds of propagation, deep-space communications, stationkeeping and a different radiation environment. They may be useful for a specialized science or relay customer, but they do not reduce the principal power, thermal and customer risks of the initial program.

6. Rev A Power and Thermal Model

Rev A is a screening model for transparent comparison. It is not a mission-grade orbit propagator or a released design.

6.1 Public Model Assumptions

Parameter	Rev A value	Status
Reference altitude	550 km	Working assumption
Maximum screening eclipse	35.61 min	Two-body cylindrical-shadow calculation
Net BOL array output	300 W/m² active photovoltaic area	Design assumption
EOL solar fraction	0.82	Working assumption
Direct-path efficiency	0.92	Working assumption
Battery charge efficiency	0.90	Working assumption
Battery discharge efficiency	0.95	Working assumption
Solar energy margin	0.15	Working assumption
Maximum battery depth of discharge	0.40	Working assumption
EOL battery capacity fraction	0.80	Working assumption
Battery reserve	0.15	Working assumption
Net radiator rejection	350 W/m²	Screening assumption

The 300 W/m² value is a net design assumption. It is not the solar constant and is not a published specification of a selected product.

6.2 Load and Solar Sizing

The public model represents total load as:

$$P_{load}=N P_{fixed}+P_{compute}(1+f_{variable})$$

For SI-N1, the reference case uses 1 kW compute and 0.65 kW combined fixed and variable overhead, producing a 1.65 kW constant screening load.

The required EOL solar power is calculated over sunlight and eclipse:

$$P_{solar,EOL}=\frac{1+m}{t_{sun}} \left(\frac{P_{load}t_{sun}}{\eta_{direct}}+ \frac{P_{load}t_e}{\eta_{charge} \eta_{discharge}} \right)$$

$$P_{solar,BOL}=\frac{P_{solar,EOL}}{f_{EOL}}$$

SI-N1 carries 10 kW BOL as the proposed installed reference capability, above the minimum energy-balance calculation. The excess remains subject to pointing, degradation, thermal limits and subsystem evidence.

6.3 Battery

The battery target is screened as:

$$E_{battery}=\frac{P_{load}t_e}{\eta_{discharge}DOD f_{capacity}}(1+m_{battery})$$

For the SI-N1 reference case, the model produces approximately 3.7 kWh target nameplate. A 4-8 kWh range remains a planning envelope until communications, heaters, propulsion, safe mode and battery-cycle behavior are modeled in time domain.

6.4 Active Solar Area and Radiator Area

At 300 W/m² net BOL output:

$$A_{array}=\frac{P_{solar,BOL}}{300\ \mathrm{W/m^2}}$$

Ten kilowatts BOL requires approximately 33.3 m² of active photovoltaic area. Gross deployed wing area will be larger because of packing, wiring, mechanisms and inactive edge zones.

Radiator area is screened as:

$$A_{radiator}=\frac{P_{heat}}{q_{net}}$$

Rev A produces approximately 4.6 m² for the SI-N1 base load and 33.3 m² for the SI-N10 reference load. Actual geometry depends on temperature limits, coatings, view factors, Earth infrared, solar exposure and the selected heat-transport architecture.

6.5 Illustrative Scaling

Metric	SI-N1	SI-N10	SI-N100	SI-G1MW
Continuous compute	1 kW	10 kW	100 kW	1 MW aggregate

Metric	SI-N1	SI-N10	SI-N100	SI-G1MW
BOL solar, Rev A	10 kW installed	30.0 kW	288.4 kW	2.884 MW aggregate
Active photovoltaic area	33.3 m² installed reference	99.9 m²	961.5 m²	9,614.7 m² aggregate
Target battery nameplate	3.7 kWh base	26.9 kWh	259.3 kWh	2.593 MWh aggregate
Effective radiator area	4.6 m² base	33.3 m²	320 m²	3,201 m² aggregate

Values beyond SI-N1 are analytical extrapolations for architecture comparison. They are not completed spacecraft designs, launch manifests or commercial commitments.

7. Compute and Data Architecture

7.1 Modular Compute

The architecture does not commit v0.1 to a processor brand or processor count. A future hardware selection should be based on:

- performance per watt on the selected workload;
- memory capacity and error behavior;
- measured thermal response;
- recovery behavior following faults;
- lifecycle and availability;
- integration constraints.

7.2 Radiation-Aware Operations

The proposed runtime treats faults as expected operational events. At a public level, the architecture includes integrity checks, monitoring, checkpointing, controlled restart, telemetry and safe separation between customer workloads and spacecraft control.

7.3 Workload Flow

A proposed workload flow is:

1. Define a representative dataset and acceptance metric.
2. Reproduce the workload on the ground.
3. Profile time, memory, power and heat.
4. Test controlled fault and recovery behavior.

5. Package an authenticated workload artifact.
6. Execute first in an isolated observation mode.
7. Run in a controlled production window.
8. Deliver results and a reproducible evidence package.

Customer workload execution must remain isolated from spacecraft command and safety functions.

---

8. Communications Strategy

The proposed architecture separates:

- low-rate, highly reliable TT&C;
- mission-data downlink for results and selected source fragments;
- an optional experimental high-rate link;
- a potential partner-sensor interface.

Loss of a high-rate data path must not remove command or safe-mode capability.

Very high optical burst rates have been demonstrated from LEO, but such demonstrations require terminals, pointing, suitable ground infrastructure, weather conditions and operational access. The initial value proposition therefore remains data reduction and faster useful results, not unlimited downlink.

---

9. Proposed Mission Success Criteria

9.1 Hosted Pathfinder

Proposed success evidence includes:

- successful boot, update, watchdog and recovery behavior in orbit;
- measured workload performance per watt;
- logged memory, storage and accelerator events;
- checkpoint and restart evidence;
- measured thermal response;
- verified output from a space-originated dataset.

9.2 Node 1 kW

Minimum proposed success:

- establish two-way command and telemetry;
- reach a power-positive and thermally safe state;
- demonstrate controlled deployment;
- execute a useful workload and recover its result;
- preserve a credible disposal path.

Full proposed success:

- demonstrate 10 kW BOL-equivalent array capability under a frozen reference condition;
- sustain 1 kW continuous compute through repeated sunlight and eclipse cycles without violating approved energy or temperature limits;
- demonstrate customer-relevant data reduction while meeting the agreed quality threshold;
- demonstrate controlled recovery from representative fault classes;
- reconcile the power and thermal model against measured data;
- execute the approved end-of-mission plan when commanded.

Acceptance thresholds remain subject to SRR and PDR review.

---

10. Development Gates

Gate 1, Public Concept

- mission purpose and exclusions documented;
- Rev A model independently checked;
- representative customer workloads defined;
- preliminary licensing map established.

Gate 2, Hosted Pathfinder Review

- hosted interface proposal available;
- payload power, thermal, data and operating limits defined;
- compute candidates benchmarked;
- characterization plan approved;
- customer experiment defined.

Gate 3, Node 1 kW Preliminary Design

- Hosted Pathfinder evidence available, or a documented rationale to proceed in parallel;
- coupled orbit, power and thermal model closed with margin;
- launch and integration path identified;
- licensing schedule established;
- customer capacity requirement supported by evidence.

Gate 4, Node 10 kW Decision

- SI-N1 demonstrates continuous compute through sunlight and eclipse;
- thermal margin is supported by measured evidence;
- a customer outcome is demonstrated;
- the proposed service closes under current launch and operations assumptions.

---

11. Key Risks

Risk	Why it matters	Public mitigation path
Heat rejection	Nearly all compute power becomes heat	Coupled thermal modeling, breadboard testing, thermal-vacuum evidence, workload throttling
Radiation effects	Compute errors may be silent or disruptive	Component evidence, monitoring, integrity checks and controlled recovery
Solar deployment	Large deployables introduce single-event mission risk	Staged deployment, independent state evidence and partial-power safe modes
Battery life	LEO cycling and temperature can reduce usable capacity	Conservative depth of discharge, cycle testing and energy reserve
Attitude and flexible modes	Arrays and radiators affect pointing and control	Coupled structural-control analysis and deployment testing
Communications availability	A useful result still requires a viable data path	Independent TT&C, data reduction and multiple ground-access options
Customer interface	A flight demonstration needs real data and acceptance criteria	Ground replay, a defined dataset and a signed experiment definition
Launch and schedule	Rideshare date and orbit may change	Interface portability, schedule reserve and multiple compatible opportunities
Disposal	A failed disposal plan creates regulatory and sustainability risk	Early orbit-lifetime analysis, maneuver reserve and passivation planning

12. Regulatory and Sustainability Framework

A U.S. mission may require radio authorization, orbital-debris information, launch payload review and, depending on sensor and data architecture, remote sensing authorization. Final obligations depend on mission configuration and then-current rules.

Disposal is a design input. The proposed LEO architecture must address conjunction operations, maneuverability, passivation and reentry timing from the beginning of design.

This public summary is an engineering overview, not legal advice and not a license application.

13. Public Roadmap

Progress is defined by evidence gates rather than promised launch dates.

Near-Term Evidence

- finalize the public mission definition;
- independently review the screening model;
- select representative customer workloads;
- build and test a compute and thermal breadboard;
- define the Hosted Pathfinder experiment;
- complete initial regulatory consultations.

Hosted Pathfinder Phase

- freeze the public-safe payload objectives;
- complete environmental and functional test evidence appropriate to the selected host;
- integrate and operate the proposed hosted experiment;
- publish a public summary of measured results where permitted.

Node 1 kW Phase

- close the coupled power, thermal, communications and disposal architecture;
- complete formal design reviews;
- produce integrated ground-test evidence;
- proceed to launch only after technical, customer and regulatory gates are passed.

Longer-Range Phase

- consider SI-N10 only after SI-N1 evidence and customer demand;
- standardize modules before SI-N100 scaling;
- treat SI-G1MW, GEO and Lagrange missions as separate architectures requiring their own evidence and customer case.

14. Public Evidence Ledger

ID	Public claim	Source	Use
E-01	SpaceX publishes a smallsat rideshare entry price and an incremental mass rate	SpaceX Smallsat Rideshare	Public launch-market context, not a mission quote

ID	Public claim	Source	Use
E-02	SpaceX describes Starship as a fully reusable heavy-lift system designed for more than 100 t to orbit	SpaceX Starship	Long-range architecture context, not a price or availability commitment
E-03	Project Suncatcher publicly studies scalable space-based AI infrastructure	Google Research, Project Suncatcher	Evidence of sector direction
E-04	NASA describes spacecraft power as generation, storage and distribution	NASA SmallSat Power	Power-system framing
E-05	NASA describes thermal-control constraints for small spacecraft	NASA SmallSat Thermal Control	Thermal-risk framing
E-06	NASA describes hosted and integrated small-spacecraft platforms	NASA SmallSat Platforms	Hosted-pathfinder framing
E-07	NASA places Sun-Earth L2 at approximately 1.5 million km from Earth	NASA Webb Orbit	Deep-space distance context
E-08	NASA explains stability characteristics of Lagrange points	NASA, What Is a Lagrange Point?	Stationkeeping context
E-09	NASA's TBIRD demonstrated a 200 Gbit/s optical downlink from LEO under demonstration conditions	NASA TBIRD mission completion	High-rate communications context
E-10	NOAA documents seasonal eclipse operations for GEO weather satellites	NOAA GOES Eclipse Schedule	GEO comparison caveat
E-11	FCC rules require orbital-debris disclosures and post-mission disposal planning	47 CFR 25.114	Sustainability and licensing context
E-12	NOAA Part 960 governs applicable U.S. private remote sensing systems	15 CFR Part 960	Remote-sensing scope context

ID	Public claim	Source	Use
E-13	FAA payload review checks required licenses, authorizations and permits	FAA Payload Reviews	Launch-process context

All Rev A values are Satellite Inference screening calculations based on the assumptions disclosed in this public document. They remain subject to independent analysis and mission-specific evidence.

15. Explicit Unknowns

Important open questions include:

- exact orbit, inclination and mission duration;
- first customer dataset and acceptance metric;
- compute hardware and measured performance per watt;
- radiation behavior and recovery thresholds;
- final array, battery and thermal architecture;
- gross deployed area and pointing constraints;
- communications path and daily data volume;
- propulsion and disposal architecture;
- launch and integration interfaces;
- applicable authorizations and schedule;
- service pricing and sustained customer utilization.

These unknowns define the evidence required for the next revision.

16. Conclusion

The proposed path to large orbital computing infrastructure begins with a bounded and measurable flight experiment. Hosted Pathfinder is intended to characterize compute and recovery in orbit. Node 1 kW is intended to demonstrate the coupled power, battery, thermal, communications and continuous-compute system. Node 10 kW becomes relevant only after those results and a customer workload justify it.

Satellite Inference's central hypothesis is that orbital compute creates value when it is integrated with the power system, heat rejection, resilient software and the workflow of a specific space-based sensor. The next revision should be driven by measured evidence, not by a larger rendering or a more aggressive schedule.

Disclaimer

This document contains forward-looking concepts and early-stage engineering estimates. Actual performance, mass, cost, schedule, launch availability, customer demand and regulatory requirements may differ materially. Reference to a public company or agency source does not imply an agreement, endorsement or partnership.