Apsia

The energy layer for satellite operations

How much energy on every orbit, which tasks fit without ageing the battery, how long it has left, and whether there is energy to come home.

White paper · version 2 · September 2026 · apsia.space

1. The problem every operator shares

A satellite in low Earth orbit crosses Earth's shadow on almost every orbit: roughly fifteen times a day, about five and a half thousand times a year. Each crossing is a discharge the solar array cannot cover. The battery carries the whole spacecraft for up to thirty-six minutes, then has to recharge before the next shadow. How deep each discharge goes depends on the angle between the orbit plane and the Sun, and that angle drifts through the year; the same satellite can see no shadow at all in one season and a thirty-minute shadow in another.

Depth of discharge decides how many years the battery lasts. Published qualification data for space lithium-ion cells show life several times longer when cycled to twenty percent than to forty. Temperature and discharge rate change the damage done by the same depth. Solar arrays fade too, around a percent a year and more after a major storm; missions have ended because array power, not the battery, ran out.

Operators feel this as a budget they cannot see. Radar imaging time per orbit is bounded by power, not by demand. Optical missions lose the margin they designed in year one to long eclipses and aging arrays. Small IoT satellites with small packs live or die by how deep each cycle goes. Onboard battery management protects the pack from the next minute, not from the next five years, and it does not see rate or temperature effects. Power failures are the leading cause of small-satellite loss after the first month.

Today this is managed with spreadsheets at design time, a fixed depth-of-discharge rule in operations, and an anomaly alarm when a voltage crosses a line. None of those answers the three questions an operator actually asks.

2. Three questions, one physical model

Question	What Apsia produces	What it needs
How much energy will I have on every orbit?	A per-orbit forecast for the next day, week and month: minutes in shadow, energy generated and consumed, lowest state of charge, margin. Each figure with a range.	The orbit and a power sheet (array, battery, housekeeping). Presets for common buses.
Which tasks fit without shortening battery life?	Each activity in the plan: feasible or not, depth of discharge it adds, cost in battery life and in euros. A proposed plan that keeps committed contacts and moves or drops the rest. Before and after, side by side.	The activity plan as CSV or JSON.
When will the battery die?	State of health and internal resistance from telemetry, capacity fade per year, a knee in the fade when one appears, remaining life to a chosen threshold as a calibrated range, array health per panel, and consistency flags with a physical reason.	Battery voltage, current and, ideally, temperature and array currents. One file to start.

Will there be energy to come home?	Battery and array projected to the regulatory deadline against the disposal plan: readiness by year, the last safe start date, passivation time, and the energy share of the disposal probability.	The end-of-mission date, the rule, and the plan as a few activities.

The three modules share one physical model of the spacecraft: orbital geometry for shadow and sun angle, an array model with attitude and aging, a battery model with state of charge, resistance and an aging law that depends on depth, rate and temperature, and load profiles per activity. Telemetry calibrates the model; the model explains the telemetry. The number shown in one module is the number the others use.

Every output carries three things: the value, the range around it, and the assumptions it rests on. A recommendation always states the physical reason. Apsia recommends; it never commands.

3. The physics, in plain words

Shadow is geometry.

Given the orbit, the times a satellite enters and leaves Earth's shadow are computable to the second for weeks ahead, including the partial-shadow transitions. The energy problem is therefore mostly predictable; what is uncertain is the spacecraft, not the sky.

Depth of discharge sets the years.

Cycle life falls steeply with depth of discharge. The same energy taken late in a discharge, when the battery is already low, costs more life than the same energy taken early. Pricing tasks by the life they consume, rather than by a fixed depth rule, is what lets an operator take one more image on the orbits that can afford it and none on the orbits that cannot.

Rate and temperature matter as much as depth.

Recent flight-data studies show that identical depth-of-discharge cycles do different damage at different currents and temperatures, and that onboard battery management does not see the difference. A model that does, calibrated by the operator's own telemetry, is the difference between a counter and a forecast.

Arrays age, and space weather is in the loop.

Array current normalised by sun angle and shadow is a direct measure of array health. Geomagnetic storms raise drag, station-keeping demand and safe-mode risk; a storm forecast raises the energy reserve and the forecast says so.

Sources: published cell qualification data (Saft VES16; commercial 18650 cells at 20, 40 and 100 % depth of discharge); NASA and ESA mission reports on array degradation; peer-reviewed studies of rate- and state-dependent aging in orbit (2026). References available on request.

4. Validation on public data

Before any customer sends a byte, the models are validated on public flight data and published cell tests, always against a control arm: the best naive baseline for forecasts, measured capacity and the voltage-endpoint rule for health, and the operator's own plan for the planner. The figures below are read directly from the benchmark results file; a figure that is not yet measured is not printed.

1.2 %

capacity error against measured capacity on NASA cell-cycling data; the voltage-endpoint rule gives 29 %

2.5 %

capacity error on shallow, orbit-like discharge windows (first 1 Ah); the voltage-endpoint rule gives 12 %

11.2 %

24-hour power forecast error on a year of Mars Express telemetry, against 12.5 % for the best naive baseline (climatology)

0.0995

RMSE on the ESA Mars Express challenge metric (33 power lines, hourly) over a held-out Martian year, against 0.128 for the best naive baseline; the 2016 leaderboard's top five scored 0.080–0.089 on a different year

80 %

of held-out shallow windows (15–25 % depth) fall inside the stated ±p80 error band: the intervals are calibrated, not decorative

2/2

panel failures reported by the BIRDS CubeSat dataset found from the per-face currents, with 0 false positives on the other satellites

0.3 %

of the 57,784 LightSail 2 beacons whose panels were producing current (sunlight beyond doubt, whatever the attitude) where our shadow model still called eclipse; the spacecraft's own eclipse flag did so on 7.7 % of them. Across all 94,238 beacons the model agrees with the panels 95 % of the time, the flag 87 %; agreement with the flag itself, which carries a 121 s margin on each side, is 88.9 %

10 s

median scatter of the predicted umbra entries and exits of MinXSS-1 against the 383 crossings its panels recorded over a year in the ISS orbit (median timing error 9 s; 94.6 % of 153,288 packets agree with the panels' light, 'always sunlit' would score 62 %)

97.8 %

of 109,248 PW-Sat2 records over 26 months agree with its panels' light (2,360 whole orbits from the on-board telemetry files; umbra crossings timed to a median 25 s; 'always sunlit' would score 62 %); Wh per orbit of the next session, 11 days on, from the last session's level 6.3 % off, from its sunlit power × the orbit's geometry 7.2 % — with the sail out and tumbling, attitude sets the level, so the twin carries the last level observed

2.5 pts

state-of-charge error against measured truth over 616 NASA cell discharges, each scored with a pack learned only from what came before it — capacity from shallow windows (2.7 % one cycle ahead against 29 % for the rated value), resistance from current steps (rank correlation 0.75 with the cells' impedance spectroscopy); a Coulomb counter with the rated capacity gives 10.6 pts, a voltage table 10.9 pts. Against the on-board fuel gauge of MinXSS-1, the pack learned before each stored run and 255 days held out: 3.9 pts, the fixed nominal pack 4.4 pts, a voltage table 15.2 pts

2/2

satellites whose operators reported a battery failure flagged as falling, from pass samples alone, in a fleet of 26 open-archive satellites (109 satellite-years, 228,498 passes); they are the steepest falls of the fleet. Every sample is placed on its orbit by geometry; the same indicator ranks the measured capacity of NASA cells at 0.96 (Spearman) from six samples per pass, the raw mean of those samples at 0.56; thinned to one ground station, LightSail 2's three-year drift stays within 4 mV/yr of the full-archive value

82 %

of 389 held-out weeks on 28 satellites fall inside the band of the 30-day forecast of the voltage at the end of the eclipse (7 days ahead: 84 % of 393); the band is calibrated on the same forecast issued from earlier dates. The forecast starts from a month of the operator's own samples placed on the orbit and moves them by the eclipse length of the coming orbits: its error is 18 % below persistence and 29 % below climatology at 30 days (8 % and 31 % at 7 days); where the eclipse length moves by five minutes or more, 42 % below persistence at 30 days

179

satellites with public battery telemetry in the validation set

243

satellite-years of flight history, clipped to launch dates

10.3 yr

longest single history (MCUBED-2)

Datasets: BIRDS constellation on-orbit EPS dataset (four 1U CubeSats with NiMH packs, 5-second samples through eclipse, Data in Brief 2022; used for eclipse and energy-balance checks, not lithium-ion health); NASA Prognostics Center of Excellence lithium-ion cycling sets; ESA Mars Express power telemetry (three years); LightSail 2 mission telemetry (The Planetary Society, 94 238 beacons over 2019–2022 with per-cell voltages, ten panel channels and the on-board eclipse flag); MinXSS-1 housekeeping (LASP, one year in the ISS orbit with a fuel-gauge state of charge); PW-Sat2 (Warsaw University of Technology, 26 months of on-board telemetry files and beacons from the team's published session reports, whole orbits); decoded amateur-radio telemetry from the SatNOGS network (CC BY-SA). Each derived file carries its provenance.

How the health figures are validated.

The open-circuit-voltage model that ages with capacity was learned on the NASA cells and is scored leave-one-battery-out: the table used to grade a cell never saw that cell. Errors are reported for full discharges and for shallow, orbit-like windows separately, because shallow windows are where every method struggles; the error band around each level is calibrated on held-out cells and its coverage is printed above. The array check is chemistry-independent and uses the per-face currents only.

How the gain is measured.

The gain an operator gets is measured on their own plan, before and after, with the same tasks, seed and model, and reported with its control arm: the comparison is the control, not a projection. On real orbital telemetry the models are validated on what the spacecraft itself recorded — the light the panels saw, the on-board eclipse flag, the on-board state-of-charge gauge, the per-face array currents, the cell voltages of a parallel pack — and the capacity figures come from cell tests with measured capacity. The fleet indicator works on what a ground station receives, a few samples per pass: each sample is placed on its orbit by geometry and the end-of-eclipse voltage is tracked at constant load, temperature and eclipse length. It is checked three ways — against measured capacity on NASA cells sampled as passes, against the full archives of LightSail 2, MinXSS-1 and PW-Sat2 thinned to a single station, and against the battery failures the operators of the fleet published. The pack itself is learned in flight — its capacity from the discharges the satellite already makes, its resistance from the current steps it already makes, its state of charge re-zeroed where the open-circuit voltage is known — and every such figure is scored prequentially: learned up to a moment, tested on what came after it, never on its own answer. The margin ahead — the end-of-eclipse voltage a station will see in a week and in a month — is the operator's own recent level moved by what the geometry of the coming orbits says, with a band calibrated on the same forecast issued from earlier origins; it is scored against persistence and climatology on every archive with an orbit.

5. Energy for the end of life

Regulation now asks operators to prove disposal: the FCC five-year post-mission rule for low Earth orbit (in force since September 2024), ESA's Zero Debris requirements (ESSB-ST-U-007, successful-disposal probability of at least 0.9 and clearance within five years), ISO 24113 passivation of stored energy, and the 25-year guideline elsewhere. The compliance tools on the market compute orbital lifetime. None of them asks whether the spacecraft will still have the energy to execute the plan: a drag-sail deployment, weeks of electric manoeuvres in safe mode, or a passivation sequence, on a battery that has aged for the whole mission and an array that has degraded, in the worst eclipse season of the orbit.

Apsia answers that from the numbers it already measures. The battery level and its fade, with their ranges, and the array trend are projected year by year to the deadline. The plan is a short list of activities: power, duration, repetitions. The check compares, for each year, the energy generated and needed per orbit in disposal mode, and the usable energy against the deepest eclipse plus the biggest single draw. The result is a readiness table, the last date the plan can still start with margin, the hours passivation takes, and the share of projected cases with margin at the planned start. That share is the energy contribution to the disposal probability. It is never presented as the total probability: propulsion, attitude and ground segment stay with the operator. The page prints as a “power at disposal” annex for the debris file, with every assumption and the rule it answers to.

6. How it plugs in

Apsia is a decision layer on top of whatever mission control an operator runs. It never replaces it and has no command path. There are three ways in, and operators usually move through them in order.

Door	What happens	When
A file	Drop a telemetry CSV; get the battery report. No integration, no account. The free analysis stores nothing after the run.	Day one
A connector	Read-only connection to the telemetry store (time-series database, mission-control export, object storage). Analyses run continuously.	Pilot
Inside your screens	Alerts and the recommended plan shown in the dashboards operators already watch. Recommendations go back as files in the operator's formats. People approve every change.	Operator plan

Fleets with strict data rules run the same software as a container inside their own infrastructure, with no outbound connection. A workspace keeps every analysis with its satellite: history of health with ranges, resistance, array status, and alert rules (health below threshold, a knee, a dead panel, a reserve breach, silence for N days) that reach the operator by e-mail.

7. Data handling

- We read housekeeping telemetry: voltages, currents, temperatures, orbit elements, activity plans. Never payload data, never command links.
- Processing in the European Union on certified cloud infrastructure; the on-premises container keeps everything inside the operator's perimeter.
- The free analysis keeps nothing after the run. Workspaces keep per-orbit summaries and reports; raw telemetry is discarded after processing; deletion on request within 30 days.
- No model training on customer data without written consent. Fleet-level learning uses aggregate, anonymised parameters, with opt-out.
- Incident notification within 72 hours. Security questionnaires answered within five working days.

8. Who it is for

Radar imaging	Imaging time per orbit is power-bounded. Every extra minute the battery can afford is revenue; every minute taken from a battery that cannot afford it is years lost.
Optical Earth observation	Long eclipses at high latitudes and aging arrays quietly eat design margins. Array and battery health in one report.
IoT and relay constellations	Small packs, many cycles, many satellites. Depth of discharge decides fleet lifetime; the planner enforces it per satellite.
Orbital compute nodes	A GPU through a 35-minute eclipse is a battery problem before it is a compute problem. Duty-cycle planning for nodes is the next module.

9. Where this goes

- Now: battery health with calibrated ranges, array health per panel, energy forecast, energy-aware planning and the end-of-life energy check, by file, by connector (Yamcs, OpenC3 COSMOS, Grafana) or inside the operator's dashboards; a workspace with fleet history and alerts; validation on three real missions and published cell tests.

- Next: the pack learned in flight (capacity, resistance and the state-of-charge zero from the operator's own telemetry), attitude-aware array health, energy forecasts weeks ahead.
- Then: compute duty cycles for orbital compute nodes; the planner running on a partner's node in orbit.

10. About Apsia

Apsia is built in Talavera de la Reina, Spain. It grew out of work on constellation energy and orbital compute and narrowed to the one problem every operator we spoke to shares. Pricing: the battery analysis is free; pilots start at €5,000 per constellation for 90 days; operator plans are priced per satellite and month.

apsia.space · apsia.space/contact

This document describes capabilities and validation results at the date shown on the cover. Figures are computed from the public datasets named in section 4 and from the benchmark results file that ships with the software; they are not projections.