The physical form of an orbital data centre

Architecture selection, coupled sizing and two-dimensional concept definition

250 kW Earth-serving reference mission. Design horizon: 2029–2030. Revision C, 25 September 2026. Units: SI unless a drawing explicitly states millimetres. This is an original analytical engineering study with reproducible calculations. It is not an institutional publication or a manufacturing release.

Abstract

The leading architecture in this study is a compact planar mosaic of photovoltaic tiles and separate, two-sided radiator tiles carrying distributed computing equipment. Its rationale is physical: solar collection rewards projected area, heat rejection rewards exposed emitting area, and moving the computers to their heat-rejection surfaces shortens thermal transport. Neither a conventional central spacecraft bus nor an exotic silhouette should be imposed before these functions are jointly sized. A partially coupled photovoltaic/compute sandwich remains a serious competing architecture, subject to measured panel mass, interlayer conductance, local package temperature and cold-transient performance. A circular outline is not inherently thermodynamically superior to a rectangle at identical surface normals.

The comparison evaluates eight principal topology families, analytical projection and radiation bounds, 135 deterministic geometric scenarios, heterogeneous launch packing, network topology and module duplication. A provisional four-spacecraft reference allocates 62.5 kW of IT per spacecraft on a 24 × 28 m tile grid. Across the fleet it provides 1,760 m² of PV and 800 m² of two-sided radiator planform. At 320 K, the radiator allocation rejects 341.8 kW in the most demanding combination of the stated 48 optical-property, environmental and emitting-area cases, compared with a 330 kW design load. That is a 3.57% corner margin, not the much larger nominal margin. The compact form is selected for the next development stage; a unique minimum-cost flight design is not established. Connected deployment, dynamics, cold survival, radiation-qualified compute, complete service hardware and full launch mass remain open design variables.

1. The decision and its accounting boundary

The engineering question is not what a satellite normally looks like. It is which distribution of collecting area, heat sources, emitting area, structure and interconnect delivers useful computing with the smallest lifecycle burden while satisfying deployment and operational constraints. The mission is Earth-serving batch inference and other separable work. Processing only data generated by other satellites is not assumed. Tightly synchronized model training is a separate workload that may reverse the preferred fragmentation and network topology.

Three decisions must remain distinct: the functional topology of a tile, the outline of a deployed spacecraft, and the number of autonomous spacecraft. A flat surface may contain a central computer with remote radiators, separate PV and compute/radiator tiles, or a fully integrated sandwich. Those are different thermal and electrical systems even if their exterior silhouettes are identical. Similarly, four independent vehicles are not merely four quarters of one vehicle: they duplicate control, protection, communication and disposal functions.

The previous central-core drawing set is retained as a reference, not as a constraint. Its approximately 93 m PV span, 66 m radiator span and long liquid paths create avoidable coupling between deployment and cooling. Its revised launch-mass estimate was approximately 22.46 t, with important carrier and thermal-survival items still unclosed. It would be misleading to compare that complete-but-incomplete system estimate directly with the bare surface mass of a new panel. This revision therefore compares like accounting boundaries and explicitly leaves unmeasured terms symbolic.

Quantity	Controlled study value	Interpretation
Useful IT electrical load	250 kW fleet total	Architecture sizing target, not measured throughput
Operating main-bus load	300 kW	IT plus 50 kW total service/conversion allocation
Routed thermal design load	330 kW	Conservative 10% allowance over operating bus load
Circular reference altitude	650 km	Ideal dawn-dusk Sun-synchronous geometry
Battery design eclipse	25 min	Retained despite favourable illumination
Required sunlit bus power	414.582 kW	Inherited recharge and operating requirement
Available generation target	435.311 kW	5% headroom above that requirement
Radiator skin screening ceiling	320 K	To be reconciled with local component limits
PV temperature screening ceiling	373 K	Study constraint, not a supplier rating
Tile grid pitch	2,000 mm	Includes future seams within the allocated cell
Launch allocation cylinder	6.5 m diameter × 12 m height	Geometric study envelope, not a verified launcher interface

The optimization variables are collected in a design vector containing topology, outline, tile pitch, platform count, thermal conductance, emitting fraction, bus voltage, joint design and material architecture:

$$\mathbf{x} = \{z, A_P, A_R, \chi, a, N_s, g, f_e, V_{DC}, \mathbf{m}, \mathbf{j}\}.$$

A defensible objective is vector-valued until the commercial mission fixes priorities:

$$\min_{\mathbf{x}} \{M_{launch}, C_{life}, L_{deployment}, P_{loss}, t_{job}\},$$

$$P_{available} \geq P_{target}, \qquad Q_{reject} \geq Q_{design}, \qquad T_{j,i} \leq T_{j,i,limit}.$$

The deployment, stress, modal, collision and failure constraints are equally real, even though this model cannot yet evaluate them at flight-design fidelity. Samples that pass a geometric proxy are not automatically feasible spacecraft. A weighted aesthetic score is not used to announce an optimum.

2. Shape families and reasons to retain or reject them

Eight principal families are compared in drawing OC250-100. Sphere and tether limits supplement that set. The architecture shortlist is deliberately wider than conventional wings around a box.

Family	Physical attraction	Decisive unresolved issue
Central core with separated wings	Concentrated electronics; distinct PV and radiator temperatures	Long heat transport and coupled wing deployment
Integrated PV/compute/radiator plane	Shared footprint, possible structural sharing	Required thermal coupling and all-in surface mass
Mixed PV and compute/radiator plane	Local heat rejection; different tile temperatures	Many interfaces and thick-tile deployment
Fleet of compact mixed planes	Shorter span, independent failure domains	Duplicated buses, links and operating power
Long ribbon or chain	Repeated modules and simple one-dimensional adjacency	Long-span dynamics, serial deployment and low network bisection
Integrated disk	Minimum perimeter at fixed planar area	Wedge geometry, finite-thickness folding and same thermal coupling problem
PV disk with edge-on radiator skirt	Collector and radiator normals can differ	Cylindrical self-view, rim structure and deployment
Curved PV cylinder	Potential shell stiffness and attitude tolerance	Collection projection penalty and packaging

For an arbitrary photovoltaic surface, the usable collecting geometry is its Sun-normal projection:

$$A_{proj} = \int_{A_P} \max(0, \mathbf{n} \cdot \mathbf{s})\, dA.$$

A tracked plane gives one unit of material area per unit projected area. A fully populated sphere gives four. A hemisphere populated only on its illuminated half gives two. A fully populated cylindrical side, with its axis perpendicular to the Sun, gives a ratio of pi. These are geometric results under equal cell properties, not full-system mass rankings. A cylinder with a flat PV cap avoids that particular collection penalty and must be assessed as a different architecture.

$$\frac{A_{PV,\,sphere}}{A_{proj}} = 4, \qquad \frac{A_{PV,\,cylinder}}{A_{proj}} = \pi.$$

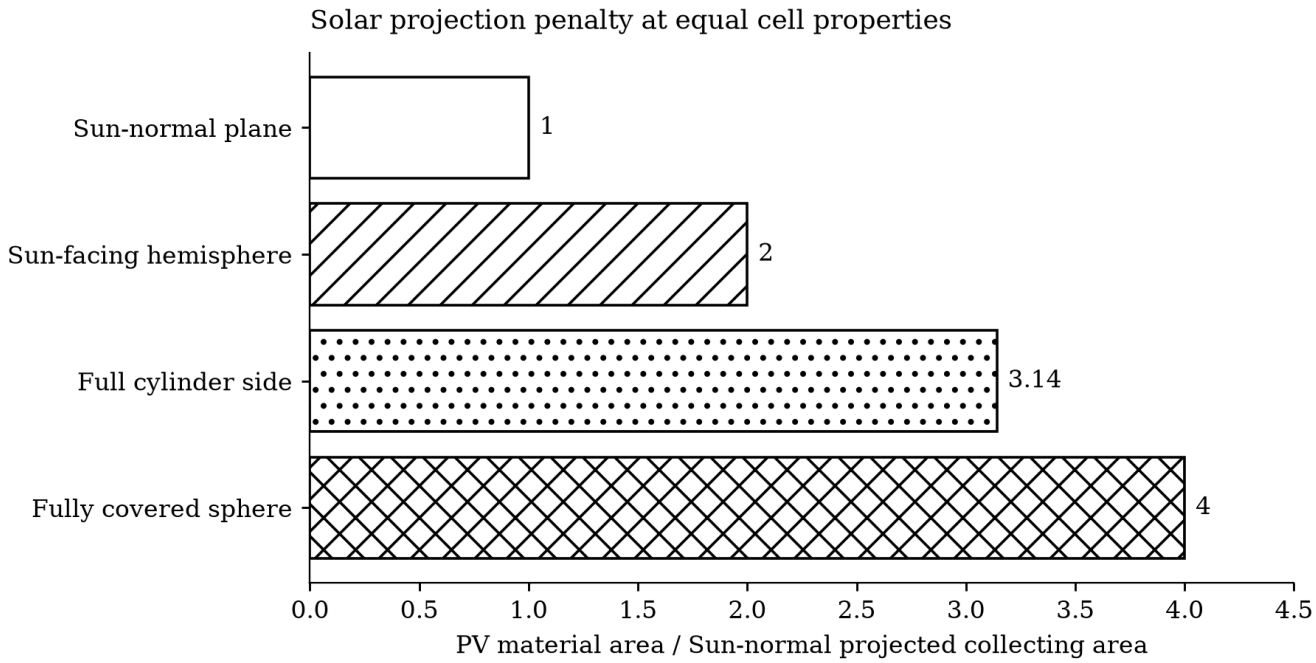


Figure 1. Exact collection-area penalties for the specified ideal geometries. Curvature may buy stiffness or attitude tolerance, but that benefit must pay for the additional PV material and supporting system.

A disk and a square with the same exposed area and surface normals have the same ideal radiative capacity. The disk reduces perimeter by approximately 11.4% relative to a square. For uniform loads served from a central point, its mean squared route distance is only 4.51% smaller. This is too small an advantage to predetermine the answer before folding, manufacture and service routing are designed. A cap-and-skirt can improve radiator solar incidence, but it does not make its inner surface an unobstructed radiator; the skirt sees itself and the cap. It remains a research alternative until a full view-factor and deployment model exists.

Research on distributed orbital computing includes clustered spacecraft, integrated panel concepts and tethered chains. These establish alternatives worth testing, not qualified performance. Google's Project Suncatcher discusses closely spaced compute satellites and optical connectivity; its system priorities are not interchangeable with independent batch-processing modules. An April 2026 integrated-panel preprint explores co-located solar, compute and radiation surfaces. Its projected thin-panel mass and illumination assumptions are not adopted here. Google system study, integrated-panel preprint, tether concept.

3. Thermal geometry determines the useful form

3.1 Radiation to space is a surface-exchange problem

At these dimensions, vacuum convection contributes no useful heat rejection. A surface emits according to its temperature and infrared emissivity, while absorbing direct solar, reflected solar and terrestrial infrared loads. A radiator's solar absorptance and infrared emissivity belong to different wavelength bands. Treating them as one black-body parameter can materially distort the result. NASA's thermal-control survey documents this distinction and the separate roles of coatings, thermal straps, heat pipes, active loops and variable heat rejection. It does not qualify the proposed tile. NASA thermal-control survey.

$$Q_{net} = \sum_i A_i \left[\varepsilon_i \sigma T_i^4 - q''_{solar,i} - q''_{albedo,i} - q''_{IR,i} - q''_{mutual,i} \right].$$

For a small diffuse surface above a spherical Earth, the directional Earth view factor is

$$F_E(\mathbf{n}) = \frac{1}{\pi}\int_{\Omega_E} \max(0, \mathbf{n}\cdot\mathbf{k})\, d\Omega.$$

With Earth radius 6,371 km and altitude 650 km, the apparent Earth angular radius is 65.164 degrees. Independent analytic checks give a nadir view factor of 0.82358, an edge-on factor of 0.24069 and an anti-nadir factor of zero. The numerical model integrates rays intersecting the Earth, then evaluates whether the intersected ground normal faces the Sun. Albedo is therefore phase-dependent rather than a constant heat load appended to every attitude.

$$F_{nadir} = \sin^2\alpha, \qquad F_{edge} = \frac{\alpha - \sin\alpha\cos\alpha}{\pi}.$$

$$q''_{env} = \alpha_s G \max(0, \mathbf{n}\cdot\mathbf{s}) + \varepsilon M_E F_E + \alpha_s a_E G F_a.$$

The nominal case uses 237 W/m² Earth outgoing longwave exitance, albedo 0.30, radiator emissivity 0.90 and solar absorptance 0.15. The annual screening model correlates solar irradiance with ideal dawn-dusk orbit geometry rather than pairing arbitrary orbital phase and season. It samples 24 seasons and 13 half-orbit phases, rejecting occulted direct-Sun states. This is a low-order spherical, diffuse-Earth model; it is not a high-resolution operational environment model. The view-factor formulation follows standard radiation exchange practice. ECSS thermal design handbook.

At 320 K, one clean surface with emissivity 0.90 emits 535.12 W/m². A two-sided coplanar radiator receives direct sunlight on its front face and sees Earth through both hemispheres. The worst nominal absorbed load in the aphelion screen is approximately 330.85 W per square metre of planform. The clean two-face net flux is therefore approximately 739.39 W/m², requiring 446.31 m² for 330 kW. This is an area bound before blocked emission, gradients, losses and reserve.

$$A_{R,\,ideal} = \frac{330000}{2(0.90)\sigma(320)^4 - 330.85} \simeq 446.31\ \mathrm{m}^2.$$

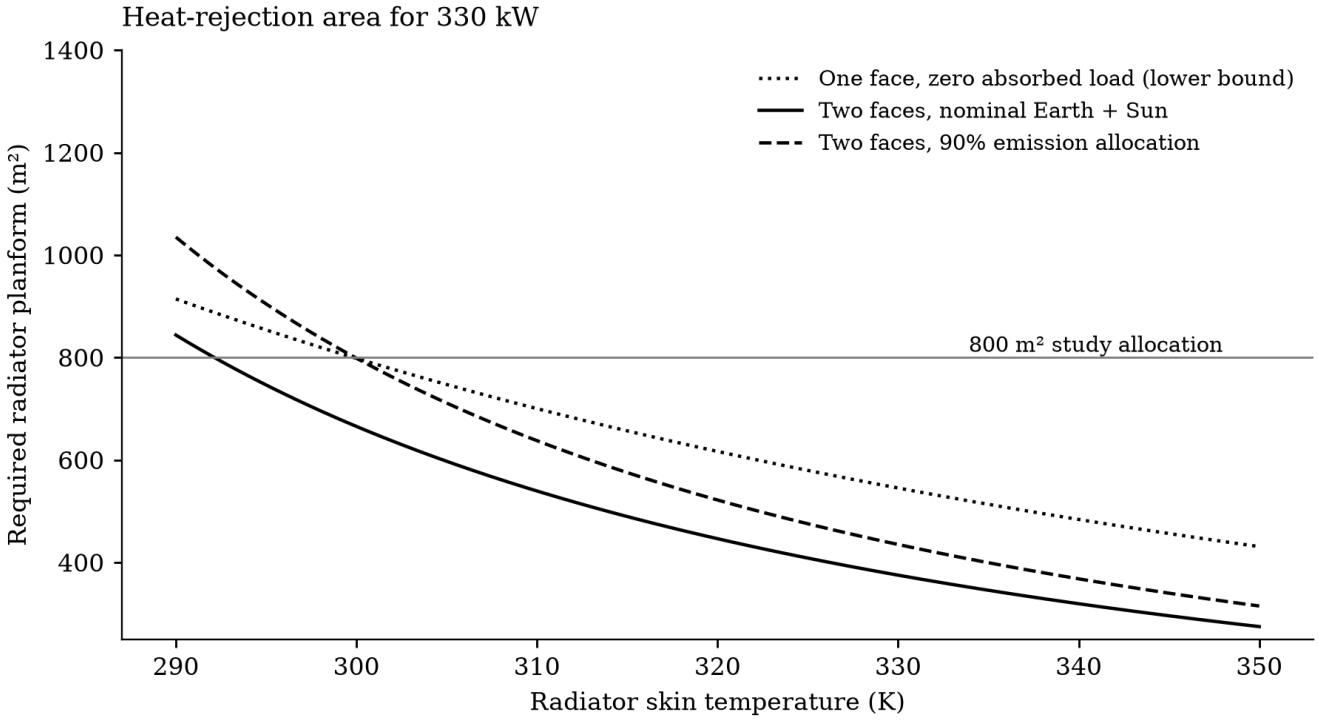


Figure 2. Area depends strongly on usable faces and permitted skin temperature. The one-face zero-absorption curve is an ideal lower bound for that one-face case, not a common orbital attitude shared with the other curves.

3.2 Emitting area must survive packaging and degradation

Let the usable emission fraction account for blocked surface, uncredited box faces, seams and other view losses. Conservatively retaining absorption over the gross radiator planform gives

$$Q_R = A_R\left[2f_e\varepsilon\sigma T_R^4 - q''_{env,\,gross}\right].$$

This reduction is an allocation, not a substitute for a radiosity model of the actual frames and electronics. The radiator drawing explicitly shows four electronics footprints totalling 0.30 m² per 2 × 2 m tile. Against 8 m² of two-face surface, those footprints alone consume 3.75% of gross emission if their own radiation is uncredited. Hinges, heat-pipe ribs, brackets, wiring and adjacent structures consume additional view and area.

The 48-case deterministic stress grid combines two solar absorptances, two emissivities, two outgoing-longwave values, two albedos and three emission fractions. Every case also samples the correlated seasonal geometry. No probability distribution is assigned to these choices.

Parameter	Screened values	Status
Radiator solar absorptance	0.15, 0.20	Coating/degradation design cases
Radiator infrared emissivity	0.85, 0.90	Design cases requiring measurement
Earth longwave exitance	237, 300 W/m²	Environmental sensitivity
Lambertian albedo	0.30, 0.40	Environmental sensitivity
Usable two-face emission fraction	0.85, 0.90, 0.95	Geometric/property allocation
Radiator skin ceiling	320 K	Conditional operating target

Radiator allocation	Nominal seasonal capacity at 320 K	Most demanding stress capacity	Consequence at 330 kW
528 m²	333.84 kW	225.58 kW	Nominally marginal; fails stress cases
640 m²	404.66 kW	273.44 kW	Credible lean variant with derating or narrower qualification envelope
800 m²	505.83 kW	341.80 kW	Passes this case set with 3.57% corner reserve

The 800 m² reference needs approximately 318.62 K in the most demanding sampled combination. Its 1.38 K skin-temperature margin is small enough that unresolved gradients matter. Approximately 773 m² would close the same algebraic corner with essentially no extra area reserve. The 800 m² selection rounds this up to 200 repeated tiles. This is aggregate surface capacity, not proof that every localized heat source reaches that surface. It does not cover all possible environments, specular reflections, failures or ageing histories.

The separate PV calculation retains fully exposed emitting faces at emissivity 0.85. Across the stated Earth-environment sensitivity, its minimum available power is 447.22 kW, only 2.73% above the generation target. Rear-face obstruction by real support and service hardware has not been applied to that PV thermal model. It can increase cell temperature and consume this reserve even when the electrical cell-packing factor is unchanged; it is a separate closure item.

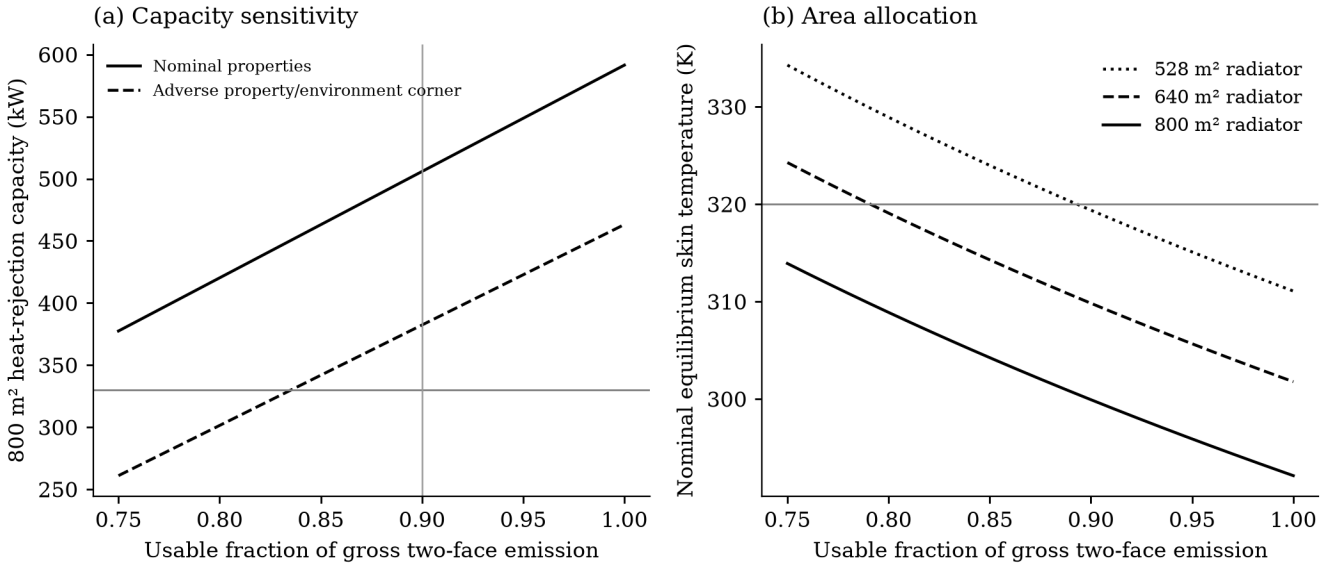


Figure 3. Sensitivity to emitting fraction and optical properties. These curves hold irradiance at 1,317 W/m² to isolate variables. The seasonal values in the table include the separately correlated annual screen and govern the reference decision.

3.3 Why a PV-over-compute sandwich is a competitor, not an automatic winner

An integrated panel saves projected footprint but couples solar waste heat and compute temperature. The PV conversion model includes temperature dependence and explicit delivery factors:

$$\eta_{del}(T_P) = 0.30(0.90)(0.85)(0.98)(0.96)[1 - \gamma(T_P - 298)].$$

The temperature coefficient is varied from 0.0015 to 0.003 per kelvin. These are study assumptions rather than a selected solar-cell specification. At fixed exported bus power, unused absorbed sunlight remains heat. Two modes are included: 414.58 kW recharge export and 300 kW battery-full export. The 330 kW routed design heat remains a conservative thermal sizing load, not an assertion that a 300 kW bus continuously generates 330 kW of physical waste heat.

$$\varepsilon_P \sigma T_P^4 = \alpha_P G - \frac{P_{export}}{A_I} + q''_{E,P} - q''_{PR},$$

$$\varepsilon_R \sigma T_R^4 = \frac{Q_{design}}{A_I} + q''_{E,R} + q''_{PR},$$

$$q''_{PR} = \frac{\sigma(T_P^4 - T_R^4)}{1/\varepsilon_{int,P} + 1/\varepsilon_{int,R} - 1} + g(T_P - T_R).$$

The conductance is an engineered component requirement. It cannot be freely tuned after construction. The sweep uses 0.05–10 W/(m² K) and interior emissivities 0.05, 0.15 and 0.80. Nearly isothermal panels can drive the compute skin too hot. Strongly insulated layers can overheat the PV and overcool the rear. Partial coupling provides a feasible region between them.

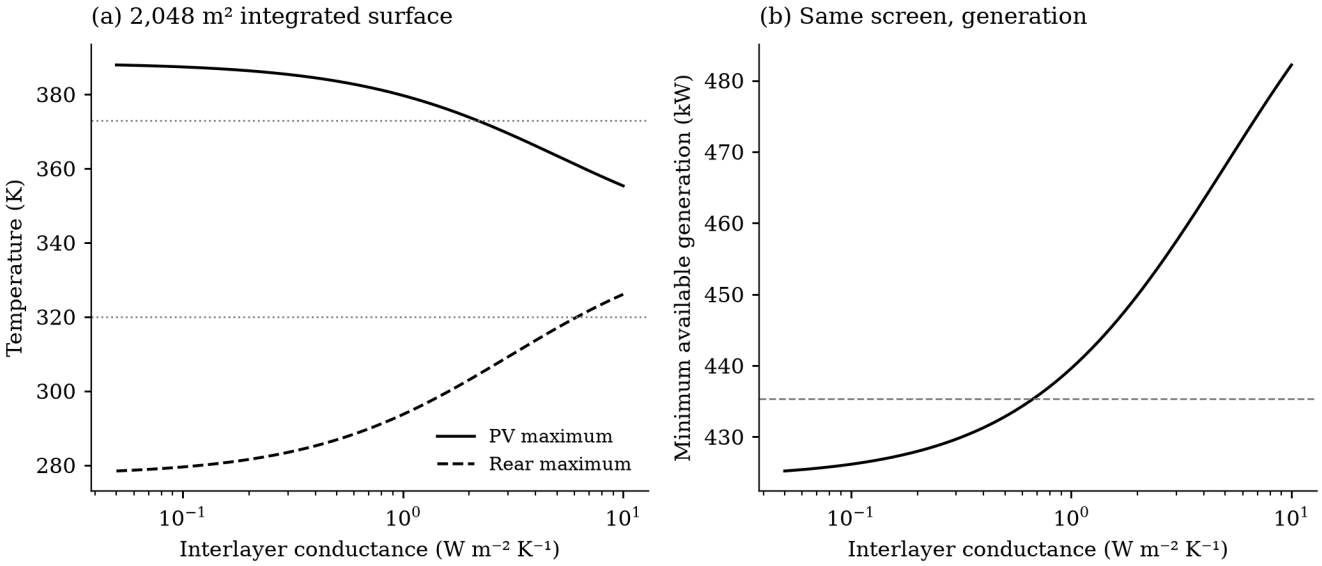


Figure 4. The 2,048 m² integrated candidate under the aphelion screen, including both export modes. Crossing one temperature threshold is insufficient; generation and the other temperature must also satisfy their constraints.

In that aphelion model, 1,760 m² integrated cases fail the combined constraints. A 1,920 m² point barely meets the available-generation target, by approximately 0.017%. A 2,048 m² point with interior emissivity 0.05 and conductance about 3.96 W/(m² K) gives approximately 366.3 K maximum PV temperature and 313.5 K maximum rear temperature. A separate correlated annual screen for a 2,112 m² candidate, using conductance about 5.01 W/(m² K), gives approximately 369.4 K PV and 316.7 K rear maximum, with minimum available generation of 482.7 kW. This remains a conditional illuminated-state result; it has not passed the same full degradation/blockage matrix as the allocated mosaic.

A metre-wide vacuum gap does not by itself thermally isolate broad facing sheets. For two equal coaxial disks,

$$F_{12} = \frac{2 + (d/R)^2 - \sqrt{[2 + (d/R)^2]^2 - 4}}{2}.$$

A 1,700 m² disk has radius 23.26 m. Opposing-face view remains about 0.958 at a 1 m gap and 0.807 at a 5 m gap. Adding neighbouring-tile views is essential: modelling only each tile's opposing counterpart falsely creates a view to cold space. Low-emissivity interior surfaces and deliberately controlled thermal bridges are necessary if broad layers are retained.

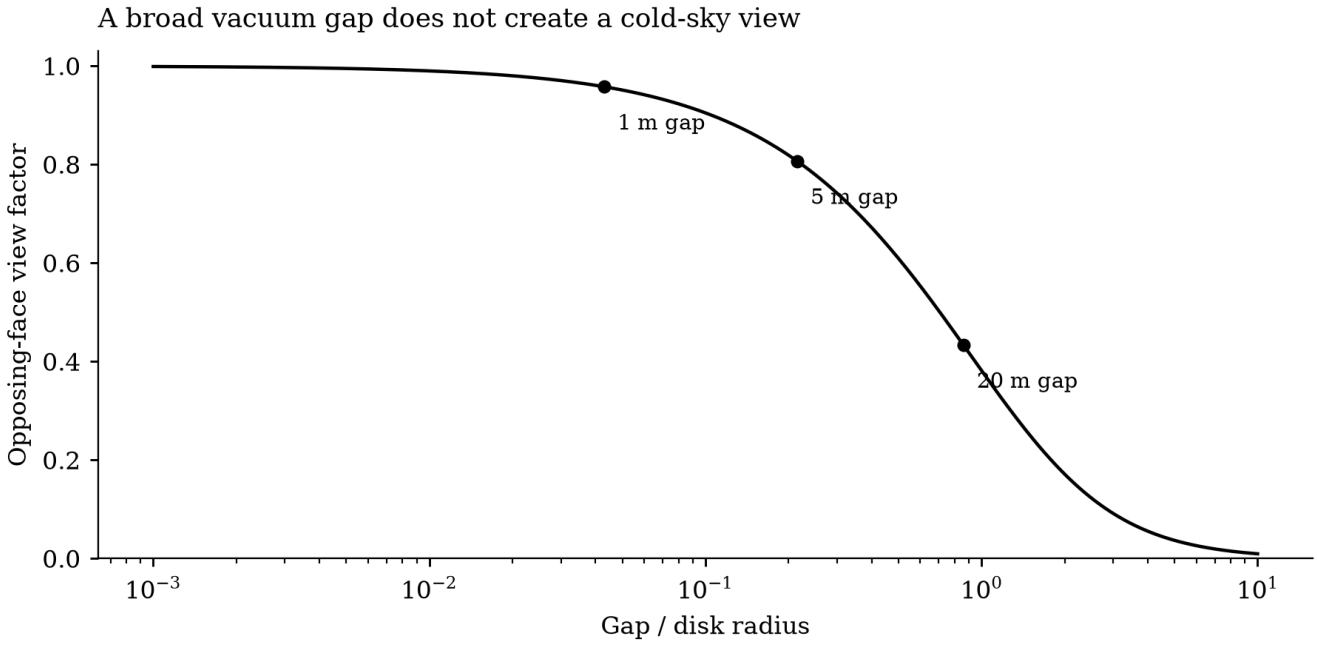


Figure 5. An exact diffuse geometric view factor for equal parallel disks. The labelled gaps use the 1,700 m² reference disk. A gap is not equivalent to a thermal insulator.

3.4 Local thermal paths set compute granularity

The selected radiator carries its compute heat source. No working fluid is intended to cross a deployment hinge. That requires a new compute package: the previous 6.25 kW, 170 mm-high tray cannot silently be retained within a 120 mm stack-pitch allocation. With 200 radiator tiles, a uniform allocation gives 1.25 kW IT and 1.65 kW design heat per tile. Four distributed pickup sites divide this into 312.5 W IT and 412.5 W design heat per site. These are allocations, not claims that a particular accelerator product fits them.

Service-pod losses are localized and cannot automatically be spread over every compute tile. Batteries, communications and protected avionics need an explicit heat budget and either their own counted emitting surface or a qualified local path to designated radiator cells. If those cells come from the 800 m² allocation, the remaining compute-tile loads must be rebalanced; the same area cannot be credited twice. The current drawings define the distributed-compute concept, while this service thermal interface remains open. A bare aggregate heat balance does not close it.

Thin aluminium alone is not a satisfactory metre-scale heat spreader. For a simplified radial conduction path,

$$\Delta T = \frac{Q}{2\pi k t} \ln\left(\frac{R}{r_s}\right).$$

A 312.5 W source spreading from 0.05 to 0.50 m through 2 mm aluminium at 167 W/(m K) develops approximately 343 K in this idealized equation. The complete heat allocation is higher. Embedded heat pipes, vapour chambers, distributed sources or a local pumped loop are therefore functional requirements, not optional refinements. The local temperature chain must be measured:

$$T_j = T_{skin} + Q_{site} R_{package} + \Delta T_{contact} + \Delta T_{spreader}.$$

Cold survival is also unresolved. The illuminated nominal equilibrium of the oversized radiator is around 300 K. During eclipse or low-utilization operation, heat rejection can exceed generation and chill local electronics or working fluid. The transient equation must include thermal mass and controllable rejection:

$$C_{th} \frac{dT}{dt} = Q_{IT}(t) + Q_{heater}(t) + Q_{absorbed}(t) - Q_{radiated}(t).$$

Thermal switches, variable-conductance devices, louvers, acceptable cold operation or duty control are candidates. None is yet sized. Heater energy must fit the battery budget; a reserve battery is not unlimited freeze protection.

4. Mass, outline and structural consequences

4.1 Compare equivalent surface systems

The controlled reference surface coefficients are 2.45 kg/m² for a complete PV surface allocation and 6.0 kg/m² for a radiator surface allocation. They originate in the previous study and are not supplier quotes for the new tiles. Local deployment structure is already represented in these coefficients; adding a complete independent hinge budget on top without first decomposing them would double-count mass.

For integrated and mosaic alternatives,

$$M_I = A_I \mu_I + H_I + M_{common},$$

$$M_M = A_P \mu_P + A_R \mu_R + H_M + M_{common}.$$

The integrated architecture is lighter only when

$$\mu_I < \frac{A_P \mu_P + A_R \mu_R + H_M - H_I}{A_I}.$$

Under the common nominal ideal-area calculation, a mosaic requires 1,694.97 m² PV and 446.31 m² radiator at 320 K. Its surface allocation is 6,830.53 kg. A 2,048 m² integrated plane must therefore achieve less than 3.335 kg/m² at equal outside-surface hardware mass. A 500 kg difference in backbone, shielding, local thermal hardware or electrical interconnect shifts this threshold by 0.244 kg/m². Two aluminium skins each 1 mm thick already weigh 5.4 kg/m² before cells, core, heat transport or joints. Meeting a low integrated mass target requires a materially different construction, not just a favourable spreadsheet entry.

The selected mosaic allocation, with 1,760 m² PV and 800 m² radiator, has a surface term of 9,112 kg. This is not its dry or launch mass. Its algebraic integrated break-even is 4.449 kg/m² at 2,048 m², but that is not an equal-feasibility comparison: the 800 m² radiator passes the stated stress grid whereas that integrated candidate has not yet been screened against the same property, blockage and cold-transient cases. The nominal 3.335 threshold and the allocated 4.449 threshold must not be interchanged to manufacture a winner.

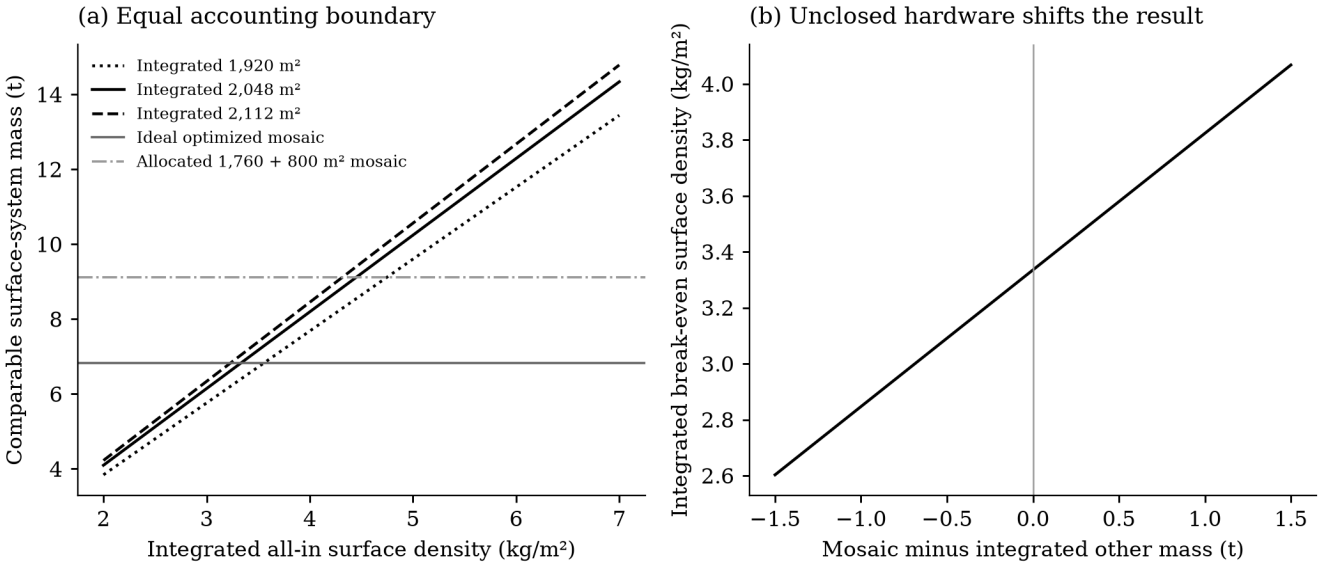


Figure 6. Comparable surface mass, excluding common compute and unclosed differences. The ideal and allocated mosaic lines have different reserve assumptions. The graph exposes a measurement target, not a full-spacecraft mass ranking.

4.2 Compact outlines shorten central services

For a rectangle of area A and aspect ratio chi, uniform areal loading gives

$$\langle r^2 \rangle_{rectangle} = \frac{A}{12}\left(\chi + \frac{1}{\chi}\right), \qquad \langle r^2 \rangle_{disk} = \frac{A}{2\pi}.$$

The rectangle's perimeter relative to an equal-area square is

$$\frac{p(\chi)}{p(1)} = \frac{\sqrt{\chi} + 1/\sqrt{\chi}}{2}.$$

At aspect ratio 10, the mean squared central route is 5.05 times the square's value. This is relevant to service distributions whose mass or loss grows with route length. It is not a proof that all decentralized ribbon networks have that penalty. Local regulation and local heat rejection change the routing problem, but cannot eliminate long-span deployment dynamics.

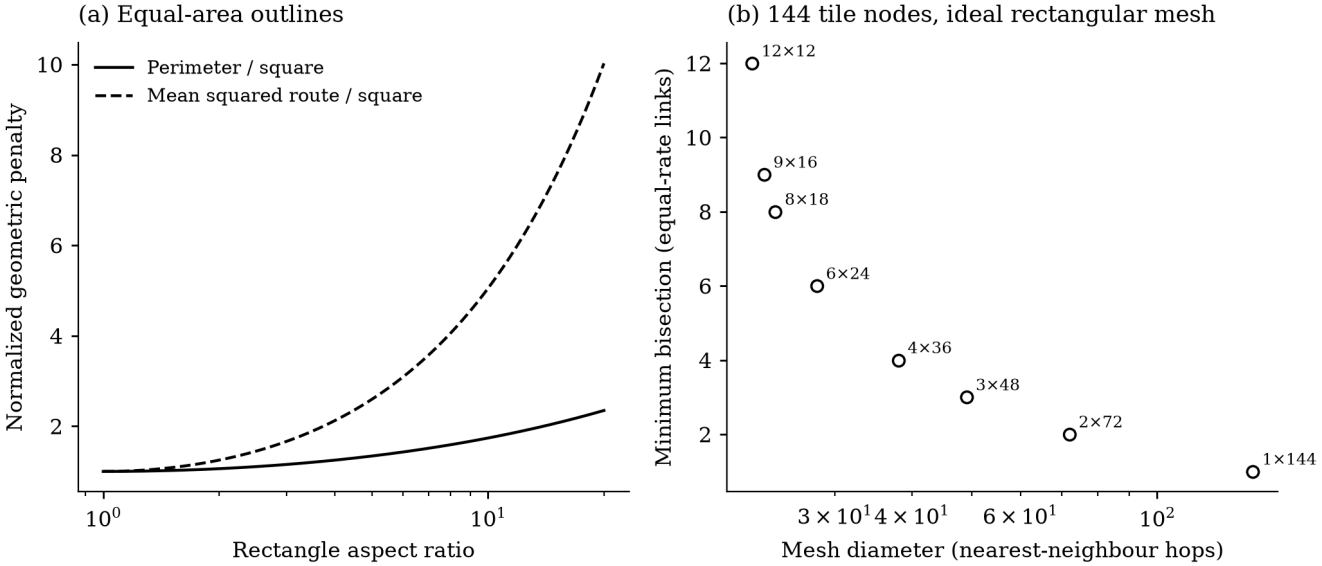


Figure 7. Left: exact outline metrics at equal area. Right: graph diameter and minimum cut for 144 equally connected tile nodes. These graph quantities illustrate communication topology; they are not achieved network throughput.

4.3 Stiffness cannot be purchased with area alone

A beam-like deployable member has a first bending frequency that scales approximately as

$$f_1 \simeq \frac{\beta_1^2}{2\pi L^2}\sqrt{\frac{EI}{m'}}, \qquad \beta_1 \simeq 1.875.$$

At fixed section and distributed mass, doubling span reduces this frequency by roughly four. Real tiled plates contain joint compliance, preload changes, thermal distortion and coupled torsion. Their modes cannot be validated by this beam surrogate. A membrane under tension, a framed mosaic, a deployable truss and a spinning tether have distinct load paths; they need distinct models. The attractive low mass of a solar sail is not evidence that the same structure can carry batteries and computing modules with precise optical pointing.

The selected concept therefore uses a repeated panel framework with distributed thermal sources and a separate service-pod allocation. Whether its global stiffness comes from a deployable perimeter truss, tensioned diagonals, a backing lattice or locked panel edges remains an open trade. The drawings do not invent bolt sizes or joint preload values before launch loads, material allowables and mode requirements exist. Structural-load analysis and qualification must use the actual coupled vehicle environment; the current NASA structural-load standard is NASA-STD-5002B, dated March 2026. NASA-STD-5002.

5. One large plane or several independent planes?

The geometric model enumerates nine platform counts, three tile pitches and five requested aspect ratios, producing 135 deterministic scenarios. The plotted frontier uses only maximum module dimension and an assumed duplicated service-mass term. It must not be interpreted as an LCOC frontier or as a flight-feasible set. It intentionally reveals the geometric incentive to fragment alongside the opposing service-duplication penalty.

$$M_{service,\,extra} = (N_s - 1)m_{fixed}.$$

Fleet decomposition	IT per spacecraft	Equal productive PV / radiator area per spacecraft	Compact 2 m grid example
One spacecraft	250 kW	1,760 / 800 m²	26 × 26 cells, 52 × 52 m
Four spacecraft	62.5 kW	440 / 200 m²	12 × 14 cells, 24 × 28 m
Ten spacecraft	25 kW	176 / 80 m²	7 × 10 cells, 14 × 20 m

The examples preserve total productive area and include one service bay per spacecraft; remaining grid cells are empty reserved apertures. With a hypothetical fixed service mass of 250 kg per vehicle, four spacecraft add 750 kg and ten add 2,250 kg relative to one. The model varies this coefficient from 100 to 500 kg. These values are sensitivity assumptions, not a mass estimate for a completed bus. Large optical terminals or propulsion hardware can change the result materially.

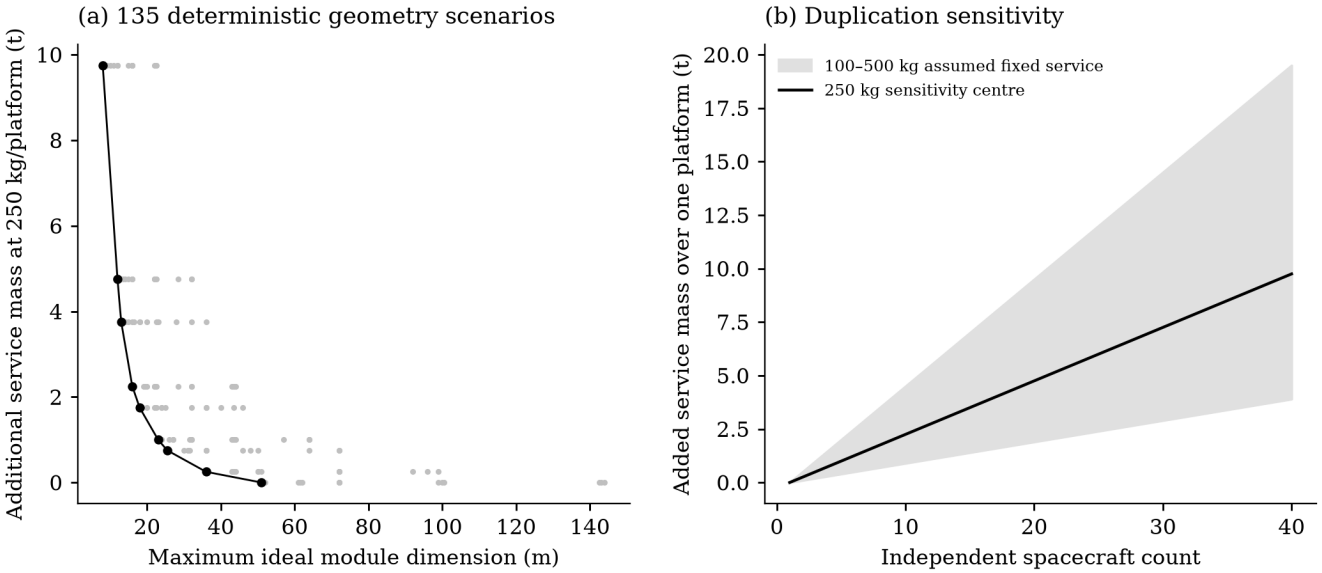


Figure 8. A Pareto set in two stated proxies only. Thermal placement, connected folding, actual fixed equipment, service power and dynamics are not feasibility-tested here. The four-spacecraft reference is a development choice balancing span and duplication, not the mathematical optimum of a complete mass model.

All four vehicles must share the existing 50 kW non-IT operating allocation, including conversion, avionics, communication and thermal-control loads. That is 12.5 kW per spacecraft on average. This budget must be itemized and verified. Duplicated service equipment cannot be added to the mass model while its power is silently left out. If the allocation fails, PV, storage, heat rejection and geometry must be resized together.

Independent jobs can be scheduled to separate vehicles with no assumption of tight formation flight or shared memory. Synchronizing large model training across them is different. A rectangular nearest-neighbour mesh has graph diameter and bisection proportional to

$$D_{mesh} = n_x + n_y - 2, \qquad B_{cut} = \min(n_x, n_y) B_{link}.$$

For a ring all-reduce of message size S over p ranks, a simplified transfer term is

$$t_{ring} \simeq 2(p-1)\alpha + \frac{2(p-1)}{p}\frac{S}{B_{eff}}.$$

Propagation over 28 m is negligible compared with many protocol and collective-operation effects. Switch count, bisection, actual delivered bandwidth, routing after failure and tail latency matter more than quoting the speed of light. Fiber and electrical data links across mechanical seams still need bend-radius, connector-cycle and radiation qualification. One thermal tile is not necessarily one network rank.

6. The reference geometry and what the plans actually define

Each reference spacecraft is allocated a 24 × 28 m plane with 168 cells at 2 m pitch. There are 110 PV cells, 50 compute/radiator cells, one service aperture and seven reserved empty apertures. The fleet has 440 PV tiles and 200 radiator tiles. Productive area is 2,560 m²; total grid footprint is 2,688 m². Empty reserve cells do not magically remove the need for bracing or boundary load paths. The footprint excludes protruding antenna fields of view, thruster plume zones and any later external truss.

Feature	One spacecraft	Four-spacecraft fleet
Nominal grid envelope	24 × 28 m	Four separate envelopes
PV tiles / planform	110 / 440 m²	440 / 1,760 m²
Radiator tiles / planform	50 / 200 m²	200 / 800 m²
Gross two-face radiator surface	400 m²	1,600 m²
IT / operating bus / thermal design	62.5 / 75 / 82.5 kW	250 / 300 / 330 kW
IT / design heat per radiator tile	1.25 / 1.65 kW	Same repeated tile
Proposed stack pitch, PV / radiator	25 / 120 mm	Packaging assumptions
Sum of active-panel stack thickness	8.75 m	Four parallel stacks total
Service and restraint axial allowance	1.60 m	Included once per stack
Free-stack total height	10.35 m	Before connected-fold verification

The tile map in OC250-101 distributes radiators through the plane as an initial placement. It is not a solved topology optimization of wiring, mass centroid or thermal distortion. Empty cells, local service access and the central service bay are explicit. The 2,000 mm pitch is the allocated cell dimension. Actual panel edges must leave finite joint gaps inside it, and the resulting cell/thermal active-area reductions must fit the generation and emission allocations. Four drawings of identical modules do not establish collision-safe orbital separation; the operational fleet does not require a rigid shared array.

6.1 Launch packing: a bound, followed by a mechanism problem

For one module stacked without a connected hinge constraint,

$$H_{stack} = 110(0.025) + 50(0.120) + 1.600 = 10.350 \text{ m}.$$

Four 2 × 2 m stacks on centres spaced by 2.10 m fit the cross-section of the assumed 6.5 m cylinder. The farthest corner radius is approximately 2.899 m, leaving approximately 0.351 m to the nominal cylinder and 0.201 m beyond an assumed 0.150 m radial reserve. This arithmetic includes four stacks total, one per spacecraft. It does not claim four stacks per spacecraft. The service cell is an aperture, with its packed equipment included in the 1.6 m allowance. The seven reserve cells are empty, not hidden full panels.

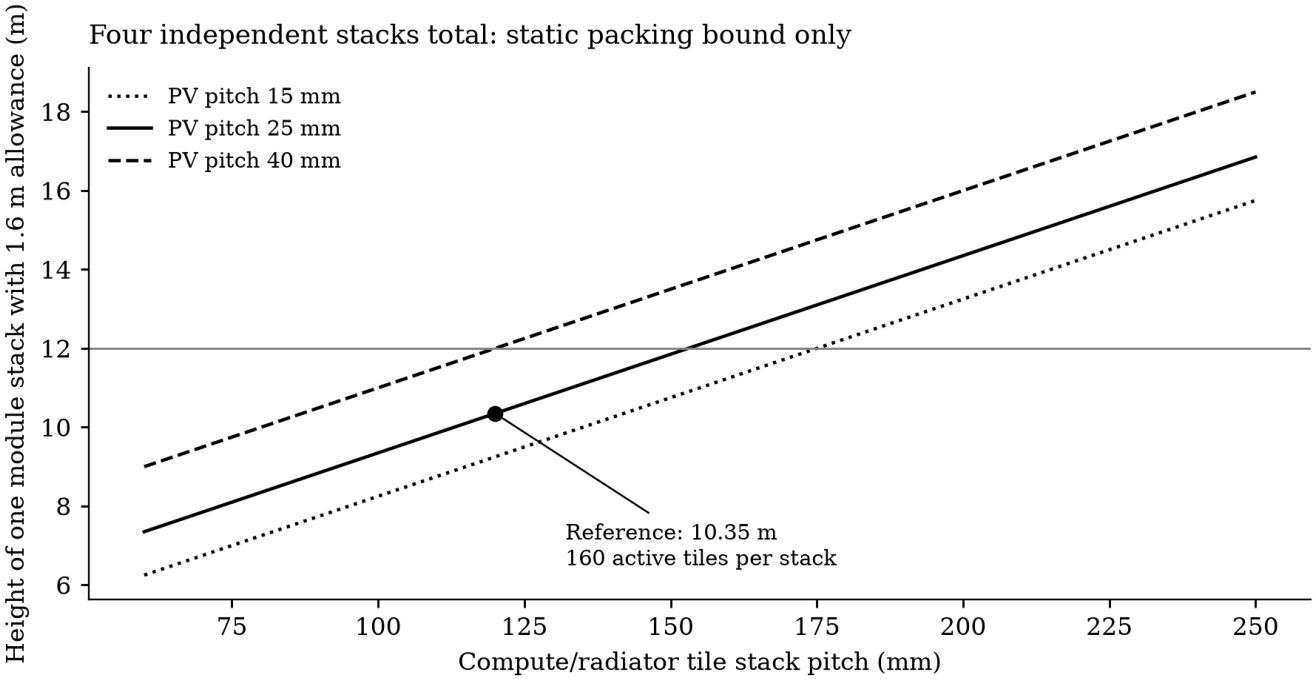


Figure 9. Static unconnected packing. Mixed PV and radiator thicknesses must eventually be accommodated by offset hinges or another deployment mechanism. A geometrically fitting pile of panels is not a flight deployer.

The available 1.65 m axial remainder cannot be credited twice to hinge offsets, launch restraints, separation hardware and margin. Fairing taper, dynamic envelope, acoustics, contamination covers, centre of mass and payload-adapter load paths are not modelled by the cylinder. An actual Starship or other launcher interface must be obtained and checked before launch compatibility is claimed.

A spanning tree of 160 active panels requires at least 159 pairwise panel connections if connected through panel joints. Four such modules imply at least 636 connection events. The actual frame may connect through service and reserved bays, or add edge latches and braces. A full 12 × 14 grid has 310 nearest-neighbour edges per spacecraft. These counts are topological bounds, not a released hinge schedule.

Finite-thickness folding is especially important because the proposed 25 mm and 120 mm stack pitches differ substantially. A simple zero-thickness origami crease pattern cannot close the mechanical design. Deployment must explicitly sequence release, separation, unfolding, restraint, damping, power enable and optical acquisition. The mechanism requirements should be developed against the actual mission and verified through analysis and test. ECSS mechanisms standard, NASA/JPL thick-fold engineering brief.

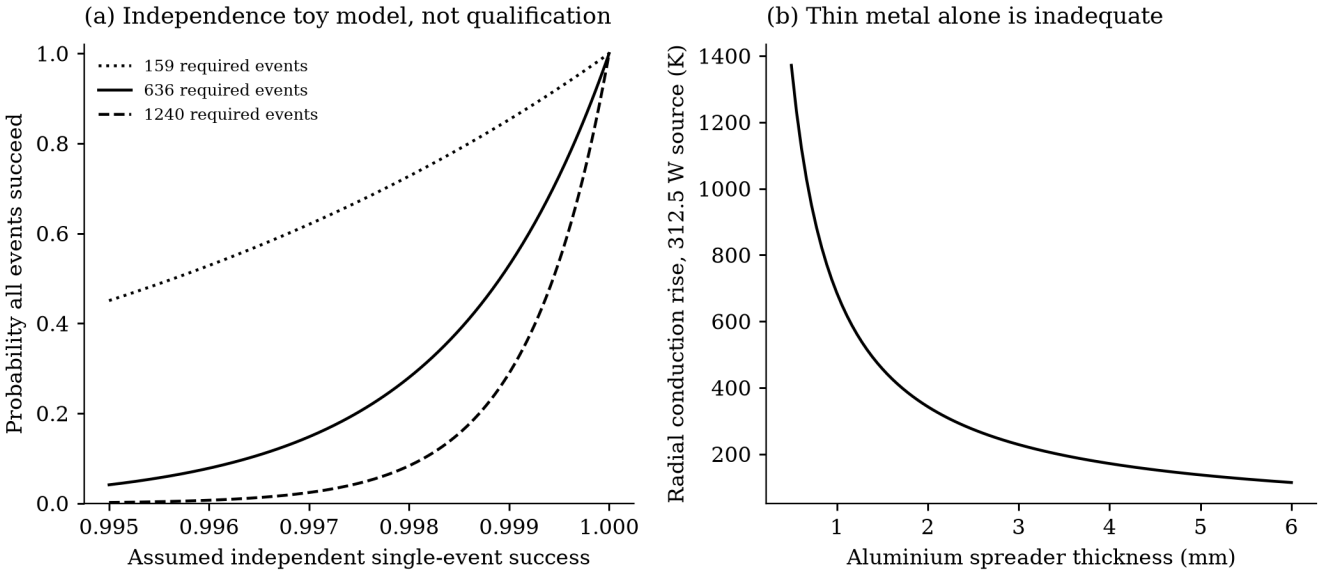


Figure 10. Left: an independence illustration showing why many required events deserve fault-tolerant deployment design; no reliability is assigned to the proposed mechanism. Right: an idealized IT-only conduction rise; the thermal design allocation per site is larger and requires a real heat-transport device.

6.2 Construction hypotheses and material choices

The following construction choices are candidates for coupon and subassembly testing. They are not flight-qualified specifications. Material selection is driven by the functions it must close, rather than an arbitrary collection of aerospace materials.

Element	Candidate construction	Required evidence before design release
PV tile	Space-compatible cell/coverglass/interconnect assembly on lightweight sandwich or tensioned support	Delivered W/m² versus temperature and degradation; complete mass including joints
Radiator tile	Metallic or composite sandwich carrying bonded or embedded heat-transport paths and low-absorptance coating	Two-face temperature map, coating optical properties, bond durability and view loss
Local heat pickup	Package cold plate plus vapour chamber or heat pipes; local pumped alternative retained	412.5 W design heat per site, startup, capillary/pressure limits and cold survival
Frame	CFRP truss or sandwich elements with metallic interface inserts	Joint stiffness, thermal distortion, outgassing, atomic-oxygen protection and galvanic isolation
Joints	Offset hinges and positive deployed restraints, materials selected after load analysis	Release torque, friction, wear, preload, thermal cycles and partial-deployment recovery
Electronics housing	Local conductive enclosure with dose-based shielding architecture	Total dose, single-event effects, charging, EMI and actual shielding mass
Harness	Segregated high-voltage power and optical/data routes with compliant service loops	Insulation, bend radius, arc/plasma environment, fatigue and fault isolation
Service pod	Batteries, protected avionics, communications and propulsion on a separate load path	Complete mass/power/volume budget, plume and sensor exclusions, disposal capability

Aluminium can provide useful contact conductance and manufacturable thermal passages; its coefficient of thermal expansion and mass must be reconciled with composite supports. CFRP can improve stiffness per mass but is anisotropic, may complicate electrical grounding and demands compatible joints. Titanium may be appropriate at particular loaded or thermally isolating interfaces; blanket use is not justified. Radiator coatings require both initial and aged spectral properties. Multilayer insulation belongs on protected non-radiating equipment surfaces, not over a surface credited with emitting heat.

Radiation and debris do not disappear because electronics are distributed. Local shielding mass follows areal density times the protected surface, and semiconductor reliability needs component and system evidence. Replication can tolerate isolated compute-node failures, but a common power arc or a deployment failure can defeat that redundancy. A radiation-aware scheduler, checkpointing and error correction improve resilience only after measured failure modes and useful-work throughput are known.

6.3 Electrical distribution remains a first-order design problem

The 110 PV tiles per spacecraft feed 50 compute/radiator tiles and service loads. Localizing heat does not eliminate energy routing. For a two-conductor path of length L to a load P, resistivity rho and allowed fractional resistive loss delta, the ideal required conductor cross-section is

$$A_c = \frac{2\rho L P}{\delta V^2}, \qquad M_c = 2 L A_c \rho_m .$$

At fixed load and fractional loss, ideal conductor mass scales as inverse voltage squared and length squared. The advantage of high voltage must be balanced against insulation, switching, surface charging, arcing, connectors and fault containment. The illustrative 75 kW, 13 m feeder calculation is not a recommended voltage selection. A branched distributed collector has a different exact optimum.

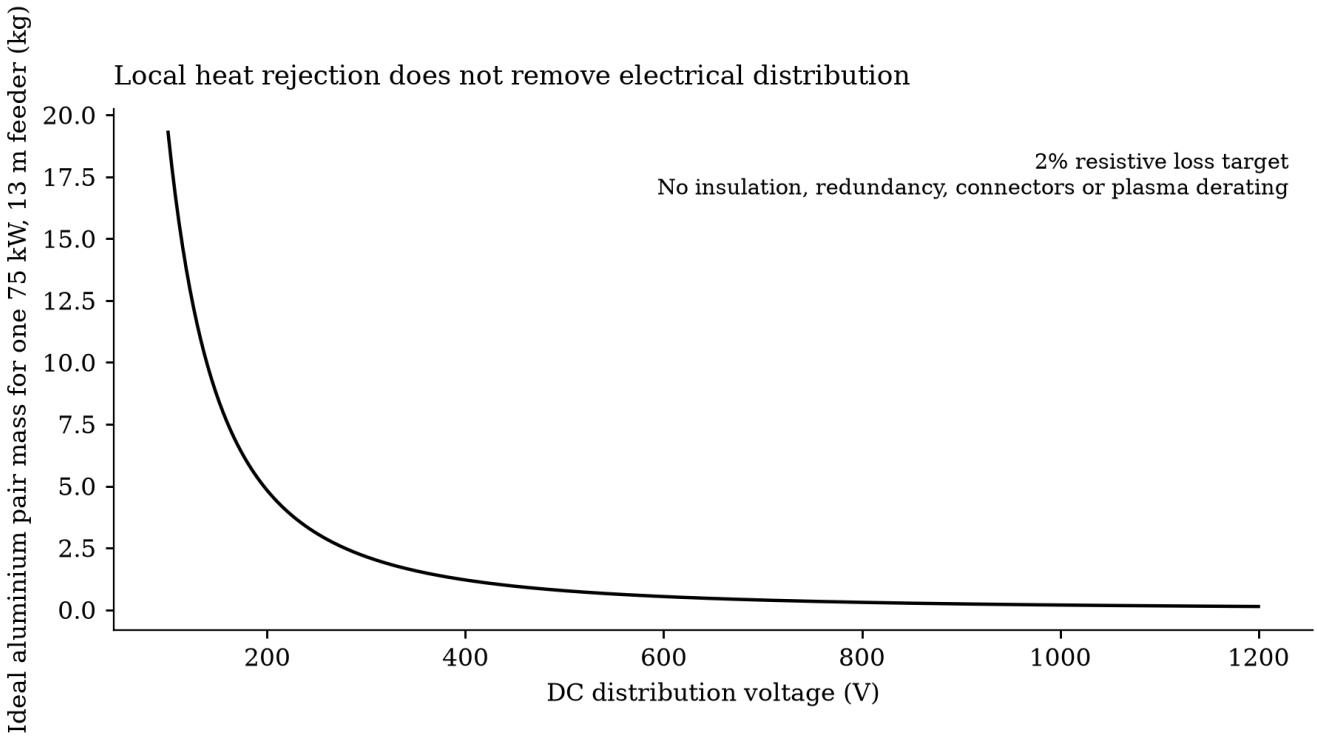


Figure 11. Ideal aluminium conductor-pair mass only, with a 2% resistive-loss allocation. Flight harness mass includes substantial additional terms. No voltage on this curve is selected or qualified by the graph.

The power architecture should isolate each PV segment, prevent a tile fault from collapsing the common bus, and provide independently switchable compute islands. Power and data connectors crossing a hinge must have a defined cycle and strain-relief design. Optical and RF terminals require operational fields of view that the flat plane and service structure do not block. No ground-to-orbit availability claim is made without a ground-station and cloud-coverage model.

7. Scale to 1 MW and 10 MW without hiding a different problem

If irradiance, cell performance, cooling temperature and useful load fraction are unchanged, both collector and radiator area scale approximately linearly with IT power. A geometrically similar single plane therefore grows in span with the square root of power. Flexible-mode penalties then grow faster than the visual change suggests.

$$A \propto P_{IT}, \qquad L \propto P_{IT}^{1/2}, \qquad f_1 \propto L^{-2} \quad \text{for fixed section.}$$

Repeating the 62.5 kW reference would require 16 vehicles for 1 MW and 160 for 10 MW. This keeps individual span bounded but moves cost into launch handling, bus duplication, autonomous operations, communication coverage and fleet replacement. It does not demonstrate a commercially viable 10 MW system. Combining the modules into a rigid giant plane or tightly synchronized optical cluster would be another architecture and must be re-evaluated.

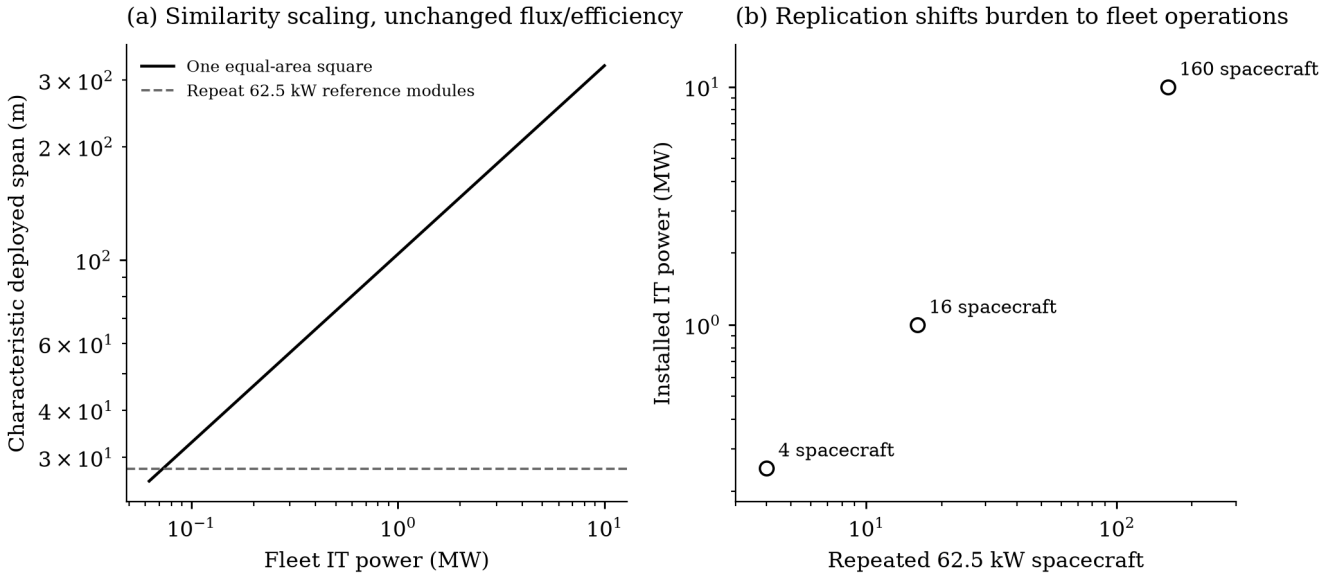


Figure 12. Similarity and replication bounds. A repeated-module programme avoids increasing each vehicle's span but does not avoid proportional productive surface area or fleet overhead.

The business metric must count completed useful jobs, not installed peak floating-point operations. Shape influences cost through launch mass, deployment yield, operating power, failure domains and service throughput:

$$LCOC = \frac{\sum_{t=0}^{N} (C_{cap,t} + C_{op,t})(1+r)^{-t}}{\sum_{t=1}^{N} F_{useful,t}(1+r)^{-t}},$$

$$F_{useful} = F_{peak}\, u\, \eta_{workload}\, a_{service}\, t.$$

Here utilization, workload efficiency and service availability are separate terms. Adding more satellite modules may lower structural risk while worsening service overhead or network efficiency. No launch-price promise, future accelerator performance or uncontracted customer is used to choose the form. A 2029–2030 programme needs measured prototype performance and actual procurement interfaces before these expressions become a bankable cost model.

8. Verification gates and the experiment that can change the decision

The drawings define allocated geometry, sections, tile types and topology. They deliberately do not claim an as-built assembly, released fastener pattern, welded pressure circuit or verified folding trajectory. The next engineering release must close the following requirements with evidence. Each failure can change the form rather than merely adding a note to a drawing.

ID	Requirement / acceptance basis	Verification and decision consequence
PWR-01	Available EOL bus generation at least 435.311 kW for the fleet	Measured PV properties, rear-face blockage and trajectory model; resize if violated
PWR-02	Four complete service systems fit the 50 kW non-IT allocation	Itemized modes and measured conversion losses; recouple all sizing if violated
THM-01	Reject 330 kW with measured optical properties, blocked area and local gradients	Correlated thermal model plus vacuum test; increase area, temperature or reduce IT
THM-02	Local package and memory limits satisfied at 1.65 kW design heat per tile	Instrumented active tile; average skin temperature alone is insufficient
THM-03	Eclipse, low-load and safe-mode cold survival fit energy and material limits	Transient test/model; select variable rejection or compatible local loop
THM-04	Localized service-pod losses have a physical path to counted emitting area	Allocate service heat and designated radiators; rebalance the 800 m^2 without double counting
DEP-01	Connected mixed-thickness assembly stows and deploys without collision	Finite-thickness kinematic model and gravity-offloaded hardware deployment
DEP-02	Failed release or latch does not create an uncontrolled common failure	Fault-injection sequence and recoverability demonstration
STR-01	Launch and deployed modes, stress, stability and optical jitter close	Coupled-load analysis, joint tests and modal correlation
MASS-01	Complete dry and launch mass, including interfaces, shielding and carrier	Bottom-up bill of materials; compare equal boundaries and explicit reserve
DATA-01	Useful workload meets throughput and latency targets after node/link faults	Representative compute fabric and scheduler benchmarks
ENV-01	Radiation, atomic oxygen, charging, debris and thermal cycling are addressed	Mission-specific analysis and component/system qualification
OPS-01	Each independent spacecraft can acquire, control, communicate and dispose	Flight-dynamics and ground-segment design with actual propulsion/terminal interfaces

The highest-value first article is a full-size 2 × 2 m compute/radiator tile with representative electronics heat injection, its actual coating, heat transport, harness penetrations and mounting frame. Measure temperatures and rejected heat in vacuum across the intended hot and cold cases. In parallel, test a PV tile and a mechanically representative mixed PV/radiator fold segment. Record complete assembly mass, not just face-sheet mass. A six-to-nine-cell deployment segment with heterogeneous thicknesses can test the mechanism assumptions before committing to 160 active panels per vehicle.

An integrated sandwich tile should be built as a competing experiment with the same accounting boundary. It wins only if its measured mass advantage survives actual power/temperature coupling, package height, cold transients and deployment. The nominal 3.335 kg/m² break-even provides a clear target for a 2,048 m² integrated concept at equal other mass. The fair decision is to eliminate architectures using these measurements, not to lock a conventional satellite shape and then make the drawings more detailed.

9. Selection statement

The best-supported development form under the current mission assumptions is a compact, modular planar mosaic with distributed computing on dedicated two-face radiators. The reference uses four independent 24 × 28 m spacecraft, with 1,760 m² PV and 800 m² radiator planform across the fleet. Its advantage is the separation of PV and compute thermal requirements, short local heat paths and repeated testable modules. A simple fully curved PV body is penalized by collection geometry; a very long ribbon is penalized by span and communication topology; a disk provides only modest ideal routing benefits while complicating repeated rectangular packaging.

The evidence selects a family and a concrete candidate for development. It does not prove that four vehicles, this exact tile count or this shape has the lowest lifecycle cost. The integrated plane and cap-and-skirt remain alternatives if their complete mass, thermal and deployment evidence is better. The central-core architecture is no longer the default.

Appendix A. Drawing and data register

The following A3 vector sheets follow this report. They are also supplied individually as editable DXF and SVG. Dimensions inside each named DXF physical-view block are true millimetres; the sheet contains separate projections and is not one assembled three-dimensional solid. The printable paper-space layout preserves the annotated sheet. NTS comparison sketches must not be measured as engineering geometry.

Sheet	Subject	Definition level
OC250-100	Eight architecture families	Functional comparison, NTS
OC250-101	Reference spacecraft plan and tile allocation	Dimensioned grid and numbered PV/RAD/service/reserve cells
OC250-102	Mixed PV/radiator construction section	Allocated layers, thermal separation and mounting constraints
OC250-103	Local compute/radiator tile	Pickup footprints, heat paths and blocked-area accounting
OC250-104	Four-stack launch packing bound	Study cylinder, stack heights and explicit unconnected status
OC250-105	Deployment graph and sequence	Topology and unresolved mechanism interfaces

The source package contains the model inputs, numerical results, all figure scripts, thermal quadrature/coupling model, structural screening model, compute-topology bounds, drawing generator and an independent design review. The geometric scenario table has 135 rows. The thermal stress combinations are deterministic cases, not Monte Carlo confidence intervals. No FEA, flight component qualification or launcher approval is implied by reproducibility of these calculations.

Appendix B. Reproduction and interpretation

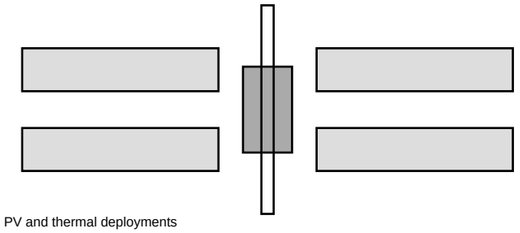
Use a recent Python 3 environment with numpy, matplotlib, reportlab, ezdxf, pypdf and markdown-it-py. The delivered source and data freeze the exact assumptions used for the figures. Run the shape model to regenerate the plots and configuration, then the drawing generator. The standalone thermal model reproduces its quadrature and sensitivity data; its detailed loop is substantially larger than the outline model. The included README provides the execution order and distinguishes generated data from explanatory review notes.

The drawing dimensions are design allocations. Before fabrication, replace them with a configuration-controlled three-dimensional assembly, engineering drawings with justified tolerances, complete interface-control documents, traceable material/process specifications, signed analyses and a qualification plan. Interface management must track geometry, electrical power, thermal transport, data and operations together. ECSS interface-management standard.

ARCHITECTURE-FIRST SHAPE COMPARISON

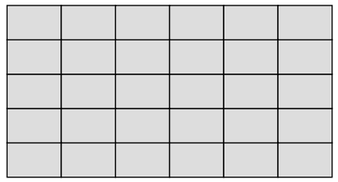
Eight topology families; NTS silhouettes expose functional geometry, not validated products

A PREVIOUS CENTRAL CORE



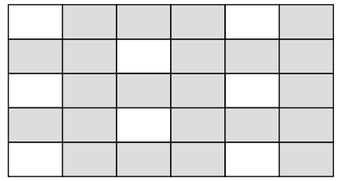
Separate large PV and thermal deployments

B INTEGRATED DUAL-FACE PLATE



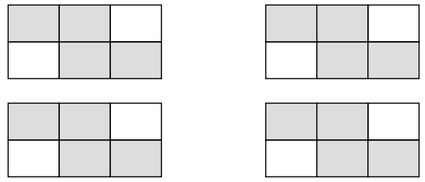
PV and radiator share every tile footprint

C MIXED PV / RADIATOR GRID



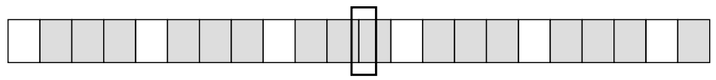
PV-only and two-face radiator-only tiles

D INDEPENDENT MODULAR FLEET



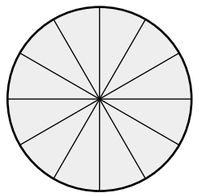
Repeated spacecraft; independent jobs / control

E LONG FOLDING RIBBON



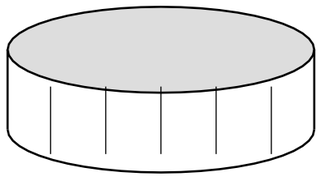
Simple chain; severe long-span stiffness penalty

F INTEGRATED FLAT DISC



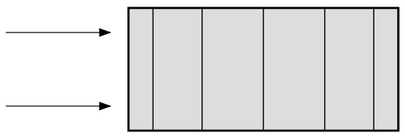
Compact radius; wedge / thick-fold problem

G PV DISC + RADIATOR SKIRT



PV cap; local compute on peripheral radiator

H CURVED PV CYLINDER



PV area / projected area = pi for side cylinder

DESIGN / INTERFACE NOTES

01. Leading candidate is C implemented as D: separate PV-only and radiator-only tiles on a few independent flat spacecraft. This selection is conditional on the thermal and mass study, not a claim of global optimality.

02. B and C are physically different. In B, an illuminated PV layer sits in front of every rear radiator. In C, no PV tile overlays a radiator; each radiator can expose both faces, subject to Sun/Earth loading and obstructions.

03. Reference mission: 250 kW IT for Earth-serving cloud jobs. Independent modules do not require precision formation flying when jobs and model state remain local.

04. Current configuration: 4 spacecraft, each 12 x 14 bays at 2 m pitch. Each uses 110 PV bays, 50 radiator bays and 1 service bay.

05. The previous shape was driven by functional separation and a central rack. Here projected solar area, usable radiation area, thermal paths, launch packing and flexible structure determine the shape.

06. Disc/skirt G remains a real alternative. It must not be confused with H, where PV coating wraps around a surface with poor average solar incidence.

07. Ribbon E may simplify one-dimensional deployment but pays a strong span/stiffness penalty. Thin membrane solar-sail heritage does not qualify a compute-loaded ribbon.

08. No silhouette proves reliability, launcher fit, useful FLOP per joule or a construction schedule. Retain the competing families in the documented trade space.

OC-250 / ORBITAL COMPUTE PLATFORM

ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE

All dimensions nominal unless stated. No approval or qualification is implied.

DRAWING	SCALE / FORMAT	REVISION / UNITS	SHEET
OC250-100	**NTS**	**C / mm**	**01 / 06**
ISSUED 2026-09-25	A3 420 x 297	3RD ANGLE *	* axes on views

MIXED-TILE SPACECRAFT PLAN AND CELL IDENTIFICATION

Representative module M01; M02 onward repeat the allocation independently

MODULE PLAN / 24 x 28 m

SCALE 1:150 | DIMENSIONS mm

RESERVE	RESERVE	R048	P106	P107	R049	P108	P109	R050	P110	RESERVE	RESERVE
P097	P098	R045	P099	P100	R046	P101	P102	R047	P103	P104	P105
P089	R041	P090	P091	R042	P092	P093	P094	R043	P095	P096	R044
R037	P081	P082	P083	R038	P084	P085	R039	P086	P087	R040	P088
R033	P073	P074	R034	P075	P076	R035	P077	P078	R036	P079	P080
P064	P065	R030	P066	P067	R031	P068	P069	R032	P070	P071	P072
P056	R026	P057	P058	R027	P059	P060	P061	R028	P062	P063	R029
P048	P049	P050	R023	P051	**S001**	P052	R024	P053	P054	R025	P055
P040	P041	R019	P042	P043	R020	P044	P045	R021	P046	P047	R022
P032	R015	P033	P034	R016	P035	P036	R017	P037	P038	P039	R018
R011	P024	P025	R012	P026	P027	P028	R013	P029	P030	R014	P031
P015	P016	P017	R008	P018	P019	R009	P020	P021	R010	P022	P023
P007	P008	R004	P009	P010	R005	P011	P012	R006	P013	P014	R007
RESERVE	P001	R001	P002	P003	R002	P004	P005	P006	R003	RESERVE	RESERVE

28,000 / 14 BAYS

24,000 / 12 BAYS

P — PV ONLY

R — RAD ONLY

S — SERVICE

O — RESERVE

DESIGN / INTERFACE NOTES

01. This is a proposed leading shape, conditional on the quantified thermal, mass, packing and reliability trade. Each spacecraft carries its own flight safety, propulsion and ground communication services.

02. Installed fleet geometry: 1760 m2 gross PV footprint and 800 m2 radiator footprint. The latter has 1600 m2 gross two-face area before obstructions and view-factor corrections.

03. PV and radiator cells occupy different planar bays. There are no PV cells over the radiator tiles. The service bay carries spacecraft functions, not an assumed centrally concentrated compute rack.

04. Radiator positions are spread through the grid to shorten local power/data runs and distribute compute mass. This deterministic pattern is a routing proposal, not an optimized thermal distortion solution.

05. The 2,000 mm nominal bay is a structural allocation. Cell packing, edge frames, hinges and harnesses reduce effective collecting/emitting area and must appear in the separate fill-factor model.

06. P/R/S numbers are unique within M01. Prefix equipment labels with M02, M03 etc. for the remaining spacecraft. Reserve bays are empty apertures, with no hidden panels or credited optical/radiator area; perimeter bracing remains to be sized.

07. Independent radiator-tile compute zones avoid a large moving spacecraft-wide liquid loop. Local transport is unqualified. Service-pod losses need their own counted emitting area or qualified local path; the uniform tile budget does not close that interface.

08. Support section sizes, truss depth, hinge axes, latches, thruster keep-outs and optical-terminal viewing cones require a three-dimensional accommodation and dynamics model.

COUNTS / MODULE

P = 110

R = 50

S = 1

RESERVE = 7

IT 62.50 kW

PV 440 m2

RAD 200 m2

OC-250 / ORBITAL COMPUTE PLATFORM

ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE

All dimensions nominal unless stated. No approval or qualification is implied.

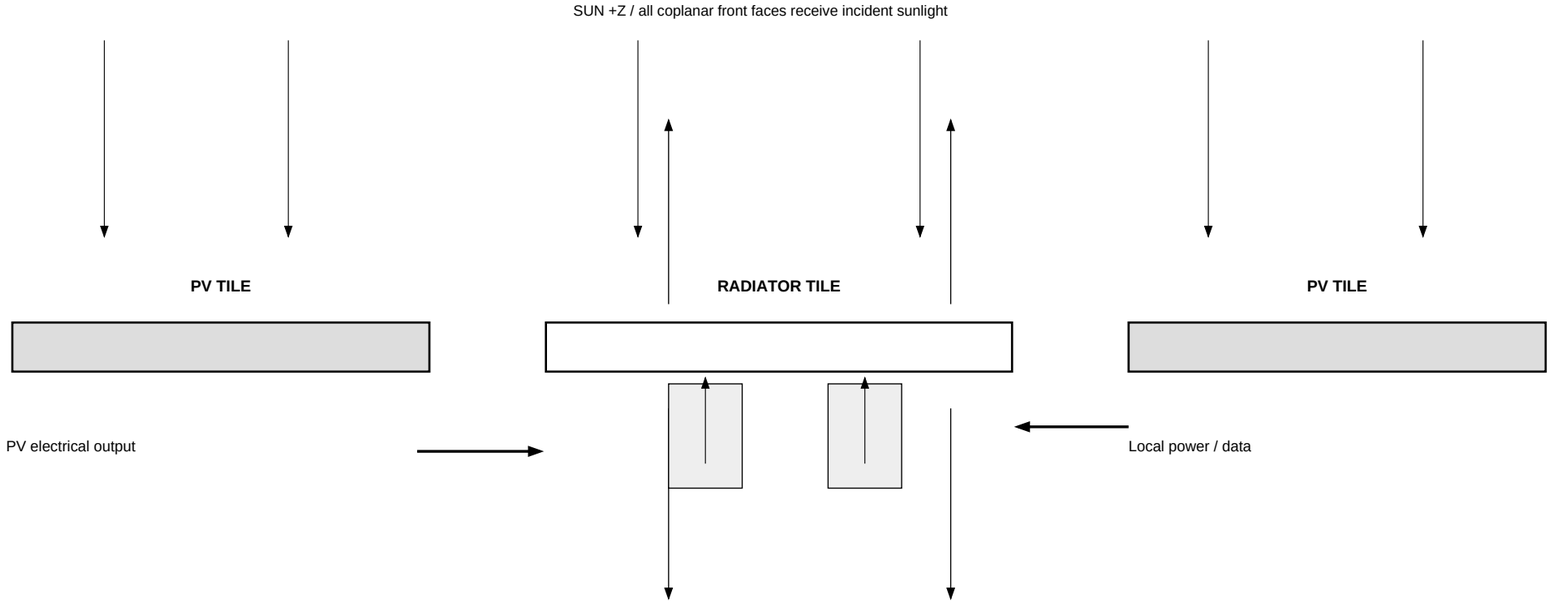
DRAWING	SCALE / FORMAT	REVISION / UNITS	SHEET
OC250-101	**1:150**	**C / mm**	**02 / 06**
ISSUED 2026-09-25	A3 420 x 297	3RD ANGLE *	* axes on views

MIXED-SURFACE CROSS SECTION AND THERMAL INTERFACES

Coplanar PV-only and radiator-only tiles; both radiator faces require environmental accounting

FUNCTIONAL SECTION / PV - RAD - PV

SCALE SCHEMATIC NTS | DIMENSIONS mm



DESIGN / INTERFACE NOTES

01. Separate surfaces are the point of this candidate: a PV bay generates electricity; a radiator bay carries local compute and rejects heat from two exposed faces. They share array support but not the same tile footprint.

02. A coplanar radiator front face is illuminated. Deduct solar absorption alpha_solar x solar irradiance, Earth infrared and albedo as appropriate. Two-face area cannot be multiplied by a deep-space-only heat flux.

03. Rear electronics, frame members, harness and neighbouring service hardware obstruct view. Face temperatures, emissivities and obstructions may differ; integrate the two faces separately.

04. Electrical/data harnesses cross between tile bays. No large shared fluid supply and return is assumed across deployment hinges; each radiator tile requires its own qualified heat-transport solution.

05. The 25 mm PV and 120 mm radiator packed pitches are replaceable accommodation hypotheses. They are not released laminate thicknesses or qualified assembled product dimensions.

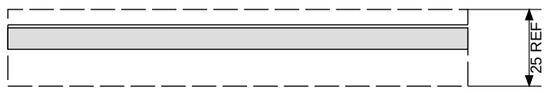
06. The large central drawing is deliberately NTS so the layer functions are legible. The lower true-size boxes show only packed-envelope dimensions; unlabelled internal layer positions are schematic.

07. A warm local compute package requires a low-resistance conduction interface to heat pipes, a vapour chamber or another validated local spreader. The heat-transfer device cannot be replaced by a bare thin plate.

08. When compute is off, tile temperature, working-fluid freeze/restart, battery survival and safe attitude remain design gates. Distributed heat rejection does not remove the cold-survival requirement.

PV PACKED ENVELOPE

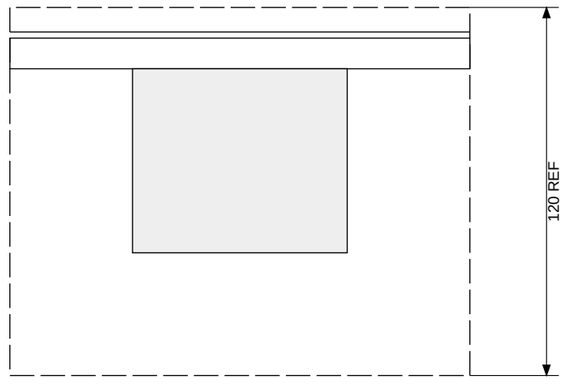
SCALE 1:2 | DIMENSIONS mm



Internal layer positions schematic. Packed pitch includes surface protection and separation.

RAD / COMPUTE ENVELOPE

SCALE 1:2 | DIMENSIONS mm



No PV layer over this radiator

OC-250 / ORBITAL COMPUTE PLATFORM

ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE

All dimensions nominal unless stated. No approval or qualification is implied.

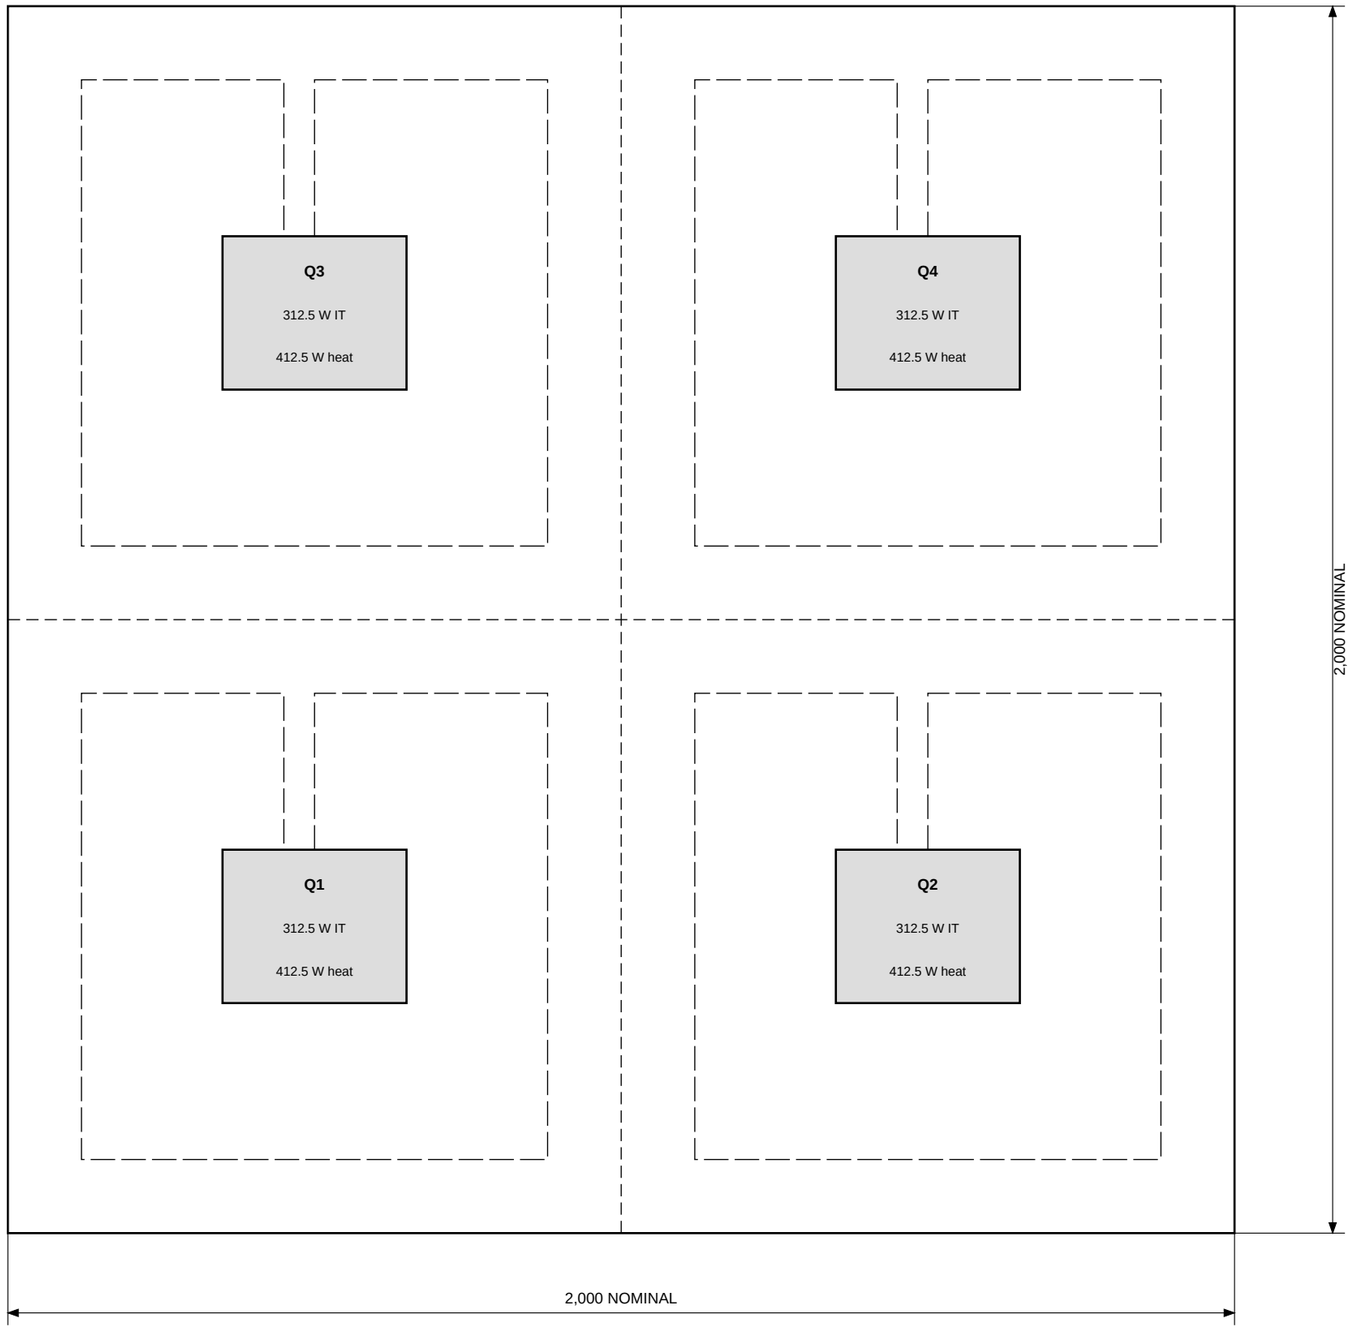
DRAWING	SCALE / FORMAT	REVISION / UNITS	SHEET
OC250-102	**NTS / 1:2**	**C / mm**	**03 / 06**
ISSUED 2026-09-25	A3 420 x 297	3RD ANGLE *	* axes on views

RADIATOR COMPUTE TILE AND LOCAL HEAT DISTRIBUTION

Four heat-pickup zones shown; heat-transport technology and capacity are unqualified

RADIATOR TILE / REAR FACE

SCALE 1:10 | DIMENSIONS mm



PER TILE

IT 1.250 kW

Heat 1.650 kW

4 local heat pickups

PROJECTED AREA

Front 4.00 m2

Rear 4.00 m2

Shown boxes:

4 x300 x250

=0.30 m2

Gross two-face

8.00 m2

After boxes only

<= 7.70 m2

More deductions:

frames / hinges

harness / views

DESIGN / INTERFACE NOTES

01. At 62.50 kW IT and 82.50 kW design heat per spacecraft, 50 radiator tiles each carry 1.250 kW IT /1.650 kW heat. Each of four zones is allocated 312.5 W IT /412.5 W heat.

02. The four 300 x250 mm projected hardware footprints are proposed rear-face exclusions. They remove 0.30 m2 from the gross two-face area before accounting for edge frames, joints and wiring.

03. For the current 2 m tile, geometry alone leaves at most 7.70 m2 after those four footprints. No radiation credit is taken for the electronics boxes without surface-temperature and coating evidence.

04. Q1-Q4 identify heat-pickup zones, not a selected GPU count. Actual processor, memory, converter and shielding packaging may increase the footprint or stack pitch.

05. Dashed routes show possible local heat-transport coverage. They are not manufacturing pipe paths, wick specifications, validated loop lengths or demonstrated transport capacity.

06. Candidate technologies: embedded heat pipes, independent small vapour chambers or a local closed pumped circuit. Down-select using orientation/microgravity, capillary limit, start-up, freeze/restart and leak tests.

07. Power and data cross bay boundaries through protected flexible harness. A failed thermal tile should be isolated without opening coolant to neighbouring tiles; local fault containment requires detailed interfaces.

08. Needed test article: one full-size powered tile under hot/cold thermal vacuum with representative electronics footprints, coating, shielding, deployment connection and shutdown/restart sequence.

OC-250 / ORBITAL COMPUTE PLATFORM

ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE

All dimensions nominal unless stated. No approval or qualification is implied.

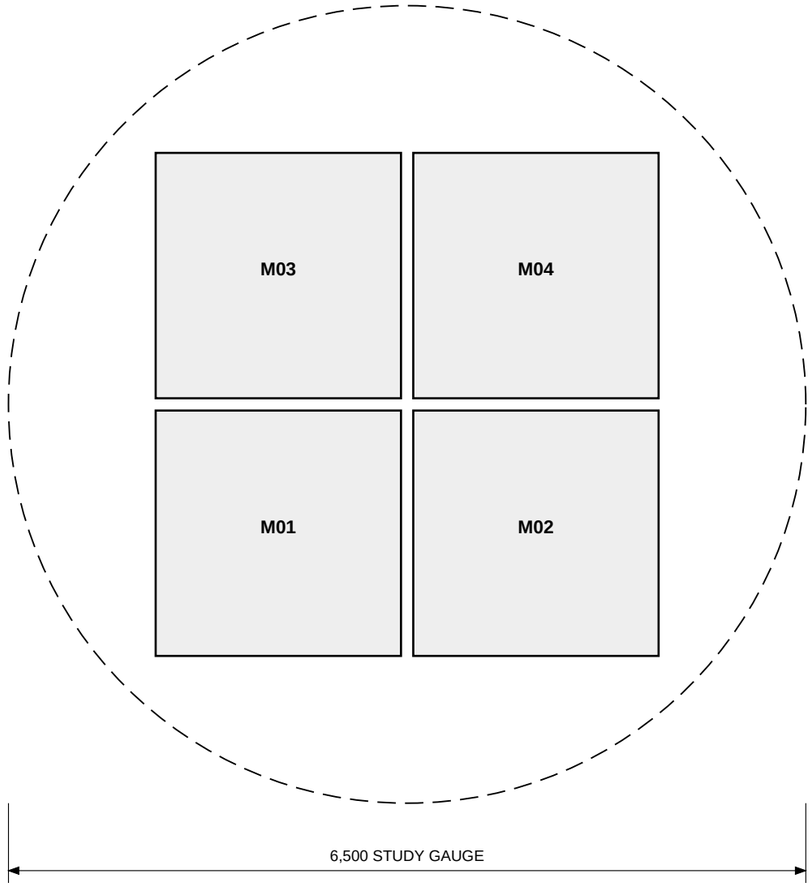
DRAWING	SCALE / FORMAT	REVISION / UNITS	SHEET
OC250-103	**1:10**	**C / mm**	**04 / 06**
ISSUED 2026-09-25	A3 420 x 297	3RD ANGLE *	* axes on views

FOUR-STACK LAUNCH ACCOMMODATION AND WEIGHTED PITCH

Study gauge only: 6,500 diameter x12,000; no launcher interface or connected-fold fit established

FOUR SEPARATE STACKS / X-Y

SCALE 1:50 | DIMENSIONS mm

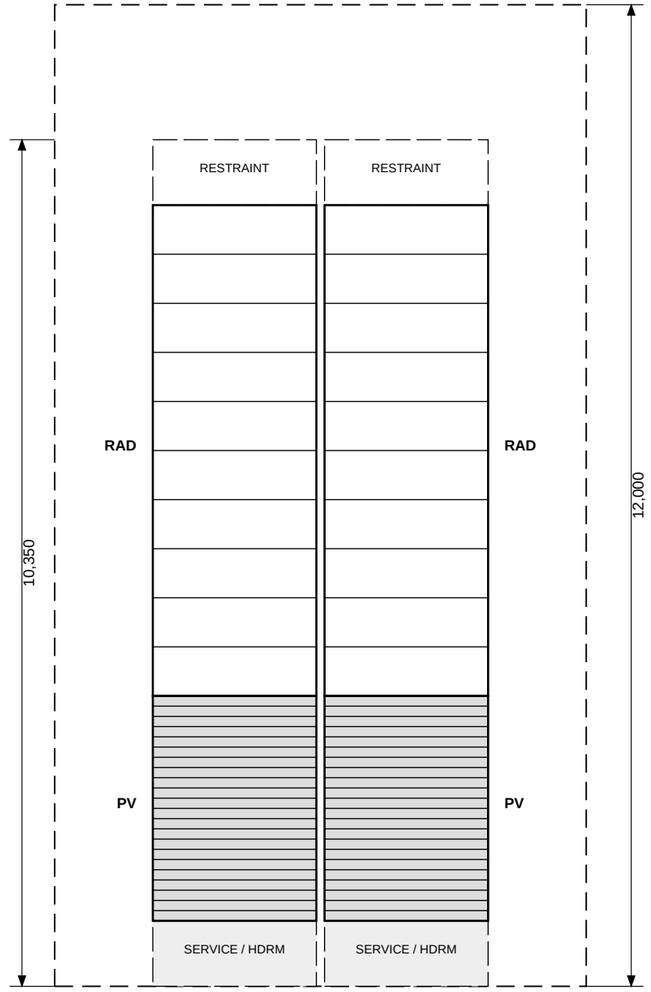


Static outer-corner clearance: 351 mm for this square-stack arrangement.

CONTENTS	PITCH mm	HEIGHT mm
110 PV	25	2,750
50 RAD	120	6,000
Service bay	In common	allowance
Service / restraints	-	1,600
Total / stack	-	10,350

AXIAL PROJECTION / X-Z

SCALE 1:75 | DIMENSIONS mm



Two stack rows overlap in projection

DESIGN / INTERFACE NOTES

01. Exactly four identical spacecraft are illustrated. If module count changes, the four-stack drawing is a reference and a new multi-payload accommodation must be generated; no automatic packing benefit is claimed.

02. Each of the four stacks contains 160 active tiles. Weighted height = 110 x25 +50 x120 = 8,750 mm. Including 1,600 mm for service hardware and restraints gives 10,350 mm.

03. The solid rectangles are unconnected flat stacks. They show one volume arrangement; they do not demonstrate that the permanently wired deployed grid folds into the same arrangement.

04. PV 25 mm and radiator/compute 120 mm are sensitivity inputs. The service package is included in the 1,600 mm common allowance. It is not an extra 500 mm panel. Reserve plan cells are empty apertures, not secretly stacked blank panels.

05. Grouping all PV plates and all radiator plates is used only for the area/thickness ledger. A real preconnected mixed grid constrains layer order, offsets, wiring folds and deployment path.

06. 100 mm between nominal square stacks is a static separation. Adapter, restraints, carrier rails, fairing taper and coupled launch deflection require additional integration geometry.

07. No mass-to-orbit capability, specific launcher acceptance or separation reliability is implied by the 6,500 x12,000 cylinder. Four separate spacecraft need verified dispensers and independent release corridors.

08. If a sequence requires robotic repacking, fluid reconnection or transfer around occupied payloads, explicitly add its hardware, development cost and failure consequences before calling the architecture feasible.

OC-250 / ORBITAL COMPUTE PLATFORM

ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE

All dimensions nominal unless stated. No approval or qualification is implied.

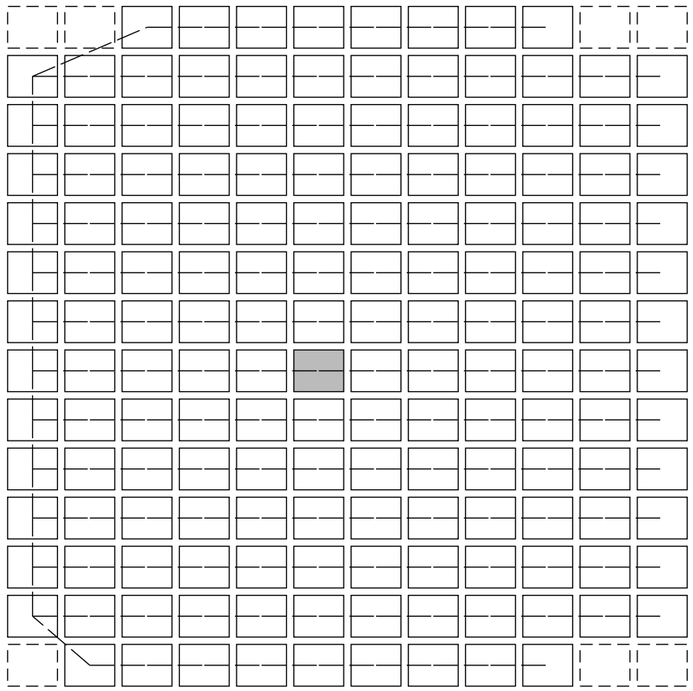
DRAWING	SCALE / FORMAT	REVISION / UNITS	SHEET
OC250-104	**1:50 / 1:75**	**C / mm**	**05 / 06**
ISSUED 2026-09-25	A3 420 x 297	3RD ANGLE *	* axes on views

DEPLOYMENT TOPOLOGY AND FINITE-THICKNESS GATES

Connectivity is not kinematics; both candidates require live-harness and thick-panel deployment trials

A ROW-PACK / BOOK TREE

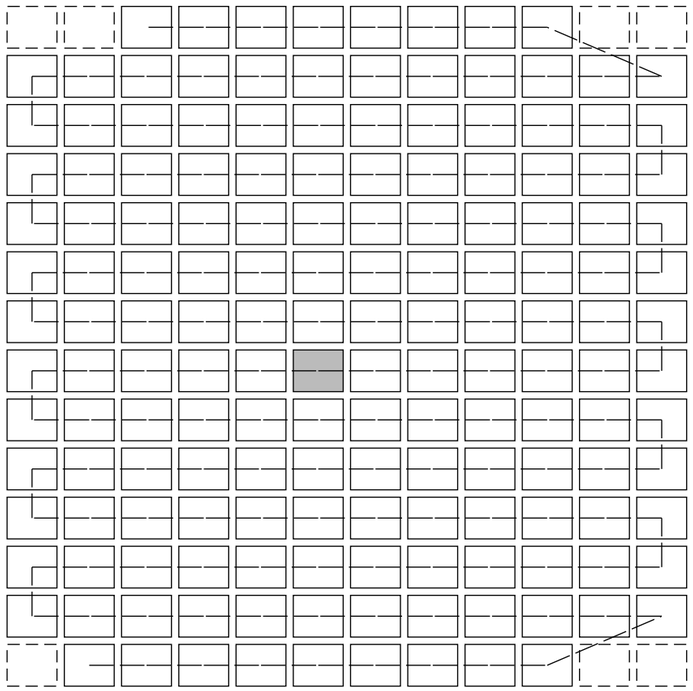
SCALE SCHEMATIC NTS | DIMENSIONS mm



160 tree links / 161 occupied bays

B SERPENTINE CHAIN

SCALE SCHEMATIC NTS | DIMENSIONS mm



160 tree links / 161 occupied bays

DESIGN / INTERFACE NOTES

01. The diagrams compare mechanical connection graphs, not trajectories. A loop-free tree helps avoid unwanted closed-chain constraints during deployment, but does not prove that finite-thickness panels avoid collision.

02. The 12 x14 footprint has 161 occupied bays including service and 7 empty reserve apertures. Either illustrative tree has 160 links; 296 occupied-neighbour edges exist. These are connectivity counts, not a hinge bill of materials.

03. A folds row panels into row packs before manipulating thicker row packages. B is a single long kinematic chain arranged into a compact final rectangle. Root jamming and recovery consequences differ.

04. Mixed pitches are significant: 25 mm PV and 120 mm radiator packages differ by a factor of 4.8. An ordinary zero-thickness map fold cannot be copied directly. Hinges must accommodate layer order and changing offsets.

05. The service bay holds a fixed safety/communication/propulsion package, included within the common launch allowance. It is not a folded thin tile. Its deployed central location and launch position require a real attachment/motion plan.

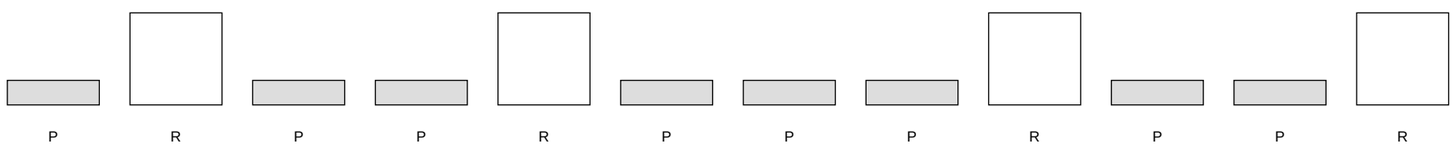
06. No common moving coolant loop is assumed between tiles. Power/data harness still require bend-radius, strain-relief, insulation, wear and cycle-life qualification under vacuum and temperature extremes.

07. Recommended next article: 3 x3 or 4 x4 mixed-thickness tiles with representative heat hardware, active electronics, latches and working harness. Perform restrained launch-environment tests before full deployment trials.

08. OPEN gates: collision-free motion; recovery from failed release; coupled structural modes; tip deflection; local thermal leakage; safe-mode heat survival; actual payload restraints and independent spacecraft separation.

UNEQUAL PACK THICKNESS / ONE ROW

SCALE SCHEMATIC NTS | DIMENSIONS mm



Unequal PV / radiator thickness changes hinge offsets and the fold path.

296 occupied-neighbour edges; 136 extra connections could add bracing/latching. Empty bays receive no hidden panels.

OC-250 / ORBITAL COMPUTE PLATFORM **ARCHITECTURE STUDY - NOT RELEASED FOR MANUFACTURE** All dimensions nominal unless stated. No approval or qualification is implied.	DRAWING **OC250-105** ISSUED 2026-09-25	SCALE / FORMAT **NTS** A3 420 x 297	REVISION / UNITS **C / mm** 3RD ANGLE *	SHEET **06 / 06** * axes on views