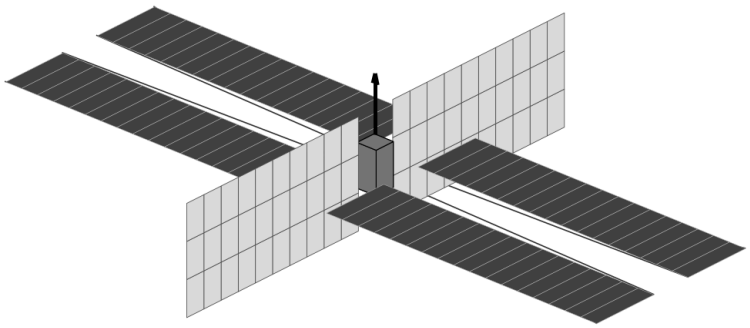
Engineering a 250 kW Orbital Compute Satellite

Preliminary spacecraft definition, materials, construction and verification

Reference design OC-250, revision A. 25 September 2026. English engineering companion to the orbital-compute thesis. The intended customers remain on Earth: workloads are uploaded, executed in orbit, and results returned. This document is a Phase-A reference concept with explicit sizing and verification requirements. It is not a manufacturing drawing set, a completed preliminary design review, a flight qualification or a supplier quotation.

250 kW IT / conceptual deployed arrangement

Arrow: Sun direction. Independent radiator orientation remains a design trade.



Four solar wings: 1,700 m²
Dark surfaces, nominally Sun-facing

Two radiator wings: 900 m²
Light surfaces, edge-on to sunlight

Figure 1. Proposed deployed arrangement. Dark wings collect sunlight; light wings reject heat. The illustration is a parametric engineering schematic, not detailed flight CAD. Radiator view factors and articulation remain open design items.

1. What the satellite would look like

The reference spacecraft is a large deployable power-and-thermal platform surrounding a compact unpressurized computer module. Its central structure is approximately 3 m by 3 m by 6 m. Four dark photovoltaic blankets extend to a total solar span of about 93 m. Two light-colored radiator wings form an approximately 66 m by 15 m heat-rejection structure oriented nominally edge-on to sunlight. The photovoltaic collection area is 1700 m^2 and the effective single-emitting radiator area is 900 m^2. An allocated folded envelope is a cylinder 6.5 m in diameter and 12 m long. These dimensions define a geometry to study; they do not establish a launcher interface or a collision-free deployment sequence.

This is a spacecraft carrying purpose-built, conduction-cooled electronic trays. There is no pressurized room containing ordinary air-cooled server racks, and no astronaut maintenance system is included. The trays are modular for manufacture and fault isolation. Replacement after launch is not assumed.

The mass budget is approximately **16.02 tonnes current-best-estimate dry**, plus a **3.20-tonne dry contingency**, **2.36 tonnes of propellant** and **0.40 tonne of launch-adapter allocation**. This yields approximately **22.0 tonnes injected with the adapter**, or 21.6 tonnes wet spacecraft mass. Every line is a design allocation. The estimate could rise after detailed structures, thermal-survival, radiation and packaging analyses. It is nevertheless much more complete than a hardware-only area calculation.

1.1 Correction to the earlier conceptual estimate

The previous 5.4-tonne illustration used favorable illumination, lightweight arrays and a 350 K radiator. Detailed engineering changes all three. A 650 km dawn-dusk sun-synchronous orbit can have roughly twenty-minute seasonal eclipses. End-of-life generation must both run the server and recharge its battery. A radiator compatible with a roughly 315 K coolant supply needs a colder, nonuniform emitting surface than an arbitrarily uniform 350 K panel. This document therefore supersedes that earlier physical-sizing example. It does not carry forward its cost estimate as if the hardware had remained unchanged.

Design quantity	OC-250 allocation	Meaning
IT electrical capacity	250 kW	Accelerators, hosts, memory, storage and internal network
Regulated bus design load	300 kW	IT plus auxiliaries, downstream conversion and reserve
Thermal design load	330 kW	Radiator-routed heat, including charging/transient allowance
Orbit; inclination	650 km; 97.986 deg	Circular first-order dawn-dusk SSO candidate
Orbital period	97.73 min	Equatorial-radius model
Analytical seasonal eclipse; storage design case	19.73 min; 25 min	Exact mission ephemerides still required
PV geometric area; sunlit EOL output	1700 m^2; 435 kW	Conditional supplier and degradation assumptions
Eclipse battery; separate avionics reserve	235 kWh; 20 kWh	BOL nameplate, approximately 255 kWh total
Radiator area; effective temperature	900 m^2; 320 K	One effective emitting face; nonuniform skin
Coolant inlet/outlet	315 K / 335 K	Candidate PAO loop; component branches vary
IT partition	8 zones × 31.25 kW	Each zone has electrical and thermal isolation
Ground assembly partition	40 trays × 6.25 kW	Five trays per zone; vendor-neutral allocation
Nominal mission design study	5 years	Materials, radiation and aging must be modeled over this life
Folded envelope allocation	6.5 m diameter × 12 m	Larger-fairing interface to be procured and validated
Deployed geometric spans	93 m × 66 m × 15 m	Clearances, gimbals and plume envelopes additional

2. Physical architecture and subsystem interfaces

The primary frame transfers launch forces from a base adapter into longitudinal members and transverse decks. Computer trays, battery modules and power converters attach close to this backbone. Radiator packs and photovoltaic roll cassettes have independent ascent restraints into the frame. Deployment hinges carry controlled deployment and on-orbit loads; they must not become the sole load path for tonne-class folded wings during launch.

Two independent regulated power domains each serve four compute zones. Eight sealed thermal circuits remove heat from their assigned zones. Independent flight avionics command power, thermal safety, deployment and propulsion. Optical terminals on isolated mounts track gateways without rotating the whole spacecraft through every contact. RF telemetry and command provide a recovery path when payload computers or optical terminals are unavailable.

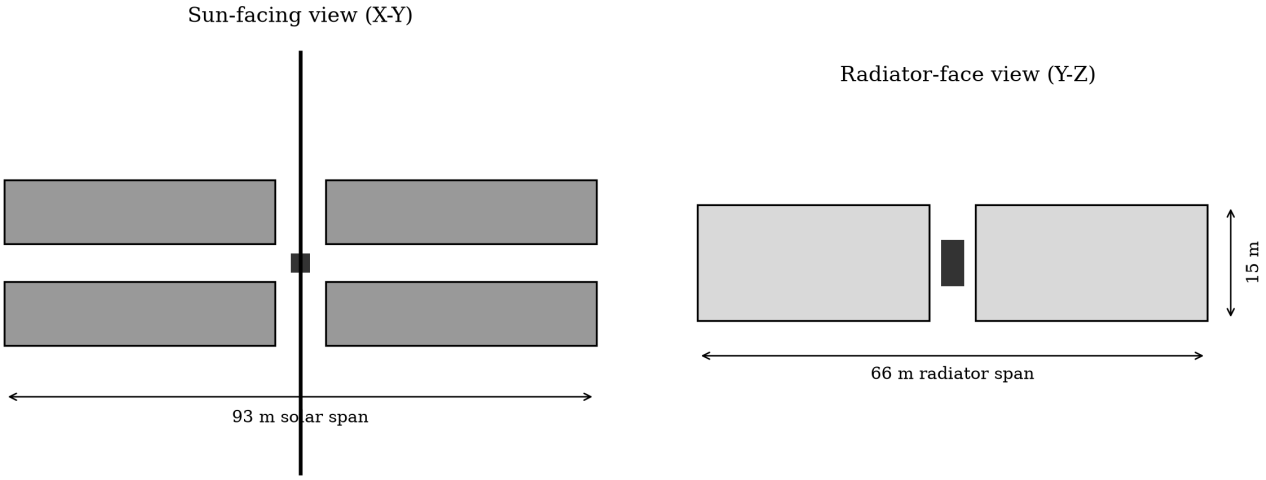
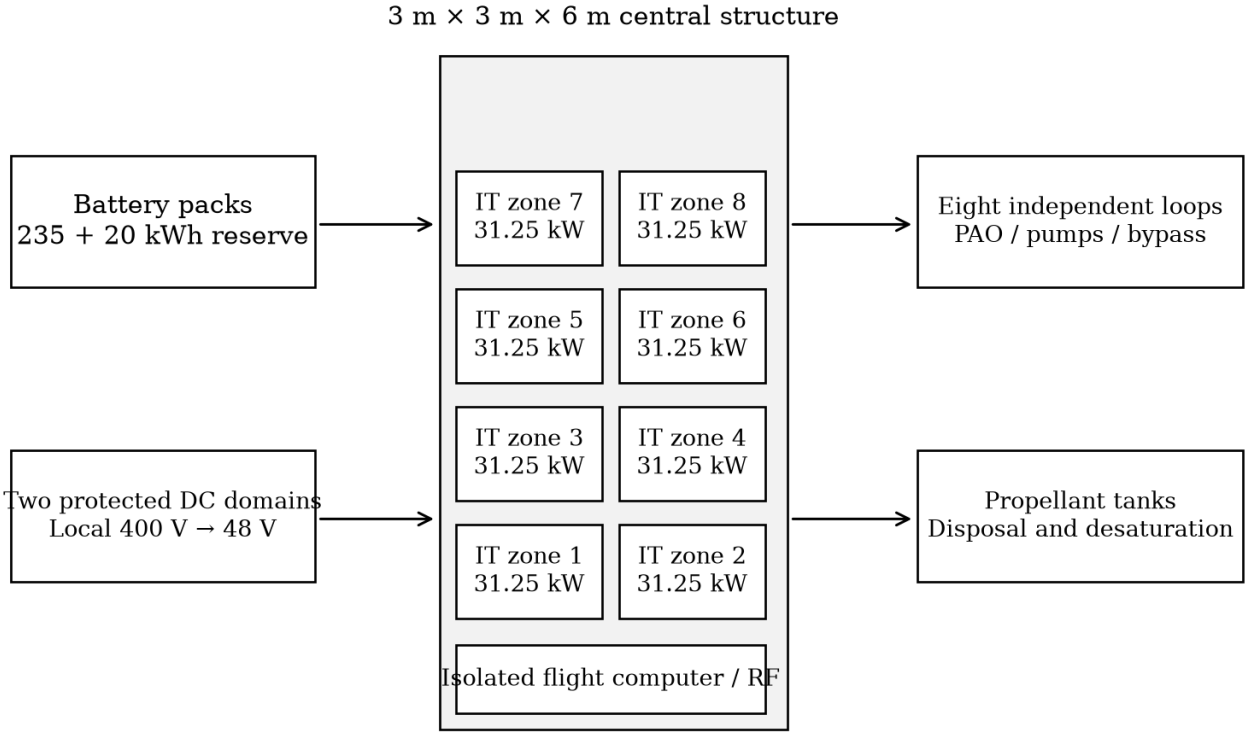


Figure 2. Orthographic allocations. The solar geometry is four 10 m × 42.5 m blankets. Each radiator wing is 30 m × 15 m. The dimensions do not include dynamic deflection, protection margins or an articulated-radiator envelope.



Functional placement only; harness, tank clearances and access require detailed CAD.

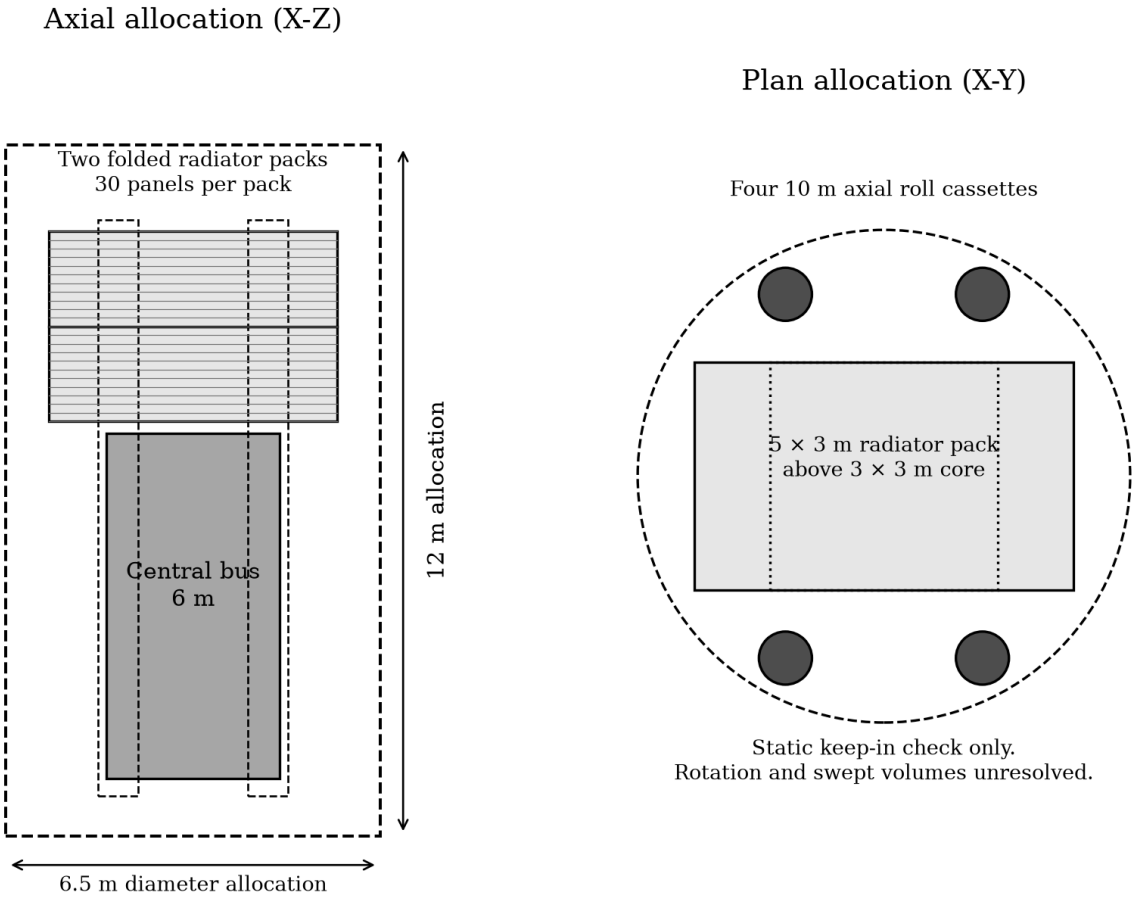
Figure 3. Central hardware functions and eight independently isolatable IT zones. The battery, propulsion and thermal blocks are functional placements; tank clearances and three-dimensional routing require detailed CAD.

An interface control document should identify every structural joint, power feeder, fiber link, fluid coupling, grounding connection and software authority. For each IT tray it must state maximum and transient power, input-voltage range, dissipation by location, allowed coolant and inlet conditions, pressure drop, mechanical mounting points, reset behavior and error-reporting format. A processor board is accepted only when its entire support chain, including memory, regulators, storage and optical interfaces, passes the same mission envelope.

3. Folded packaging and deployment

The 900 m² radiator is divided into sixty 3 m × 5 m panels. Each wing contains thirty panels, conceptually arranged as three ten-panel strips. A 40 mm panel-thickness allocation plus 15 mm packing gap gives a 1.65 m stack per wing. Two stacked 5 m × 3 m packs occupy 3.30 m above the central bus. This fits inside the *allocated* cylindrical cross-section: the pack corner radius is 2.92 m against a 3.25 m allocation. Locks, flex hoses and local supports still need a physical layout.

Four nominal 10 m axial roll cassettes can be allocated at transverse coordinates $(x, y) = (\pm 1.3, \pm 2.4)$ m, with a 0.35 m outer-radius allocation. Their outer radial extent is approximately 3.08 m, leaving only 0.17 m to the allocated cylinder. This is a static volume check, not a launch clearance margin. The true fairing taper, spacecraft deflection, thermal insulation and launch adapter determine the usable volume. SpaceX's public Falcon guide illustrates why payload integration must use mission-specific keep-in and load interfaces; no Falcon compatibility is claimed here. The 2020 Starship guide is preliminary historical accommodation information, not a current mission acceptance.



Preliminary 6.5 m × 12 m stowed allocation. The current launcher interface and deployment path are unverified.

Figure 4. Static volume allocation. The X-Z view projects solar cassettes in front of/behind the radiator packs; the X-Y view shows their separate positions. Cassette rotation, translation and full deployment swept volumes remain unresolved.

The photovoltaic blanket width is 10 m, so its spool axis is 10 m long. Its stowed axis is axial; its deployed axis must lie in the solar-wing plane. Each cassette therefore needs an articulated cradle that translates clear of the core and rotates before rollout. Adjacent cassettes cannot rotate in place into the same occupied region. A sequenced swept-volume model must prove clearance from radiators, other cassettes, harnesses and the launcher separation trajectory.

For a 42.5 m blanket winding length, a simple packing bound is

$$t_{eq} \leq \frac{\pi(R_o^2 - R_i^2)}{L}.$$

At a 0.35 m outer radius and 0.15 m mandrel radius, the equivalent winding thickness bound is 7.39 mm, or 5.90 mm when usable outer radius is reduced to 0.32 m for housing. This does not prove minimum cell bend radius, interconnect durability or room for folded booms. The four roughly tonne-class array packages need distributed ascent restraints along their ten-metre length.

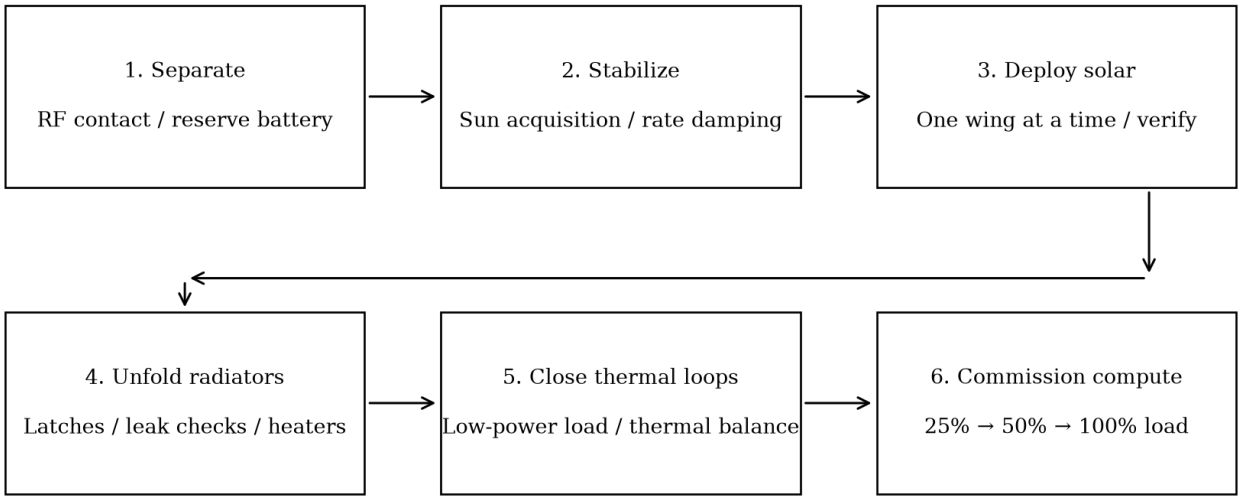


Figure 5. Proposed deployment and commissioning logic. The sequence must be revised if the final mechanism geometry imposes a different order. Each step has telemetry and an explicit abort/reduced-power state.

4. Candidate materials and why they are used

The following table is a proposed trade baseline. It does not assert that any listed commercial stock, adhesive or coating is automatically qualified for the mission. Property allowables depend on product form, thickness, heat treatment, manufacturing process, operating temperature and statistical basis. ECSS requires establishment and verification of applicable mechanical and physical material properties, including environmental effects. ECSS-E-ST-32-08C Rev.1.

Component	**Proposed material or construction**	**Engineering reason**	**Principal failure modes**	**Required evidence**
Primary bus frame and equipment decks	Machined aluminium 6061-T6/T651 candidate, with alloy/temper changed where joint or strength trade requires	Inspectable load paths; manufacturability; electrical and thermal continuity	Yield, buckling, fatigue, corrosion before launch, thread pull-out	Certified stock; tempe rature-dependent allowables; stress analysis; proof/modal tests
High-load attachment inserts	Titanium Ti-6Al-4V candidate or qualified corrosion-resistant steel	Compact attachment and wear-resistant contact where justified	Fretting, fatigue, galvanic interface, differential expansion	Joint coupons; preload retention; fracture-critical assessment

Component	Proposed material or construction	Engineering reason	Principal failure modes	Required evidence
Bus shear panels and radiator substrates	Aluminium face sheets on vented aluminium honeycomb; localized solid inserts	Bending stiffness with low mass; conductive heat spreading	Face wrinkling, core shear/crush, adhesive disbond, trapped gas	Coupon and panel tests; ultrasonic/thermographic inspection; vent verification
Solar deployment booms	Space-qualified carbon-fibre/polymer laminate with tailored layup and environmental protection	Low mass and tailored bending/torsional stiffness	Matrix cracking, delamination, deployment jam, thermal distortion	Laminate allowables; ageing; full-scale deployment and dynamic tests
Photovoltaic blanket	Qualified flexible dielectric or mesh substrate with cell islands, flexible interconnects and coverglass	Compact stowage; articulated blanket with protected cells	Cell fracture, interconnect fatigue, insulation damage, arcing	Roll/unroll fatigue, electrical continuity, plasma and thermal-vacuum tests
Radiator heat-spreading and fluid circuits	Aluminium panels and compatible alloy channels; fluid-compatible welded/brazed manifolds	Distributed heat rejection and minimized fluid interfaces	Leak, corrosion, fatigue, blockage, freeze damage	Fluid/material compatibility; pressure proof; helium leak and flow tests
Radiator optical finish	Qualified high-emittance, low-absorptance coating or optical solar reflector system	Reject internal heat while limiting absorbed sunlight	UV darkening, contamination, adhesive failure, erosion	Beginning/end-of-life optical properties; radiation/UV/AO exposure
Thermal insulation	Conductive, grounded multilayer insulation with qualified polymer films and protective outer layer	Reduce unwanted heat exchange outside dedicated radiators	Tearing, charging, trapped gas, contamination	Venting, bonding, contamination and exposure tests
Sensitive computer enclosures	Aluminium boxes with locally optimized shielding	Structural housing plus some particle attenuation	Residual single-event effects, thermal penalty, excessive shielding mass	Radiation transport and device tests; fault-injection validation
Bearings, hinges and release devices	Qualified metallic mechanisms with application-specific lubricant/coating	Controlled deployment and retention in vacuum	Cold welding, lubricant loss, wear, latch failure	Vacuum life test; cold/hot deployment; end-of-travel and release-margin tests

The cold-plate alloy and coolant chemistry must be selected together. A high-conductivity copper insert may be justified locally under a concentrated heat source, but widespread copper penalizes mass and introduces dissimilar-metal interfaces. Likewise, carbon fibre is not a universal substitute for metal: low through-thickness conductivity and difficult bonded interfaces can undermine a thermally loaded electronics enclosure. These are design trade recommendations, requiring coupons and assemblies rather than catalogue-property extrapolation.

The practical bill of materials therefore centers on qualified aluminium alloys, carbon-fibre composite deployables, protected III-V photovoltaic cells, compatible metal fluid channels, low-absorptance radiator finishes and conventional space-qualified electronic packaging. Exotic bulk materials are not required to state the architecture. Exact alloy temper, laminate, resin, adhesive, coating, lubricant and cell chemistry must be selected through a controlled materials-and-processes register. A generic material name never substitutes for a qualified production process.

For a concrete photovoltaic stack, GaInP/GaAs/Ge triple-junction cells with space-qualified coverglass, bypass diodes and welded interconnect assemblies are a credible procurement class. Spectrolab's product documentation describes this material system and packaged cell assemblies. It supports the cell technology choice, not the completed 42.5 m blanket or its radiation life. The flexible substrate and cell-island pattern still require roll-fatigue, insulation, thermal-cycle and atomic-oxygen verification.

For non-radiating equipment insulation, a candidate blanket uses aluminized polyimide reflective layers, low-conductivity spacers and a protected outer layer such as PTFE-impregnated glass-fibre beta cloth where appropriate. NASA materials studies describe this construction and the importance of electrical grounding. Layer count, seams, vents, fasteners and exposed-edge protection follow the actual environment. The lithium-ion battery chemistry and pressure-vessel liner/composite system remain supplier down-selections; neither an automotive battery enclosure nor an arbitrary composite tank can be substituted without qualification.

5. What the structure must survive

The load matrix must cover transport and lifting, handling misalignment, ascent quasi-static loading, coupled vibration, acoustic loading of folded panels, separation shock, fairing depressurization, release transients, deployed manoeuvres, thermal gradients and pressure in coolant/propellant hardware. Quasi-static acceleration alone cannot qualify the vehicle. For scale only, a hypothetical 22,000 kg spacecraft at an illustrative 6 g axial acceleration produces

$$F_z = ma_z = 22{,}000 \times 6 \times 9.80665 \simeq 1.29\,\mathrm{MN}.$$

This is a worked example, not a selected launcher requirement. With lateral acceleration a_y and centre-of-mass height z_c, the root bending moment is approximately $M_x = ma_y z_c$. Actual launch combinations and dynamic amplification come from coupled loads analysis. For a load-controlled screening case, a margin may be written $MS = R_{allow}/(FS\,R_{limit}) - 1$, with the appropriate failure mode, allowable and factor prescribed by the project's verification policy. Composite buckling, bonded joints and pressurized hardware require their own treatment.

Vacuum requires vent paths in sandwich panels and equipment volumes, material outgassing control and no reliance on air convection. Mission temperatures must come from hot/cold analyses for sunlit, eclipse, safe mode, partial deployment and failed-loop states. A common blanket temperature range for every component is not an engineering requirement. The selected processor, battery, adhesive, lubricant and coolant each impose separate operational and survival limits.

Atomic oxygen remains relevant in LEO; exposure depends on altitude, attitude, time and solar activity. NASA documents erosion of polymers and carbon-fibre polymer matrices, and the sensitivity of protective coatings to geometry and defects. Candidate protective systems require testing of seams, cut edges and damage sites as well as pristine coupons. NASA atomic-oxygen protection study. Durability calculations should use environment-specific fluence and measured erosion yield, $d = E_y \Phi_{AO}$, with geometry and shielding accounted for. NASA's handbook provides flight-derived polymer durability data. NASA-HDBK-6024.

UV exposure and contamination change radiator absorptance and photovoltaic transmission. Adhesives, coatings, cable insulation and composite resin therefore enter a controlled materials list. Outgassing screening is necessary but insufficient near optical terminals: deposition depends on actual temperature,

emission kinetics and viewing geometry. ECSS explicitly allows more stringent acceptance near sensitive hardware. ECSS-Q-ST-70-02C. Bake-out and witness samples should precede final optical integration; long-term contamination analysis can use kinetic outgassing measurements. ECSS-Q-TM-70-52A.

For a first sizing screen, combined axial and bending stress is $\sigma = N/A + Mc/I$, while an ideal slender compressive member has Euler load $P_{cr} = \pi^2 EI/(KL)^2$. Neither equation replaces a panel-buckling or coupled-launch analysis. Aluminium honeycomb may fail by face wrinkling, core shear or insert pull-out before a simple beam stress limit is reached. Dynamic amplification and joint compliance are usually decisive for tall folded payloads.

After deployment, low stiffness is as consequential as strength. A simplified modal frequency scales as $f_n \sim \sqrt{k/m}/(2\pi)$. A structural member can survive the load yet vibrate too much for the laser terminal. Pumps, rotating actuators, wheel imbalance and release transients must be represented in the same flexible-body model used by the attitude controller.

Thermal expansion is also an interface load. For an unconstrained member, $\Delta L = \alpha L \Delta T$; differential strain between bonded materials is $\Delta\varepsilon = (\alpha_1 - \alpha_2)\Delta T$. Flexures, compliant joints and controlled fastener preload accommodate the mismatch between titanium inserts, aluminium panels and anisotropic composite booms. ECSS mechanism guidance provides a framework for mission-tailored mechanism and life verification.

6. The orbit and the electrical design point

The reference orbit is circular at 650 km. Using equatorial Earth radius $R_E = 6378.137$ km, $\mu = 398600.4418$ km³/s² and $J_2 = 1.08262668 \times 10^{-3}$, first-order sun-synchronous precession gives

$$\dot{\Omega} = -\tfrac{3}{2} J_2 \sqrt{\tfrac{\mu}{a^3}} \left(\tfrac{R_E}{a}\right)^2 \cos i, \qquad i \simeq 97.986^\circ.$$

For solar declination δ, solar right ascension α and ascending-node longitude Ω,

$$\sin\beta = \sin i \cos\delta \sin(\Omega - \alpha) + \cos i \sin\delta.$$

An ideal 18:00 local ascending node has $\Omega - \alpha \simeq 90^\circ$. At the unfavorable solstice, the minimum beta is about 58.58 degrees. The cylindrical-shadow zero-eclipse boundary is $\beta_c = \sin^{-1}(R_E/a) \simeq 65.16^\circ$. Below it,

$$t_e = \tfrac{P_o}{\pi} \cos^{-1}\left(\frac{\sqrt{1 - (R_E/a)^2}}{\cos\beta}\right).$$

The result is approximately 19.73 minutes of eclipse in a 97.73-minute revolution. The 06:00 case exchanges seasons. The model uses ideal local-time tracking, a cylindrical shadow and an idealized annual declination; launch-epoch propagation, penumbra, maneuvers and orbit uncertainty are mandatory next steps. NASA orbital power analysis supports treating illumination as an orbital design calculation.

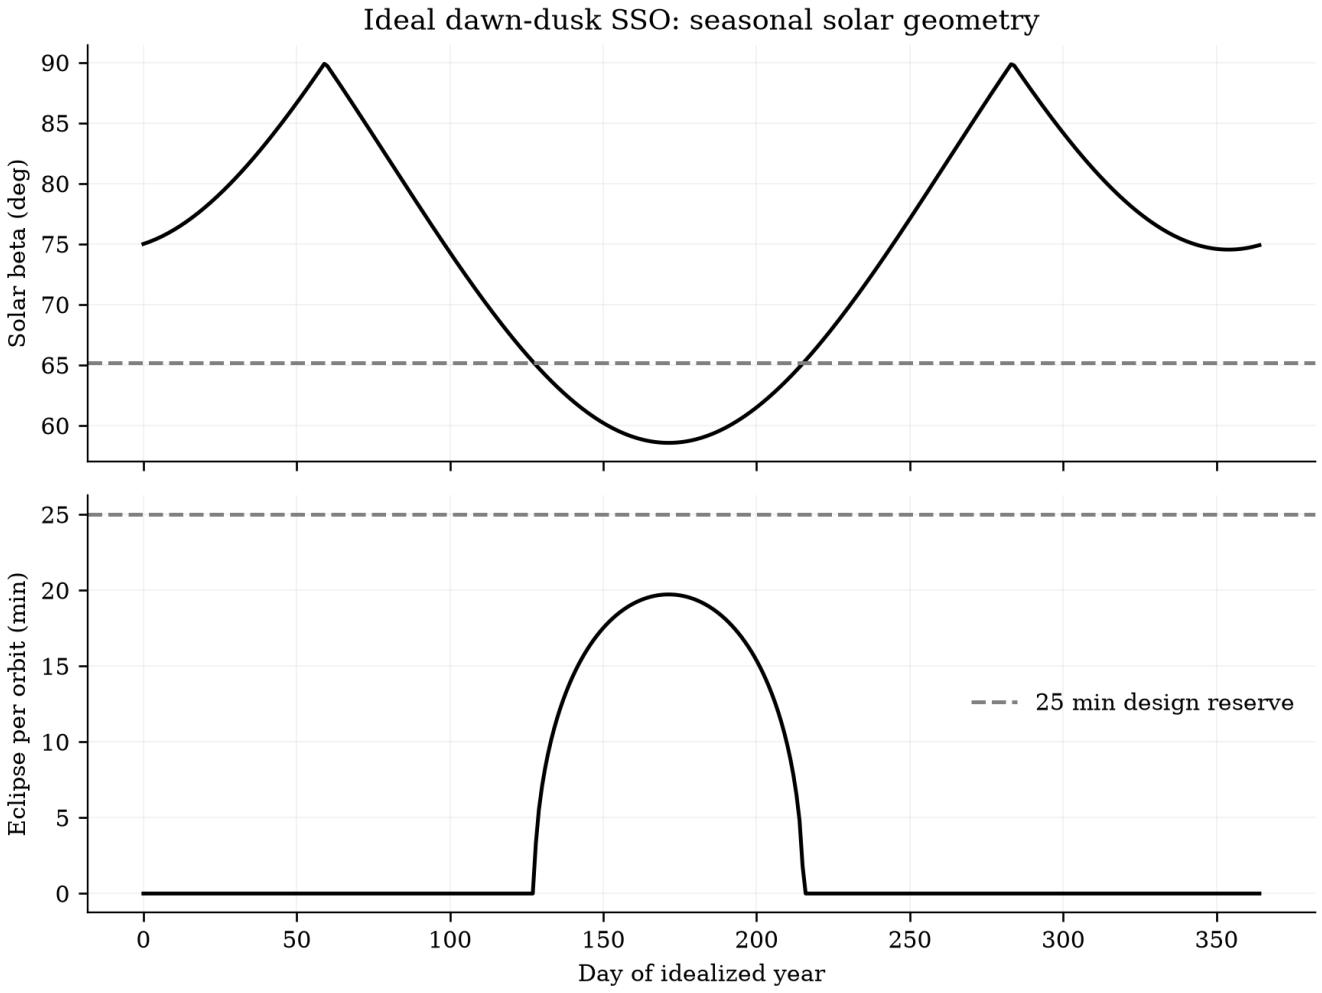


Figure 6. Seasonal beta and cylindrical eclipse for the idealized 18:00 SSO. Twenty-five minutes is a sizing allocation, not a claim that the exact mission has that maximum eclipse.

6.1 Electrical boundaries and load allocation

The 300 kW design load is defined at the regulated main bus. The 0.96 array-delivery factor below accounts for upstream array regulation and harness losses. Downstream conversion from the main bus to local computer rails remains inside the following 300 kW budget, avoiding double-counting the same loss.

Main-bus load	**Allocation**	**Interpretation**
IT payload	250 kW	Eight zones, including memory/network/storage
Pumps and thermal actuation	8 kW	Compared with a 7.73 kW hydraulic screening case
Optical/RF communication hardware	6 kW	Terminal selection and duty cycle unresolved
Flight avionics and attitude hardware	4 kW	Operating allocation
Local heaters and ancillary loads	4 kW	Nominal operation only
Downstream power conversion	14 kW	Device efficiencies must replace this allocation
Unallocated operating reserve	14 kW	Consumed explicitly during later design
Total regulated-bus design load	300 kW	Full-power design point

6.2 Solar generation and eclipse recharge

For solar irradiance $G_{\min}$, cell efficiency η_c, cell packing fraction f_p, temperature factor d_T, end-of-life degradation factor d_R, incidence factor d_θ, and harness/MPPT delivery efficiency η_e, usable EOL array output per geometric square metre is

$$s_{EOL} = G_{\min}\,\eta_c f_p d_T d_R d_\theta \eta_e.$$

The following are explicit sizing assumptions, to be replaced by supplier and mission-environment data:

Parameter	Preliminary value	Meaning
Minimum solar irradiance	1317 W/m²	Aphelion sizing condition
Cell efficiency at reference condition	0.30	III-V multi-junction candidate
Active-cell packing fraction	0.90	Geometric panel fraction
Hot-cell temperature factor	0.90	Requires equilibrium-temperature closure
EOL radiation/contamination factor	0.85	Combined allowance; not a radiation analysis
Incidence factor	0.98	Tracking and alignment allocation
Wiring and regulated delivery efficiency	0.96	Conversion boundary must match mass estimate
Battery round-trip efficiency	0.90	Recharge sizing assumption

These assumptions give $s_{EOL} = 255.9$ W/m². NASA describes contemporary III-V multi-junction space cells around 30-34% reference efficiency and emphasizes hot EOL rather than beginning-of-life cell efficiency when sizing complete arrays. This supports the technology class, not the specific derating allocations above. NASA, Power.

Let t_e be eclipse duration and η_{rt} battery round-trip efficiency. Array energy during sunlight must supply the live bus and restore eclipse energy:

$$s_{EOL} A_{PV}(P_o - t_e) = P_b(P_o - t_e) + \frac{P_b t_e}{\eta_{rt}}.$$

Consequently,

$$A_{PV} = \frac{P_b}{s_{EOL}}\left[1 + \frac{t_e}{\eta_{rt}(P_o - t_e)}\right].$$

For $P_b = 300$ kW and $t_e = 25$ min, $A_{PV} \simeq 1620$ m². A 1700 m² geometric allocation provides approximately 435.1 kW EOL sunlight output and roughly 5% additional area headroom beyond this calculation. Do not add an undocumented second blanket margin after already counting the 300 kW bus allocation. Dividing average power by an illumination fraction alone misses battery losses and transient power requirements.

Candidate construction is multiple independently switched roll-out wings: III-V cells, radiation/contamination coverglass, electrically insulated flexible polyimide blanket, protected interconnects, and CFRP deployable booms. Wings must radiate their own unconverted solar heat. Their surfaces are not simultaneously available as central-computer radiator area. Array dimensions, shadow clearance and pointing mechanisms must be integrated with separate thermal wings.

Beginning-of-life array rating at reference incidence is approximately 625 kW, but hot end-of-life output at the regulated bus is only about 435 kW under the chosen factors. The 25-minute eclipse case needs about 415 kW during sunlight, leaving roughly 4.9% generation headroom. Excess early-life generation

must be curtailed at the arrays or rejected externally. Dissipating hundreds of excess kilowatts in a central shunt would invalidate the radiator design.

6.3 Battery, distribution and recovery reserve

Specify battery mass using pack-level nameplate energy, not a bare-cell marketing value:

$$E_{BOL} = \frac{P_b t_e}{DOD\,\eta_{dis} f_{cap,EOL}}, \qquad m_b = E_{BOL}/e_{pack}.$$

For depth of discharge 0.70, discharge efficiency 0.95 and EOL retained capacity 0.80:

Full-load bridge duration	BOL battery nameplate	Pack mass at 150-200 Wh/kg
10 min	94.0 kWh	470-627 kg
20 min	188.0 kWh	940-1253 kg
25 min	235.0 kWh	1175-1566 kg

The 150-200 Wh/kg range is a design sensitivity, not a qualified vendor promise. Life testing must reflect calendar aging, eclipse-cycle count, state of charge, temperature and peak current. Add a separately isolated 20 kWh nameplate avionics reserve to the 235 kWh eclipse bank. Total allocated storage is approximately 255 kWh. With the same EOL, discharge and depth factors, the reserve alone supplies a 5 kW essential-electronics target for about 2.1 hours. This does not establish survival of the deployed external coolant circuits. A failure that forces load shedding protects the vehicle but interrupts contracted compute. Each battery string needs isolation, autonomous management, cell-temperature sensing, propagation-resistant compartmentation and a vent path directed away from optics and blankets.

An initial electrical trade should compare 270 V and 400 V distribution with isolated 48 V cassette buses and point-of-load conversion. At 300 kW, total current is 1111 A at 270 V or 750 A at 400 V, versus 6250 A at 48 V. Partition power into multiple protected feeders: eight 37.5 kW zone allocations draw about 94 A each at 400 V, including each zone's share of auxiliaries. Design conductor section from hot resistivity, voltage drop and local cooling, not current alone:

$$P_{harness} = I^2 \rho_e L/A_c, \qquad m_{harness} = \rho_m L A_c.$$

High voltage reduces conductor mass but creates a major qualification task in LEO plasma. The baseline requires insulated conductors, encapsulated cell/interconnect edges where appropriate, controlled grounding, electrical bonding across moving structures, arc detection, fast current-limited switching and cold-open-circuit voltage analysis. NASA's LEO charging standard identifies high-voltage array interactions as potentially destructive; its stated latitude scope does not by itself qualify a near-polar SSO mission. Auroral charging and the complete operational orbit require additional assessment. NASA-STD-4005; NASA-HDBK-4006A; NASA-HDBK-4002B.

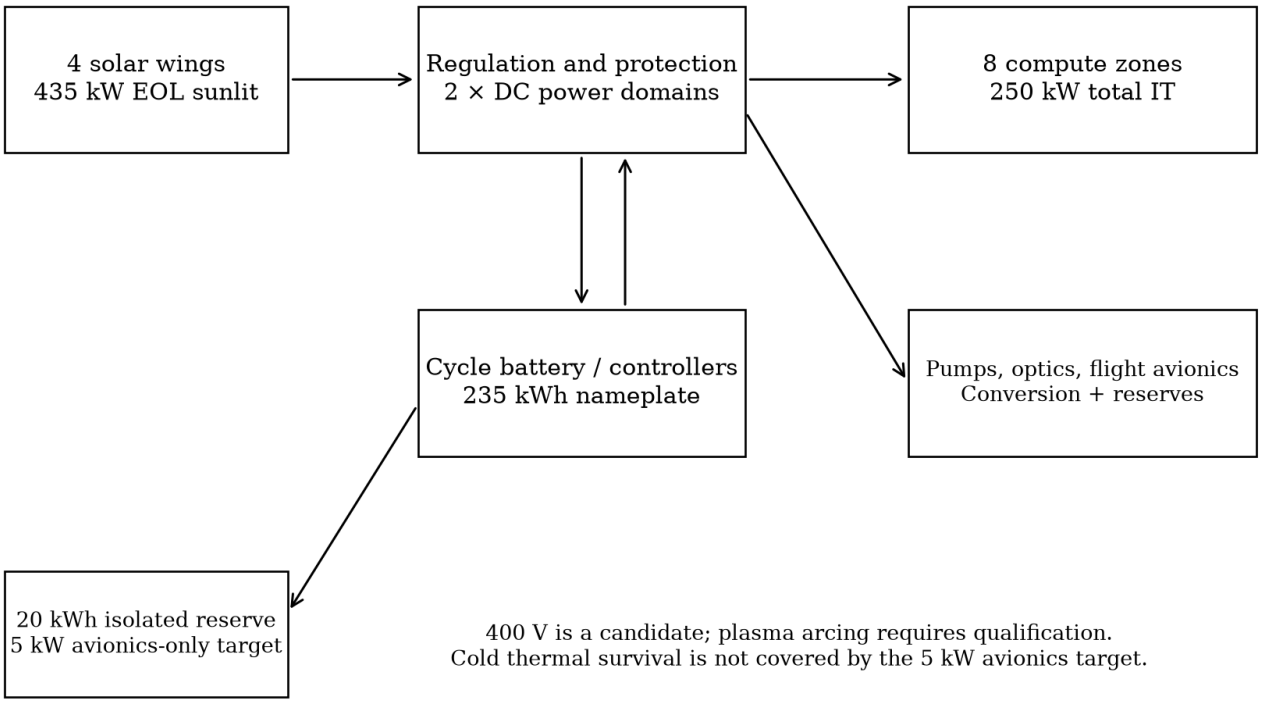


Figure 7. Proposed power architecture. Separate power domains reduce common failures; the 20 kWh reserve supports avionics recovery. The cold thermal-survival requirement is still open.

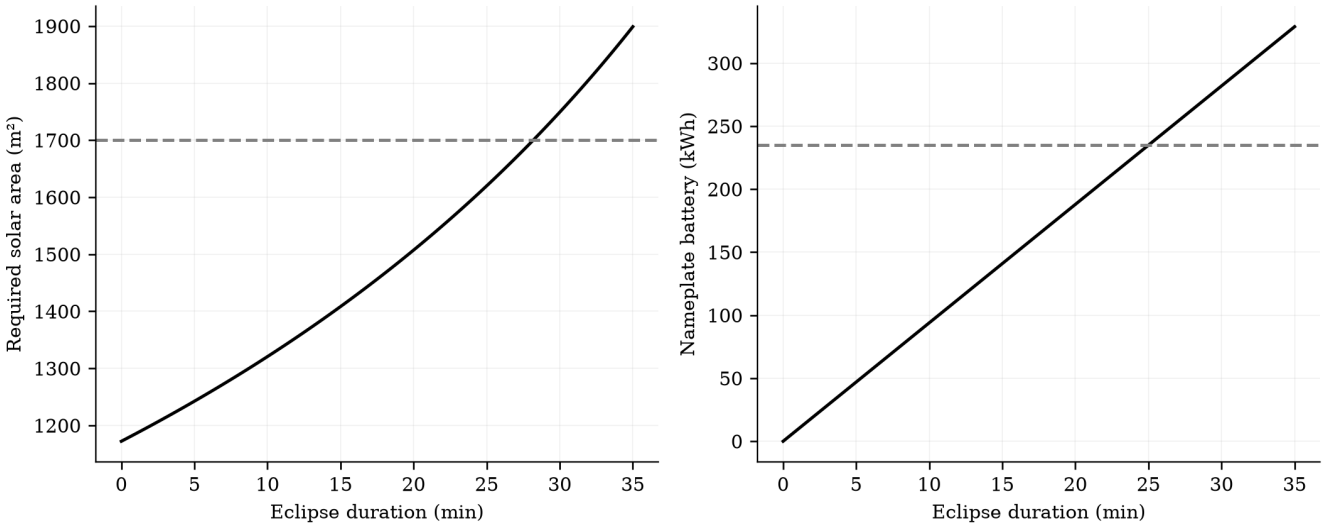


Figure 8. Required EOL solar area and BOL battery capacity versus eclipse duration, with 300 kW regulated-bus load. These curves expose the penalty for orbit and illumination uncertainty.

7. Cooling the computer without air

A radiator cannot be assigned 350 K merely because hotter radiators are smaller. With passive heat transport, the heat source must be hotter than the sink throughout every heat-transfer interface:

$$T_j = T_{fluid,\,local} + Q_{chip} R_{j \to fluid}, \qquad T_{fluid,\,local} > T_{skin,\,local}.$$

The candidate baseline is a custom single-phase, single-fluid PAO topology, subject to fluid and cold-plate qualification. It uses coolant entering the computing cassette at 315 K, leaving at 335 K, and radiator skin varying approximately from 310 to 330 K. Its effective fourth-power mean is around 320 K:

$$T_{r,\,eff} = \left(A^{-1}\int_A T_s^4\, dA\right)^{1/4}.$$

This does not mean a uniform 320 K radiator cools return fluid to 315 K. It requires a non-isothermal radiator with a cold end below 315 K. A notional 1 kW device with $R_{j\to fluid} = 0.020$ K/W and local fluid at 335 K reaches 355 K junction temperature. Both resistance and allowed junction temperature remain device-specific qualification requirements; they are not public GPU guarantees.

Supplier limits may force lower temperatures or separate branches. For example, NVIDIA's liquid-cooled Q3450 switch specifies 45°C maximum inlet and 55°C maximum return, plus PG25 coolant and an ambient operating envelope. Thus even liquid-cooled terrestrial equipment retains environmental dependencies; the switch cannot simply be placed in vacuum. NVIDIA Q3450 liquid-cooling specification. Lenovo's GB200-based SC777 V4 describes up to 45°C inlet water, but this is terrestrial product guidance and does not validate orbital operation. Lenovo SC777 V4.

All power converters, memory devices, optics and storage that originally used airflow need explicit conduction paths or cold plates. A flight cassette should expose defined thermal interfaces, bonded heat spreaders and mechanically supported boards, avoiding dependence on terrestrial rack airflow.

Eight separate circuits; one leak should remove one zone, not the spacecraft.

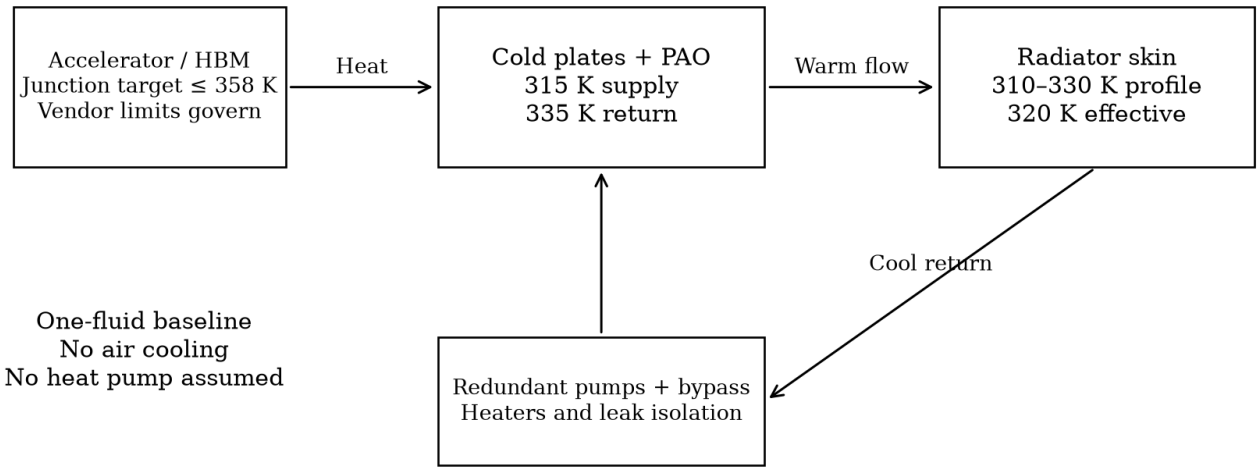


Figure 9. One of eight single-fluid loop concepts. Skin temperature varies spatially; 320 K is an effective radiation temperature. The cold end must be colder than the 315 K supply fluid.

The cited terrestrial network limits are stricter than the hottest 335 K branch return. Those devices therefore require a cooler parallel branch, more flow or a different qualified component. The diagram is a thermal interface allocation, not evidence that an off-the-shelf server can operate at these conditions. Unpressurized voltage regulators, memory packages, SSD controllers and optical electronics all need conduction paths that are absent from many terrestrial board designs.

7.1 Radiator construction and area

For one effective emitting face , define absorbed environmental flux q_{abs} consistently, including solar, albedo, Earth IR and mutual-view contributions:

$$Q_{rej} = A_r[\varepsilon_{EOL}\sigma T_{r,eff}^4 - q_{abs}], \qquad A_r = \frac{Q_{rej}}{\varepsilon_{EOL}\sigma T_{r,eff}^4 - q_{abs}}.$$

Use $\varepsilon_{EOL} = 0.90$ and conditional $q_{abs} = 100$ W/m². This absorbed-flux allowance must be verified over the entire orbit and attitude envelope; it is not a universal property of 650 km altitude. Sun avoidance, Earth view factor, plume impingement and radiator mutual shadowing drive the geometry. The reference neglects the negligible 3 K background term, not the Earth or Sun.

Effective skin temperature	**Net rejection**	**Area for 300 kW**	**Area with 15% growth**
300 K	313 W/m²	957 m²	1101 m²
310 K	371 W/m²	808 m²	929 m²
315 K	402 W/m²	745 m²	857 m²
320 K	435 W/m²	689 m²	793 m²
350 K	666 W/m²	451 m²	518 m²

A 900 m² allocated radiator at 320 K rejects approximately 392 kW under these assumptions, against a 330 kW thermal design point. At $q_{abs} = 200$ W/m² the same panels reject only 302 kW, exhausting that margin and requiring a load-budget review. Radiator redundancy cannot be asserted from area margin alone: losing a quarter of the panels leaves only 294 kW capacity at the reference environment. At 315 K, 900 m² rejects 362 kW, so an intermediate heat exchanger can be accommodated at the cost of substantially reduced margin.

Candidate panels use aluminum honeycomb or a lightweight stiffened substrate, embedded welded/brazed compatible fluid channels, high-conductivity aluminum spreaders, and a low-solar-absorptance/high-IR-emissivity coating or optical solar reflectors. CFRP supports reduce mass and thermal distortion but need qualified joints and galvanic isolation at aluminum interfaces. Coating optical properties must be measured after atomic-oxygen, UV, contamination and thermal-cycle exposure. Single-sided area is used throughout; a second radiating face must be separately assigned its own environmental views before taking credit. NASA, Thermal Control.

The 900 m² effective area is one emitting face per geometric panel. No credit is taken for the back face. Four thermal zones receive seven panels each and four receive eight, giving 60 panels total; a seven-panel zone has 105 m² and approximately 45.7 kW heat capacity in the reference environment, against 41.25 kW of allocated design heat. A failure removes its associated compute zone, rather than redistributing all heat into an undersized common loop.

The full-load envelope is constrained by the smallest independent zone, not just total radiator area. At 320 K, a 105 m² zone can reject 41.25 kW only while absorbed flux remains below approximately 142.3 W/m². The total-area calculation alone would permit 168.5 W/m² and therefore overstate the independent-loop operating envelope. At 315 K and 100 W/m², the smallest zones have only about 2.4% headroom. Exceeding these limits requires local throttling, a larger zone allocation or a different verified thermal arrangement.

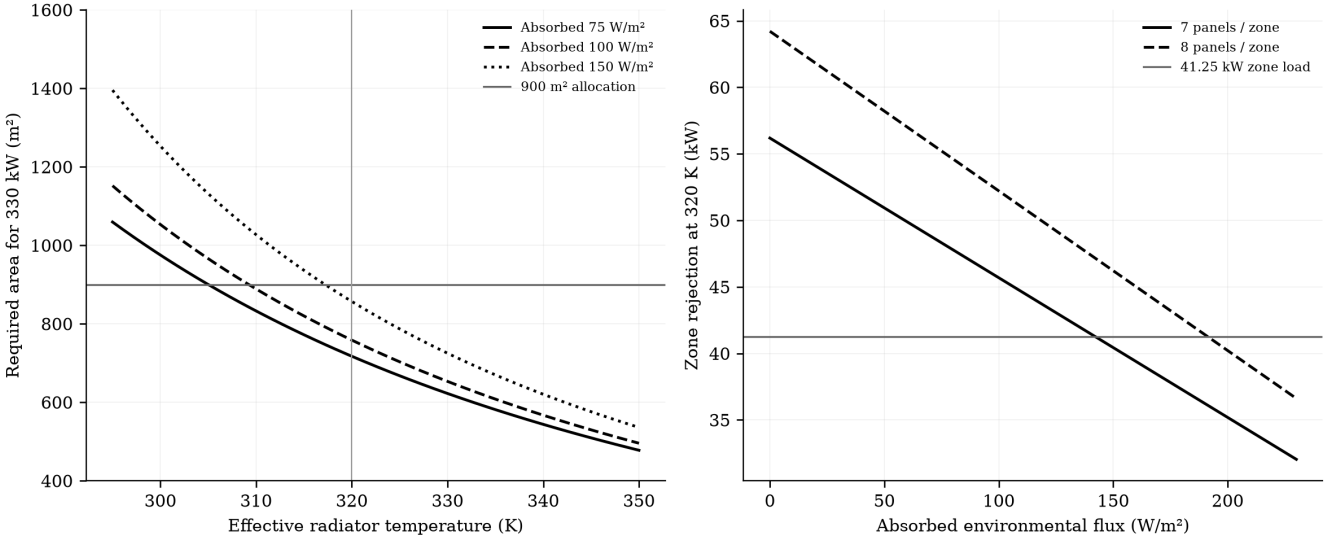


Figure 10. Aggregate required area and the independent-zone limit. The four seven-panel zones constrain full-load operation before the total 900 m² area does; local thermal protection must enforce this bound.

7.2 Fluid choice, flow and isolation

Use eight independently isolatable thermal zones, each supporting 31.25 kW of IT plus its share of auxiliaries. The cassette count per zone remains a packaging choice. Each zone has redundant pumps, pressure/temperature/flow sensing, an accumulator, bypass valve and independently isolatable radiator branches. A local leak then loses a defined capacity fraction. Pump redundancy alone does not tolerate a ruptured common manifold. Loss of one complete zone leaves approximately 218.75 kW IT; claiming uninterrupted 250 kW requires additional installed redundancy and resized budgets.

Three fluids deserve explicit trade study:

• A qualified low-freezing-point PAO formulation is a candidate for a single-fluid loop, enabling the illustrative 315-335 K coolant chain without an intermediate exchanger. Its lower heat capacity, viscosity, material compatibility, vacuum outgassing consequences after leakage, and cold-start pumping penalty must be measured. Treat approximate $c_p = 2$ kJ/(kg K) as a sizing placeholder, not a flight-qualified fluid specification.

• Ammonia has substantial spacecraft cooling heritage, including the ISS external loops. It requires compatible aluminum/stainless hardware, pressure containment, contamination controls and controlled freeze protection. At 320 K its saturation pressure is approximately 1.873 MPa; warmer liquid-only operation requires still higher pressure. It is not interchangeable with a terrestrial water cold plate. NASA ISS pump reliability; NIST ammonia properties.

• Water or qualified inhibited water/glycol gives favorable chip-cooling properties but demands freeze protection and rigorous wetted-material compatibility. A short internal loop coupled to an ammonia external loop adds exchanger mass and temperature approach. A 5 K exchanger allowance shifts the illustrative effective radiator from 320 K to roughly 315 K. At a 330 kW thermal design point and 15% area growth, required area rises from 872 to 943 m². The 900 m² allocation still rejects 362 kW at 315 K, but does not retain the full 15% growth allocation. NASA documented an Orion candidate coolant compatibility test that clogged filters within two months; naming a fluid family does not establish compatibility. NASA fluid life-test lessons.

For single-phase transport,

$$\dot{m} = \frac{Q}{c_p \Delta T}, \qquad P_{pump} = \frac{\dot{m} \Delta p}{\rho_f \eta_p}.$$

At 300 kW and 20 K rise, water with $c_p = 4.18$ kJ/(kg K) requires 3.59 kg/s total; the PAO placeholder requires 7.50 kg/s. At the 330 kW thermal design load, the PAO flow rises to 8.25 kg/s, or approximately 1.03 kg/s per thermal zone. Pressure losses, cold-plate restrictions, line diameters and redundancy determine pump power. Begin with a measured single-phase breadboard; a two-phase upgrade needs separate microgravity stability and restart validation. Safe mode must bypass/shutter radiators or heat critical loops; compute shutdown otherwise creates a severe cold-case problem.

With a sizing fluid density of 800 kg/m³, total pressure drop of 300 kPa and combined electrical-to-hydraulic efficiency 0.40, the 8.25 kg/s PAO screening flow requires 7.73 kW of electrical pump input under the simplified relation. This nearly consumes the 8 kW pump allocation. Higher viscosity, restrictive cold plates, filters, longer lines or lower efficiency require more power and a revised thermal budget. The 600 kg circuit mass allocation must also include all working fluid, redundant pumps, accumulators, valves and the three-dimensional hose routing; it has not been closed by detailed hardware.

7.3 The cold-survival problem must be solved

A shutdown of the IT load removes its heat source while the radiators still radiate. For a deliberately severe zero-absorption screen, 900 m² at 203 K with emissivity 0.9 radiates approximately 78 kW; at 250 K it radiates approximately 179 kW. A 5 kW avionics bus cannot hold the full radiator network warm. Bypassing the radiators protects the computer but leaves fluid in the external panels cold.

The project must select and demonstrate at least one workable cold strategy: freeze-tolerant circuits with verified thaw/restart; drainable or gas-displaced external sections; controllable heat rejection such as shutters; or a thermal-safe attitude using absorbed environmental heat. Each option changes mass, mechanisms and operations. A fluid's low pour point alone does not establish restart viscosity, valve behavior or fatigue life. The current mass budget carries general contingency but no closed detailed survival system. This is a major design gate.

8. Compute hardware, radiation and data integrity

Each of eight zones contains five nominal 6.25 kW trays. A tray includes accelerators, host processing, HBM or other memory, local fabric, storage interfaces and a cold plate. The 1200 kg IT allocation is approximately 30 kg per tray including shared hardware; it is not a measured product. The actual number and model of accelerator packages must follow measured performance and radiation/thermal qualification. The target product is queued terrestrial batch inference or other compute-intensive jobs with resident assets; tightly coupled training across multiple spacecraft is a later system trade.

Two radiation-tolerant flight-computer strings control the spacecraft independently of the customer payload. Customer containers have no path to propulsion, deployment or protective power control. Each compute zone has a watchdog, overcurrent/latch-up protection, resettable feeder, sensor data and a defined recovery sequence. A software restart may be routine; a destructive electrical event must be isolated before it collapses healthy branches.

Separate cumulative ionizing dose, displacement damage, recoverable single-event upsets and potentially destructive latch-up. An aluminum enclosure is not a universal radiation cure; the mission spectrum, component cross-sections and secondary-particle production determine protection. NASA describes ECC, redundancy and component-level radiation effects, but its small-spacecraft survey does not qualify a 250 kW computing platform. NASA avionics survey.

The proposed verification basis is an environment-to-device chain: propagate the orbit; estimate particle spectra and dose behind a three-dimensional shielding map; test representative production lots under proton and heavy-ion exposure; run real workloads while measuring silent corruption, hangs, destructive events and recovery time. Specify a dose margin only after selecting the mission duration and shielding

model. A provisional factor of two on predicted end-of-life dose is a design-study assumption, not universal certification.

Protect all stateful memory with available ECC and periodic scrubbing; checksum stored objects and transfers; retain verified executable images in protected nonvolatile storage; use hardware watchdogs and resettable power domains. ECC does not cover every arithmetic or logic failure. Important results require application-level consistency tests, selective duplication, or deterministic re-execution. Job accounting must report useful, verified output after recovery losses. Google reports proton testing of Trillium TPUs, including HBM sensitivity, but those results apply to its tested hardware and conditions, not every commercial accelerator or this mission. Google Suncatcher engineering study.

If checkpointing costs C seconds and recoverable failures have mean interval M, a first screening optimum is $\tau \simeq \sqrt{2CM}$, under independent exponential failures and negligible restart cost. This sets an initial software trade, not a guaranteed optimum. Common-mode radiation storms, thermal trips and communication outages require correlated-event simulation. The loss of one zone should remove at most its assigned capacity, while surviving zones and the flight computer remain functional. A hypothetical 99.9% annual spacecraft availability target corresponds to only 8.76 hours of downtime; it cannot be claimed before the measured failure and recovery budgets close.

The local shielding allocation is 300 kg. As a geometric illustration, 5 mm of aluminium across 20 m² weighs roughly 270 kg before joints. This is not a dose requirement or protection guarantee. Dose, displacement damage, single-event effects and secondary particles must be modeled behind the actual shielding geometry. Critical manifolds and avionics can receive localized particle/debris protection; blanket shielding of 2600 m² of exposed wings is not included. NASA debris-damage modelling treats material, spacing and geometry jointly.

Payload reset cannot reset guidance, survival power or the disposal controller.

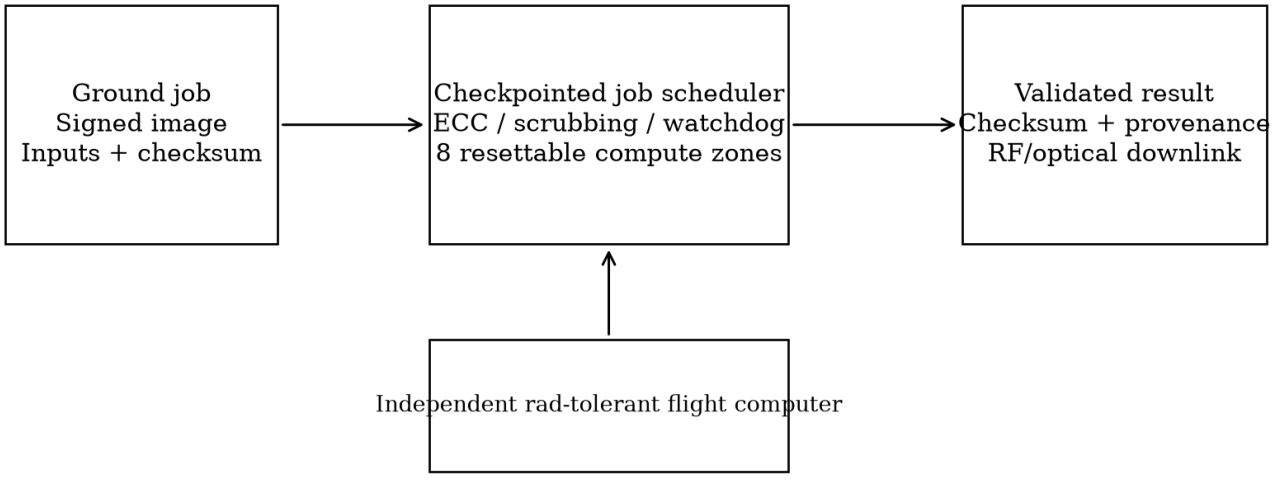


Figure 11. Workload execution and fault containment. Checksums protect transfers and stored objects; application checks are still needed for radiation-induced computation errors.

9. Serving customers on Earth

Provide two independently steerable optical terminals for gateway handover and resilience, with a proposed 100 Gbit/s usable burst-rate objective per active link at 1550 nm. Add independent RF telemetry, tracking and command, capable of bootstrapping and safe-mode control without any optical terminal. A broadband RF fallback is a separately costed service option; an RF command channel cannot carry normal customer traffic at optical rates. NASA TBIRD demonstrated 200 Gbit/s downlink and 4.8 TB delivered in five minutes; this supports high-rate burst feasibility, not symmetrical 100 Gbit/s uplink, continuous coverage or our pointing allocation. NASA TBIRD results.

The usable throughput must be calculated from contact geometry and correlated weather:

$$\bar{B} = \frac{1}{T_m}\int_0^{T_m} \sum_j b_j(t) v_j(t) w_j(t) \eta_j(t)\, dt.$$

Here v_j denotes visibility and scheduling, w_j weather availability, and η_j net coding/protocol efficiency. Summation is constrained by simultaneous terminal and ground-backhaul capacity. Six geographically diverse gateways are an initial coverage-study case, not a guarantee. Local clouds are correlated; multiplying independent availability estimates would exaggerate resilience. NASA LCRD explicitly uses ground-site selection and atmospheric monitoring to manage optical interruptions. NASA LCRD ground segment.

For an explicitly illustrative, workload-dependent sustained rate $F = 10^{17}$ useful FLOP/s (numerical precision and workload not yet qualified), an average uplink of 20 Gbit/s implies at least

$$I_{min} = \frac{8F}{\bar{B}} = 4 \times 10^7 \text{ FLOP/byte}$$

of computation per new input byte if the accelerator is to remain busy. Reusing resident model weights and data matters more than peak link rate. Even at an uninterrupted 100 Gbit/s, uploading 1 TB takes 80 s before retransmissions. A satellite directly overhead adds only $650\,\text{km}/c = 2.17$ ms one-way propagation; real job latency also includes ground routing, pass acquisition, queueing, transfer, processing and return.

A proposed 1 PB usable object store accommodates job queues, resident assets and recovery copies, but its final capacity follows the contracted workload and outage distribution. At 20 Gbit/s, one hour without egress accumulates 9 TB if output continues at that rate. Storage consumes power, dissipates heat and has radiation failure modes; it belongs in the IT allocation. Encrypt transit and stored customer objects, separate tenants, verify signed workloads and rotate keys through the protected management domain. Laser directionality alone does not provide customer security.

No access to another operator's intersatellite relay network is assumed as a free, automatic interface. A relay architecture requires an actual terminal interface, optical compatibility, routing service, bandwidth reservation and commercial agreement. The alternate direct-to-ground architecture needs its own gateways and contact scheduling. The 6 kW communication allocation is a starting point for terminal procurement; gateway stations and ground backhaul are separate infrastructure.

A proposed optical receiver is housed behind a contamination-controlled aperture with a coarse gimbal and fine-steering mirror. It is mounted through an optical bench and vibration-isolation path rather than directly on a flexible wing. Solar exclusion angles, Earth-limb stray light, body obstruction and thruster plumes restrict its field of regard. Two terminals permit handover and fault tolerance but do not guarantee continuous coverage from one orbiting vehicle.

10. Attitude control, propulsion and end of life

Large solar and radiator wings demand a coupled flexible-body control model. Use redundant star trackers, inertial sensors and coarse Sun sensors, a four-wheel redundant attitude-control set as the initial trade, and thrusters plus magnetic unloading where effective. Wheel torque and momentum capacity remain unselected until disturbance and modal analyses close. Control-moment gyroscopes may provide higher maneuver authority, but introduce gimbal singularities, vibration and failure modes; they cannot eliminate the need to reject accumulated external momentum. NASA describes momentum unloading and the challenges of flexible-spacecraft control. NASA spacecraft attitude control; NASA flexible-structure control research.

Solar-pressure torque illustrates the scale:

$$\tau_{SRP} \simeq \frac{G_\odot}{c} C_R A_{proj} \ell .$$

With $G_\odot = 1361$ W/m^2, $C_R = 1.5$, $A_{proj} = 1700$ m^2 and a 1 m center-of-pressure offset, $\tau \simeq 11.6$ mN m; a persistent 10 m offset produces 116 mN m and about 679 N m s per orbit. This is a screening case for the revised array scale; changing wing geometry and attitude change the result. Atmospheric drag adds $\tau_D \simeq \frac{1}{2}\rho v_{rel}^2 C_D A_{ram} \ell_D$, whose magnitude must use a high-solar-activity density case and the actual ram-facing projected area. Solar-facing area and ram-facing area cannot be substituted indiscriminately. Differential thermal distortion also moves centers of pressure and changes optical alignment.

At approximately 93 m solar span, 66 m radiator span and 21.6 tonnes wet mass in the rounded system allocation, this is a large flexible spacecraft. Do not import the momentum-wheel selection or controller gains of a compact 5 tonne bus. Choose low body-slew rates, balanced wings, and coarse Sun-pointing that minimizes persistent disturbance torque. Size wheel or CMG momentum capacity, magnetic dipole and thruster unloading together; magnetic torque is perpendicular to the geomagnetic field and cannot provide arbitrary instantaneous unloading. Optical terminals require independent coarse gimbals and fine-steering loops, rather than slewing the entire platform for every ground pass. Modal testing must verify that wheel disturbances, deployment oscillations, pumps and thruster firings remain inside the optical jitter budget. Thruster placement and firing constraints must prevent plume impingement on solar cells, radiator coatings and optical apertures, while keeping the line of thrust compatible with changing center of mass. Collision avoidance and desaturation can temporarily interrupt optical service, and that interruption must appear in the service budget.

For disposal from a 650 km circular orbit to an ellipse with 120 km perigee, the ideal retrograde impulse is

$$\Delta v = \sqrt{\frac{\mu}{r_a}} \left[1 - \sqrt{\frac{2 r_p}{r_a + r_p}} \right] = 149.0 \text{ m/s}.$$

This is an energy calculation, not a certified controlled-entry trajectory. A preliminary total allocation of 250 m/s can reserve 150 m/s for disposal and 100 m/s for injection correction, orbit control, collision avoidance and uncertainty, pending a mission-specific budget. The rocket equation gives $m_p = m_d[\exp(\Delta v / g_0 I_{sp}) - 1]$. As a rounded illustration, for dry mass $m_d = 19000$ kg and 250 m/s, $m_p = 2334$ kg at 220 s or 1576 kg at 320 s. An ideal 1500 s electric system needs only 326 kg, but low thrust changes maneuver duration and cannot automatically provide a targeted final reentry. Retain a high-thrust disposal capability unless a validated alternative closes the casualty-risk analysis. Disposal propellant must be protected from routine consumption; tanks, residuals and margins must be reconciled with the root mass budget rather than counted twice.

Provide autonomous collision-response logic, external conjunction screening, post-maneuver tracking, passivation and a standardized removal interface. Large batteries, tanks and dense compute assemblies require explicit breakup and ground-casualty modeling. ESA's own policy specifies disposal and reentry criteria; the applicable licensing authority and mission obligations must be established separately. ESA reentry safety; ESA debris mitigation.

The root budget uses 19.23 tonnes of dry mass including contingency and a 220 s specific-impulse allocation, giving 2.36 tonnes of propellant for 250 m/s. The 300 kg dry-propulsion allocation includes tanks, lines, thrusters, regulators and controls only; it needs supplier sizing. Titanium or qualified composite-overwrapped pressure-vessel options must be selected for the actual propellant, pressure, life and fracture-control requirements. The 250 m/s allowance must not be counted again as a separate propellant reserve.

For a spacecraft of this mass, a low perigee is not sufficient evidence of a safe disposal plan. A licensed controlled-entry analysis must determine breakup, surviving components, trajectory dispersions, failure modes and casualty risk. Batteries, coolant and propulsion vessels require passivation. Loss of customer compute must leave enough flight control and protected propulsion to carry out the planned disposal or a validated contingency.

11. Mass accounting and margins

The table below is generated from the same model as the figures. Solar mass uses 625 kW BOL reference power divided by an assumed 150 W/kg complete-array specific power. Redwire's published ROSA configurations imply about 81-124 W/kg for the listed BOL configurations; 150 W/kg for much larger wings is an engineering target, not a vendor offer. Radiator mass is an assumed 6 kg/m² including panel and local deployment hardware, with main circuits budgeted separately. A NASA deployable-radiator development project describes a below-8 kg/m² development target; it does not establish the proposed panels as available flight hardware.

Subsystem	CBE mass, kg	Basis/status
Compute trays, cold plates, local network	1,200	Preliminary allocation
Primary structure and deployment roots	1,500	Preliminary allocation
Solar wings at 150 W/kg BOL	4,165	Preliminary allocation
Radiator panels and local deployment	5,400	Preliminary allocation
Pumps, manifolds, fluid and piping	600	Preliminary allocation
Cycle and 20 kWh avionics-reserve batteries	1,500	Preliminary allocation
Power conversion, protection and harness	500	Preliminary allocation
Flight avionics and attitude hardware	220	Preliminary allocation
Optical and RF terminals	220	Preliminary allocation
Local shielding allocation	300	Preliminary allocation
Insulation and distributed heaters	120	Preliminary allocation
Dry propulsion hardware	300	Preliminary allocation
Dry CBE subtotal	16,024	Before system contingency
Dry contingency	3,205	20% of dry CBE
Dry design mass	19,229	Basis for propellant sizing
Propellant	2,362	250 m/s, 220 s ideal rocket equation
Wet spacecraft	21,592	Excludes launch adapter
Launch adapter	400	Interface allocation
Injected allocation	21,992	Approximately 22.0 tonnes

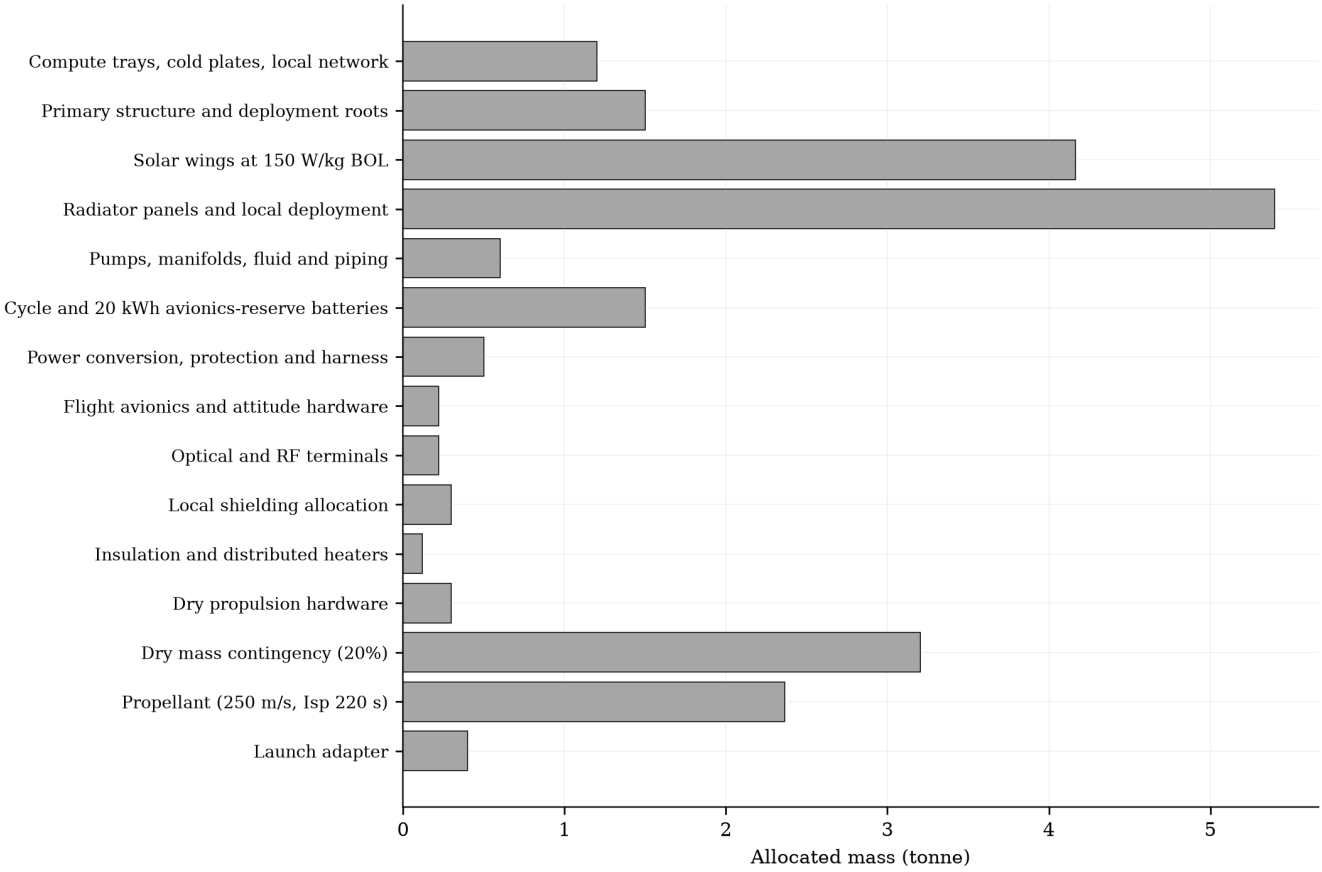


Figure 12. Component mass, dry contingency and propellant are separate. The largest hardware terms are the solar arrays and radiators. Contingency is an allowance, not a statistical confidence interval.

A heavier array at 100 W/kg BOL adds approximately 2.08 tonnes before dry contingency and propellant. Radiators at 8 kg/m² add 1.80 tonnes before those same multipliers. Conversely, 250 W/kg arrays and 4 kg/m² radiators would save several tonnes, but neither performance should be booked without representative structures and deployment evidence. The principal engineering objective is reducing mass while preserving end-of-life power, heat transport, launch survival and failure containment.

The 20% dry allowance cannot absorb every unbounded risk. Cold-survival hardware, larger launch restraints, extra radiation shielding or a more complex radiator-pointing system could consume it. A proper mass-control process assigns owner, maturity, uncertainty and growth allowance to each line; the single percentage here is a screening convention. Likewise, a 22-tonne payload mass says nothing about its available launch date, SSO delivery capability or price.

12. How it would be manufactured and assembled

Manufacture proceeds from traceable metal stock and qualified laminates to machined parts, inspected bonds, panel-level wiring/fluid assembly, subsystem acceptance, bus integration and system verification. Every radiator circuit receives documented cleanliness, pressure and leak verification. Every bonded load path needs process coupons and non-destructive inspection appropriate to the failure mode. Every deployable package receives controlled packing, restraint-preload records and a traceable release configuration. A mass-properties measurement updates the structural and attitude-control models before flight.

The proposed verification chain comprises material coupons; structural/thermal development articles; one representative powered compute module; a full-scale wing and radiator deployment article; an integrated structural/thermal model where justified; and qualification or protoflight plus flight acceptance testing under an agreed programme. Thermal-vacuum operation must reproduce

representative dissipation, coolant flow, startup, shutdown and injected failures, not merely survival of unpowered equipment. Modal testing correlates the stowed structural model; deployed dynamics testing supports optical pointing and control stability.

GEVS supplies a recognized framework for environmental verification, but is not automatically a mandatory certification for this private project. NASA GSFC-STD-7000B. ECSS distinguishes qualification, acceptance and protoflight approaches, with tailoring to project needs. ECSS-E-ST-10-03C Rev.1. At this spacecraft size, shaker capacity and over-testing risk require an agreed combination of integrated acoustic testing, modal/sine tests and equipment-level random vibration; the Falcon guide specifically discourages integrated random-vibration shaker testing of medium/heavy payloads.

Before a preliminary design review, the project must produce a three-dimensional folded/deployed layout; launcher interface model; correlated mass and stiffness model; materials/processes register; structural and thermal margins; deployment/failure analysis; radiation and debris allocations; verified electrical/fluid interface definitions; and a funded verification plan. These documents convert a visual spacecraft concept into a design that suppliers can quote and test.

Proposed production and integration route

• **Build a 6.25 kW representative compute tray.** Select accelerators and all supporting components; design cold plates, conduction paths, connectors and a mechanically restrained board stack. Run real workloads while measuring useful work, temperatures and power, then repeat under vacuum and radiation exposure.

• **Build one complete 31.25 kW thermal/electrical zone.** Five trays, an independent feeder, redundant pumps, valves, accumulator and representative radiator channels form the minimum useful system article. Demonstrate controlled leakage isolation, pump failure, restart and brownout recovery.

• **Fabricate and qualify radiator panels.** Produce traceable aluminium substrates, compatible channels, bonded/brazed interfaces and optical finish. Verify cleanliness, pressure integrity, leakage, channel flow distribution, coating properties and micrometeoroid sensitivity before panel integration.

• **Build array blanket and deployment demonstrators.** Integrate protected cells, interconnects, grounding and blanket substrate with composite booms. Test roll packing, long-axis support, hinge/cradle motion, electrical continuity and hot/cold deployment after launch-like vibration.

• **Assemble the primary frame and pressure hardware.** Machine and inspect load-bearing parts; install decks, reinforcement and attachment inserts. Assemble tanks and propulsion lines under their selected pressure-system verification plan. Protect optical and fluid interfaces from contamination during parallel work.

• **Integrate subsystem harnesses and loops.** Route high-voltage conductors away from sensitive measurement and optical hardware; control bend radii, bonding, strain relief and fluid/harness interaction through every moving joint. Record as-built mass and center of gravity after each major installation.

• **Pack, test and commission the spacecraft.** Validate releases and latch telemetry; perform representative structural, acoustic, modal and thermal-vacuum verification with a correlated test strategy. Ground tests need gravity offload for wings and a justified thermal equivalent for full-area heat rejection. A small thermal chamber cannot validate a 900 m^2 deployed radiator by assertion.

13. Requirements that must be verified before design freeze

The following are proposed design-review gates. Where a numerical value is still absent, the deliverable is the analysis and test that sets it. Quotation-ready requirements should be approved only after the hardware, launcher and operating orbit are selected.

ID	Proposed requirement	Verification evidence
SYS-01	Sustain 250 kW IT within 300 kW regulated-bus allocation	Full-zone electrical tests and integrated budget
PWR-01	Bridge 25 min at 300 kW after EOL capacity derating	Battery cycling, hardware-in-loop orbit replay
PWR-02	Provide required recharge from 1700 m^2 at worst modeled solar condition	Supplier cell/array data and time-resolved power model
PWR-03	A short in one feeder cannot collapse both main domains	Arc/short protection and fault-injection tests
THM-01	Reject 330 kW within verified coolant and component limits	Correlated thermal-vacuum/thermal-balance tests
THM-02	Demonstrate 315 K cold-end supply with nonuniform radiator	Full-path conductance and flow validation
THM-03	Survive IT shutdown and restart external circuits	Cold-soak, freeze/thaw or chosen survival-strategy test
THM-04	Isolate a punctured branch without losing healthy zones	Representative leak and valve-failure tests
MEC-01	Meet actual launcher coupled loads and keep-in envelope	Agreed interface document, structural and modal correlation
MEC-02	Deploy each wing without collision or fluid/harness overstress	Swept-volume CAD and full-scale gravity-offloaded testing
MAT-01	Materials maintain required properties after the mission environment	Traceable allowables, exposure coupons and process qualification
IT-01	Returned results satisfy application correctness requirements	Radiation tests under workload and recovery validation
GNC-01	Maintain Sun pointing and optical jitter within selected terminal limits	Coupled flexible-body model, disturbance and pointing tests
COM-01	Sold job volumes and deadlines fit actual contacts and weather	End-to-end gateway/contact simulation and link demonstrations
SAFE-01	RF command and flight control survive payload shutdown	Independent reserve-power and safe-attitude acquisition tests
EOL-01	Preserve verified disposal capability throughout life	Propellant reserve, gauging, reentry and failure analyses

14. Development path and what remains unresolved

The reference concept is sufficiently defined to commission subsystem trades and representative hardware. It is not ready for a flight production order. Five issues determine whether it becomes practical: low-temperature radiator mass; freeze-safe fluid deployment and restart; actual EOL solar specific power; very-large flexible-wing dynamics and optical pointing; and the launch/deployment interface.

A 10-25 kW orbital technology demonstrator should carry a representative compute tray family, fluid circuit, coating, deployment joint and communication path. It can measure radiation recovery, useful FLOP per joule, real thermal resistance and operations. It cannot validate a 93 m solar wing or a sixty-panel radiator mechanism solely by flying a smaller version. Full-scale deployment and structural development articles must proceed in parallel, and the planned 2029-2030 flight date depends on their results and an actual launch contract.

The first ground demonstrator should be one electrically and thermally complete zone. This resolves the highest-value question quickly: whether a real accelerator stack, its memory/network electronics, candidate fluid and radiator panel can run the intended customer workloads at the assumed temperatures. If that chain requires a materially colder radiator or a heavier fluid/pressure system, the area, mass, power and spacecraft geometry are recalculated before a factory or full spacecraft is funded.

The design also creates a clear scaling choice. Eight independent smaller vehicles might remove one large deployment challenge but duplicate buses, optical terminals and propulsion. A single larger vehicle reduces that duplication while concentrating failure and launch risk. The present document specifies the one-vehicle 250 kW case so that the trade has an explicit comparator; it does not establish a global optimum.

Appendix A. Reproducibility and engineering status

The accompanying `design_model.py` generates the twelve figures, `design_budget.json` and `mass_budget.csv`. It uses Python, NumPy and Matplotlib. Vector SVG and raster PNG versions are included. The model checks energy closure, the seasonal eclipse range, radiator heat capacity, battery usable energy, mass accounting and the disposal impulse. Numerical precision in the ledger is retained for reproducibility, not as an accuracy claim.

The equations are first-order and deterministic. No structural finite-element model, ray-traced thermal model, detailed fluid network, radiation transport, weather/contact optimization or Monte Carlo reliability simulation has yet been produced. The cold-survival architecture and deployment swept volumes remain explicitly open. The 22-tonne mass and chosen geometry are proposed engineering allocations, not measured performance or a quoted spacecraft specification.

The reference constants use equatorial Earth radius 6378.137 km; this differs slightly from the mean-radius simplification used in the earlier conceptual paper. The solar and battery degradation factors are five-year sizing assumptions pending mission-specific supplier data. All source links in the body identify technical heritage or verification frameworks. None qualifies the complete OC-250 spacecraft by extension.