ORBITAL DATA CENTERS: POWER AND THERMAL MANAGEMENT FOR SCALABLE ARCHITECTURES

REDWIRE

RDW.COM

ACKNOWLEDGMENTS

"ORBITAL DATA CENTERS: POWER AND THERMAL MANAGEMENT FOR SCALABLE ARCHITECTURES" was authored by Al Tadros, Dana Turse, Joe Footdale, Derek Hengeveld, Jeff Manson, and Nathan Mullins.

EXECUTIVE SUMMARY

Orbital data centers (ODCs) represent a transformative opportunity to expand the future of global computing beyond the limitations of terrestrial infrastructure. As demand for artificial intelligence, cloud computing, and high-performance data processing accelerates worldwide, space-based compute architectures offer access to abundant solar energy, global connectivity, and a pathway toward scalable digital infrastructure unconstrained by many of the physical, environmental, and societal challenges facing Earth-based data centers.

This paper demonstrates that the key technologies needed to enable scalable orbital computing systems already exist today. Drawing on decades of heritage in spacecraft power systems, deployable structures, thermal management, and on-orbit operations, the industry is well positioned to address the engineering challenges associated with large-scale compute in space. Rather than representing barriers to adoption, power generation and thermal management emerge as solvable architectural drivers that can guide the development of practical, efficient, and scalable systems.

The analysis shows that successful ODC architectures depend on treating electrical power generation, distribution, and thermal rejection as tightly integrated components of a unified energy system. As compute capability scales, the efficient transport of electrical power and waste heat becomes increasingly important to system performance, influencing spacecraft configuration, deployable architecture, structural dynamics, and operational design. By evaluating these factors early in the architectural process, developers can identify scalable solutions that balance performance, efficiency, mass, and operational flexibility.

The paper highlights several promising approaches for enabling near-term and future orbital compute platforms. Flight-proven deployable solar array systems, advanced passive radiator technologies, and modular spacecraft architectures already provide strong foundations for scalable systems operating at meaningful compute levels. The study also identifies modular tiled architectures as a particularly compelling long-term solution, integrating compute, power generation, and thermal transport locally within self-contained elements designed to unlock the scalability required for next-generation orbital data centers.

Importantly, the findings reinforce that orbital computing is enabled through thoughtful system-level engineering and architectural integration. With continued advances in deployable systems, launch capability, spacecraft autonomy, and modular design, orbital data centers have the potential to become a foundational layer of future digital infrastructure, supporting AI services including training and inference, high-performance computing, and entirely new classes of space-enabled applications.

Orbital data centers (ODCs) have the potential to not just complement but ultimately overtake terrestrial data centers as the dominant architecture for large-scale computing.

In space, power is abundant, and systems are freed from the physical, environmental, and societal constraints that limit growth on Earth, including land availability, infrastructure bottlenecks, and the ever-present "not in my backyard" challenge. ODCs composed of vast distributed constellations or massive, modular platforms have the potential to operate at unprecedented scale to enable an entirely new class of digital infrastructure. From orbit, these systems could deliver global AI services and high-performance computing capabilities with a reach and capacity that far exceed what is practical on the ground. ODCs will support the exponential growth in global demand for compute and fuel the next era of the data-driven economy.

Recent introduction of space-based compute payloads by NVIDIA Space-1 Vera Rubin module, AMD Virtex and Versal, Google's Trillium, SpaceX D3, and others underscore the urgency of moving beyond conceptual discussions toward practical system design. As compute capability increases, both power delivery and heat rejection become dominant drivers of system architecture, introducing constraints that are not yet widely addressed in current discourse. Two prevailing assumptions in many public concepts are that power is effectively unlimited in space and that solar arrays can simply be scaled to meet demand. While solar energy is indeed abundant in orbit, this perspective underestimates the system-level realities of generating, distributing, and utilizing large amounts of power. Similarly, thermal management is often cited as a major challenge for space-based compute. While space-based compute presents meaningful challenges, this paper will show that they are entirely addressable. The technologies and expertise needed to overcome these hurdles already exist, and continued progress is well within reach. The aim of this whitepaper is not to dismiss ODC concepts or overstate their complexity, but to contribute constructively to the conversation with a clear-eyed, solutions-focused perspective. Nor is this paper intended to discredit ODC concepts or overemphasize their difficulty. Rather, by considering the different architectures it is possible to leverage heritage and experience in this domain to develop a practical solution with existing technologies.

GENERAL CONSIDERATIONS IMPACTING POWER AND THERMAL SYSTEMS

ODC architectures should be evaluated from first principles, as an energy system, in which electrical power generation, distribution, and thermal rejection are tightly coupled to compute performance. Figure 1 illustrates this energy system at a high level, where solar energy is converted to electrical power through deployable arrays and routed through the power management and distribution system to the compute payload. Nearly all of this electrical energy is ultimately converted to heat, which must then be transported and rejected through a dedicated thermal management system and radiated to space.

While this flow may appear straightforward, each stage introduces losses, inefficiencies, and design constraints that become significant as system power levels grow. This separation of key subsystems (power and thermal) requires the system to transport power and heat across increasingly large distances and complex interfaces as compute demand increases, introducing fundamental scaling challenges. Understanding these challenges is essential to evaluating the practicality of conventional spacecraft architectures for large-scale ODCs, and these challenges are where leveraging lessons learned will address the breakdown of this paradigm.

Note that this assumes a conventional spacecraft paradigm, where a centrally located compute payload is powered and cooled using dedicated solar arrays and heat rejection hardware. This paradigm represents some of the most prominent concepts currently in the public domain, particularly those involving deployment of large constellations of spacecraft. However, alternative concepts, such as compute "tiles," have compelling characteristics from a power and cooling perspective and will be briefly discussed at the end of this paper.

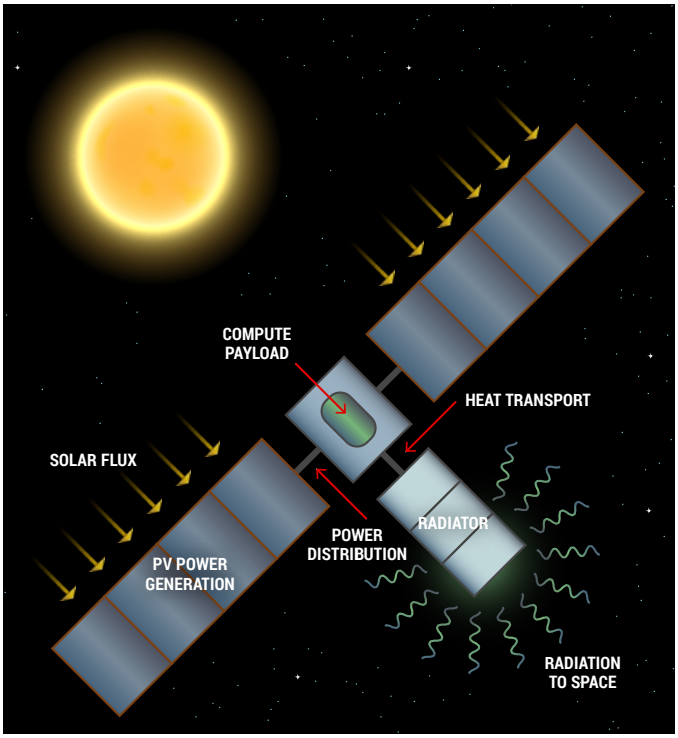

Figure 1. Illustration of ODC as an energy system.

Before sizing and scaling power generation and thermal rejection systems, it is important to consider operational orbit as a primary driver. Orbit directly determines solar availability, eclipse duration and frequency, battery dependence, thermal sink geometry, duty cycle, spacecraft pointing constraints, communication latency and connectivity, and bandwidth. As a result, it is not simply an environmental input, but a first-order design variable that fundamentally shapes system architecture. Across the range of proposed public concepts, distinct orbital regimes present different trade spaces. Low Earth orbit (LEO) altitudes of ~500–700 km offer advantages in latency, radiation environment, and responsiveness, but this regime also has high eclipse fractions, driving significant battery dependence and imposing large swings in thermal conditions as the system cycles between sunlight and eclipse. As altitude increases into the ~700–1200 km regime, eclipse fraction decreases and power continuity improves; however, systems remain power-limited and must still accommodate dynamic thermal environments, albeit with somewhat reduced Earth infrared and albedo heating. Sun-synchronous orbits (SSO), particularly near the terminator in the ~600–900 km range, are often viewed as a balanced solution. These orbits provide high solar availability with minimal eclipse duration, reduce albedo loading, and enable more stable thermal conditions. However, they introduce constraints on spacecraft orientation and require careful radiator design to mitigate direct solar loading. At higher SSO altitudes (~1000–2000 km), systems can approach near-continuous sunlight during favorable conditions, improving power availability, but at the cost of increased radiation exposure and higher latency.

The key takeaway is that no single orbit is universally optimal. Instead, each regime shifts the balance between power availability, thermal performance, operational constraints, and system complexity. Orbit selection cannot be decoupled from power and thermal design as it directly drives system sizing, architecture, and ultimately the achievable scale of an ODC platform. This point is illustrated in Figure 2, where solar array sizing is plotted as a function of altitude in SSO. Importantly, the plot shows that at higher SSO altitudes, increased exposure to radiation environments, particularly the inner Van Allen belt, can drive significant degradation in solar cell performance over time, offsetting gains from improved illumination conditions.

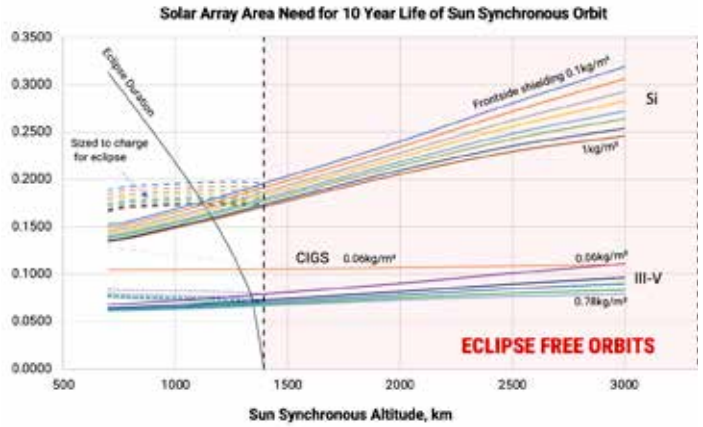

Figure 2. Solar array area needed for a 10 year life in a Sun Synchronous Orbit.

This introduces the topic of solar cell chemistry to the trade space. High-efficiency III-V (GaAs) cells minimize array size and offer strong radiation tolerance, but their cost and limited production capacity constrain scalability for ODC-class systems. Lower-cost silicon (Si) cells scale economically but require substantially larger arrays due to lower efficiency and greater radiation sensitivity, amplifying system-level challenges. Intermediate options such as copper indium gallium selenide (CIGS) and perovskite offer a balance of cost, radiation tolerance, and packaging efficiency, and are coming onto the market as a viable options.

Thermal rejection performance is also strongly driven by orbital parameters. While radiators reject heat based on temperature, emissivity, and view to deep space, they are also exposed to environmental loads such as solar radiation, Earth infrared, and albedo, all of which reduce net rejection capability. These effects vary with altitude, inclination, and geometry. In SSO, radiator orientation relative to the Sun and Earth is critical; configurations that maximize deep space view and minimize solar exposure (e.g., edge-on to the Sun) significantly improve performance. Earth-driven heating decreases with altitude, directly impacting radiator sizing and placement. As a result, thermal design must account for both internal heat loads and the external orbital environment. For SSO, albedo loads are minimized at the terminator while a flat radiator, edge-on to the Sun design minimizes solar loading. In this configuration the predominant loading is Earth infrared with either edge or face to Earth designs to help minimize environmental loading (~50 to 100 W/m^2 for LEO).

POWER AND THERMAL SCALING CONSIDERATIONS

Power Scaling

At first glance, scaling solar array power appears straightforward—simply add additional photovoltaic (PV) area to increase total power. In practice, however, the method for implementing this scaling introduces important system-level constraints that become increasingly significant at higher power levels. Different architectures are explored to address this consideration for example using a tiled design and or directly addressing the distance of PV's in a monolithic design.

One of the most direct methods for increasing solar array output is to add cells in series, thereby increasing string voltage while maintaining lower current for a given power level. This approach is attractive because resistive losses in electrical distribution scale with the square of current, making higher-voltage operation inherently more efficient. However, spacecraft power systems are not free to increase voltage indefinitely. Historically, solar array operating voltage has typically been limited to a relatively low value (e.g., less than 200 V), driven by considerations such as arcing risk in plasma environments, insulation complexity to avoid snapover conditions, fault protection requirements, and long-term reliability concerns. However, credible designs for much higher voltage solar arrays are being demonstrated by increasing cell separation and incorporating highly insulative grout between cells, approaches that decrease arcing risk and allow much higher operating voltages. This, of course, comes at the cost of increased array area for a given power level, but in the world of larger fairings, lower launch costs, and relaxed first modes of deployable arrays, this does not necessarily represent a practical design constraint.

If operating voltage is capped at a given value, scaling for additional power must instead be achieved by increasing the number of parallel strings rather than continuing to extend string length. While this maintains voltage, it introduces a second-order effect: power must be collected across increasingly large physical areas and transported over longer distances back to the spacecraft power management and distribution system.

This distributed collection architecture creates a fundamental scaling challenge. As array size increases, the distance between outermost generation points and the central bus grows, increasing harness resistance. In a more physically representative model, resistance scales with conductor length and inversely with conductor cross-sectional area, while total path length includes both supply and return conductors. Even with modest current levels, these effects can drive significant losses as distances increase. The voltage drop in the harness quickly exceeds the voltage contribution of a single cell, prompting the need for increasingly longer string lengths farther away from the bus, eventually reaching a point of diminishing return.

The impact of transmission distance and operating voltage on power loss is illustrated in Figure 3. At a fixed power level, lower-voltage systems experience substantially higher losses due to increased current. For example, at low bus voltages, losses increase rapidly with distance and quickly become prohibitive.

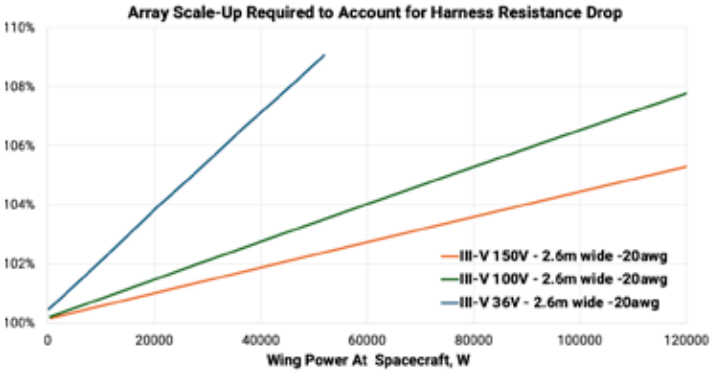

Figure 3. Array scale-up required to account for harness resistance drop.

Higher-voltage systems mitigate this effect, but remain constrained by practical voltage limits, and losses still grow with increasing transmission distance. This highlights both the benefit of operating at higher voltage and the limitations imposed by spacecraft electrical design constraints.

As resistive losses increase, voltage drop across the harness becomes significant, reducing the effective power delivered to the spacecraft bus. One potential mitigation is to increase conductor cross-sectional area and use it at a fraction its current carrying capability, thereby reducing resistance and limiting I^2R losses. While effective from an electrical standpoint, this introduces a competing penalty in the form of increased harness mass. Because conductor mass scales directly with cross-sectional area, reducing losses through larger conductors results in a linear increase in harness mass, which can become substantial for long runs and high-power systems.

The relationship between conductor size and harness mass is shown in Figure 4. For representative transmission distances, harness mass grows quickly as conductor area is increased to control losses. In large-scale systems with multiple long distribution paths, this mass becomes a significant contributor to overall system mass. These impacts extend beyond simple mass growth—larger harnesses also affect stowage efficiency, structural design, and deployed dynamic performance.

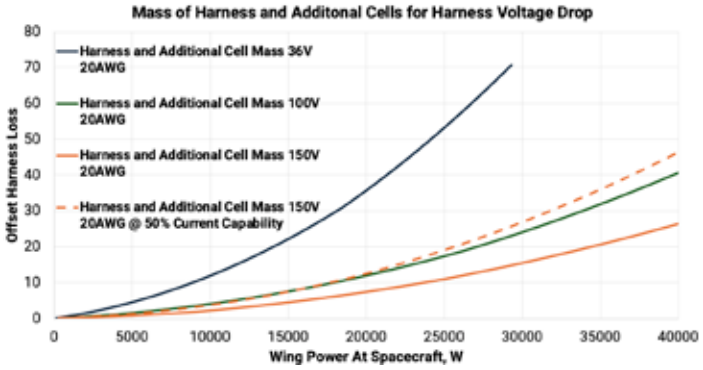

Figure 4. Mass of harness and additional cells to accommodate for harness voltage drop.

Taken together, these effects define a fundamental trade space between electrical efficiency and system mass. Reducing resistive losses requires larger conductors, but doing so imposes penalties that propagate into structural, packaging, and control subsystems. Conversely, limiting harness mass results in increased electrical losses and reduced delivered power.

Ultimately, the competing effects—of voltage limits, transmission distance, resistive losses, and harness mass introduce a point of diminishing return in the scaling of power generation via a single solar array. As array size increases, the system must transport more energy across greater distances at constrained voltage, leading to compounding inefficiencies that limit practicality. In summary, there is a practical break point for solar array sizing. For a fixed transmission distance and a representative upper-end operating voltage, maintaining a modest harness loss target requires steadily increasing conductor size and corresponding harness mass. If conductor size is held fixed instead, distribution losses rise rapidly with increasing power. Together, these trends define a practical limit at which further scaling of a single centralized power system begins to impose disproportionate penalties in both mass and efficiency.

THERMAL REJECTION SCALING

At first order, radiator area scales with compute power. For a given allowable electronics temperature, higher heat loads require larger radiating surfaces in order to maintain thermal balance. This relationship is fundamental and unavoidable. While it may be tempting to view radiator growth as a simple area-scaling exercise, the challenge is not limited to radiator sizing alone.

For a centralized compute module, extracting heat from the electronics and moving it to remote radiators presents several practical challenges. High local heat fluxes must be managed without creating unacceptable component temperatures and excessive thermal gradients that lead to thermal expansion mismatch. As radiator area grows, the distance between the compute module and the outermost rejection surfaces increase, making the thermal transport network larger, heavier, and more complex. In this sense, the thermal problem mirrors the electrical distribution problem discussed previously—energy must be transported across physically separated subsystems, which introduces penalties that grow with system scale.

Radiator performance also depends on more than just area. The efficiency of a thermal rejection system is strongly influenced by radiator operating temperature, surface coatings, view factors, orientation, thermal interfaces, and heat transport losses. Simple

radiator panels are capable of radiating ~250W/m^2 using common materials and engineering design. These radiator designs are common, thus enabling architectures like modular bays that can be tiled into large arrays. Higher radiator temperatures improve rejection capability, but they may not be compatible with compute hardware temperature limits or with the desired thermal architecture. Similarly, radiators must maintain favorable views to cold space; self-shadowing, obstruction by other spacecraft elements, and attitude constraints can all reduce effectiveness. Operational orientation therefore becomes part of the thermal design problem, particularly as radiator systems become larger.

From a practical implementation standpoint, the preferred solution is often to avoid large deployable thermal systems whenever possible since static, passive radiators are generally less complex, lower cost, and lower risk. However, for meaningful compute loads, fixed/passive radiator area may be insufficient, especially once other spacecraft surfaces are already allocated to solar arrays, communications hardware, sensors, or structural functions. This creates pressure to move toward deployable rejection surfaces.

Deployable passive radiator systems can also expand rejection area and have been successfully used in high-power space applications. However, they come with several practical drawbacks, including additional mass, packaging demands, and deployment mechanisms. As radiator size grows, impacts propagate into the broader spacecraft design, influencing factors such as stowage efficiency, dynamic performance, control stability, and mission cost.

Active pumped fluid loops (PFL) provide another path to transport heat over larger distances, but they introduce another set of challenges, including pump reliability, plumbing complexity, fluid management, startup and control behavior, and additional parasitic power. These systems are well understood and have clear heritage, but they are not trivial, particularly when extended over large deployable structures or incorporated into architectures intended to scale significantly beyond traditional spacecraft thermal loads.

These considerations show that thermal rejection is not simply a matter of adding more radiator area. As compute power increases, the spacecraft must manage a tightly coupled problem involving concentrated heat acquisition, thermal interface and transport path design, radiator sizing, operating temperature, geometry, and operational orientation. These effects compound as the system grows, ultimately creating practical limits on how much compute power can be hosted on a single orbiter using conventional thermal architectures.

STRUCTURAL AND DYNAMIC IMPACTS OF LARGE DEPLOYABLES

The scaling of power generation and thermal rejection on a single-orbiter platform inherently drives the use of increasingly large deployable structures, including solar arrays and radiator panels. While these systems enable the required surface area for energy generation and heat rejection, they also introduce important structural and dynamics considerations that become more pronounced as system size increases. In particular, large, lightweight deployables tend to exhibit low-frequency flexible body modes, which can couple with spacecraft attitude dynamics and control systems. These flexible modes present challenges for pointing stability and control, especially for missions requiring high-precision attitude knowledge and stability such as those employing optical communication links. Structural deflections, vibrations, and dynamic coupling between appendages and spacecraft bus can degrade pointing performance and increase control effort. As deployable span and mass grow, maintaining adequate stiffness without incurring significant mass penalties becomes increasingly difficult, and the interaction between structural dynamics and attitude control must be carefully managed through design and analysis.

Advances in spacecraft control systems, as well as the development of isolation platforms and precision pointing mechanisms, have significantly improved the ability to mitigate these effects. However, these solutions introduce additional system complexity, mass, power consumption, and integration challenges. As a result, structural and dynamic considerations should not be treated as secondary effects, but rather first-order drivers that can influence architecture selection, subsystem sizing, and overall system cost. In the context of large-scale ODCs, the interplay between deployable structures, dynamic response, and control performance represents an important constraint that must be considered alongside power and thermal scaling.

PACKAGING AND STOWAGE CONSTRAINTS

In addition to power, thermal, and structural considerations, the scaling of large deployable systems is also constrained by the practical realities of launch vehicle packaging and stowage. Solar arrays and radiator systems must be designed to fit within the available launch vehicle fairing volume in a stowed configuration, which imposes limits on both overall size and allowable configuration complexity. As deployable area increases, achieving efficient packaging becomes increasingly challenging, often requiring more intricate packaging schemes and additional deployment mechanisms.

Historically, launch vehicle constraints, particularly fairing diameter and usable volume, have been a primary driver

in the design of large deployable spacecraft systems. With the emergence of next-generation launch vehicles such as Starship and New Glenn, these constraints are expected to relax significantly. Increased fairing volume and mass-to-orbit capability will reduce the need for extreme packaging efficiency and enable larger monolithic structures or less tightly folded deployable configurations. However, packaging constraints do not disappear entirely. Even with larger fairings, there remain practical limits on deployable size, aspect ratio, and configuration complexity. Extremely large or highly intricate deployable systems still introduce risk in deployment reliability, integration, and test. Additionally, larger stowed configurations may complicate structural support during launch, interface design, and ground handling. As a result, while next-generation launch vehicles expand the feasible design space, launch packaging and stowage considerations continue to play an important role in shaping system architecture, particularly for very large power and thermal subsystems.

BASELINE ARCHITECTURE POINT DESIGN

This white paper presents a representative point design for a near-term orbital compute node based on a single-spacecraft architecture. This required a set of simplifying assumptions, starting with the compute payload. Recent announcements by NVIDIA provide a useful reference, including the Space-1 Vera Rubin module introduced at GTC 2026 as part of a broader space-based AI computing portfolio. While detailed SWaP specifications have not been publicly released, the platform is described as enabling data center-class AI in space. This study assumed a compute payload power of 8 kW, with an additional 2 kW allocated to spacecraft subsystems (attitude control, propulsion, and communications), resulting in a total system power requirement of 10 kW. This study also selected a 1,500 km SSO altitude as the baseline operating regime in order to appropriately size systems.

Figure 5 presents the resulting point design. The architecture adopts a flat-stack bus configuration with two deployable solar array wings and a deployable radiator platform, onto which the compute platform is directly integrated. The configuration and sizing of the power and thermal subsystems is described in the sections to follow.

A primary driver of the configuration was compatibility with existing launch vehicle constraints, specifically packaging within a Falcon 9 fairing, where vertically stacked configurations provide better packing efficiency than radial arrangements. The spacecraft structure is optimized to minimize height while providing sufficient stiffness for radiator support and accommodating core subsystems. In the Falcon 9 configuration, overall vehicle height is driven by radiator dimensions overlapping the stowed solar arrays, resulting in an approximately 1.1 m tall vehicle and enabling six-unit stacking. For larger fairings, such as New Glenn, increased lateral spacing between arrays allows the radiator to be nested between them, reducing overall height and improving packaging efficiency.

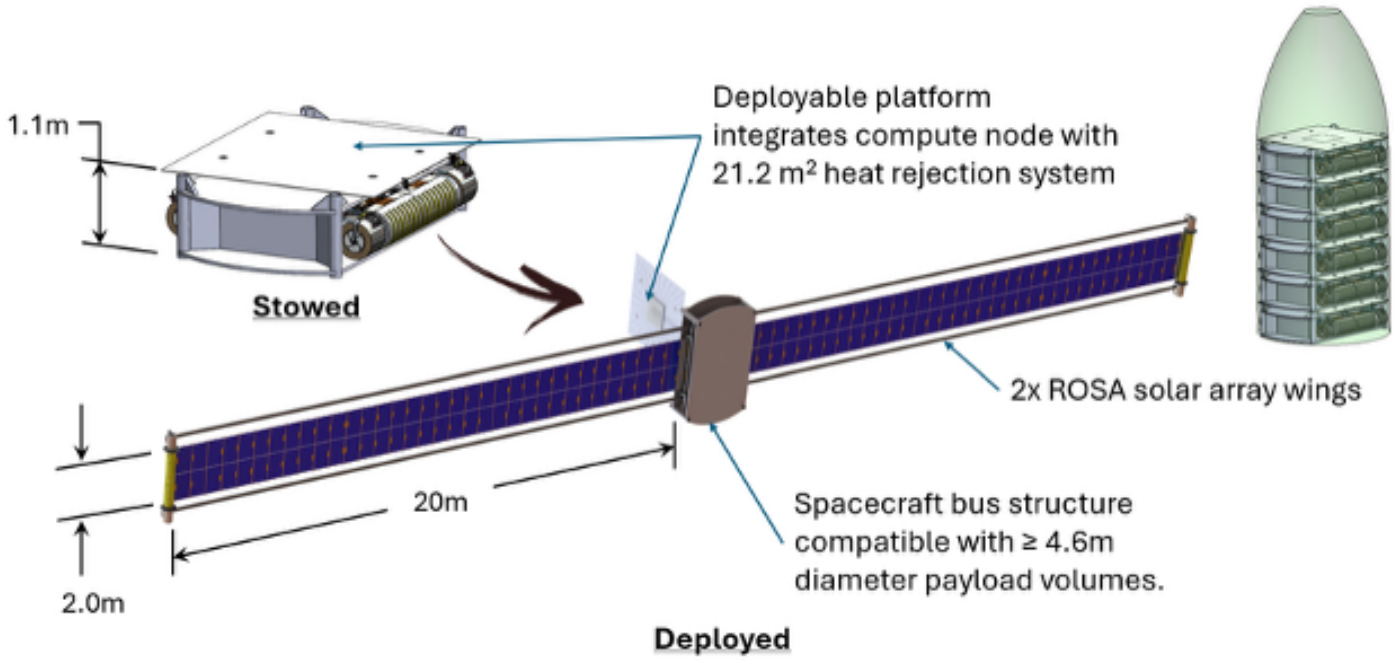

Figure 5. Initial point design.

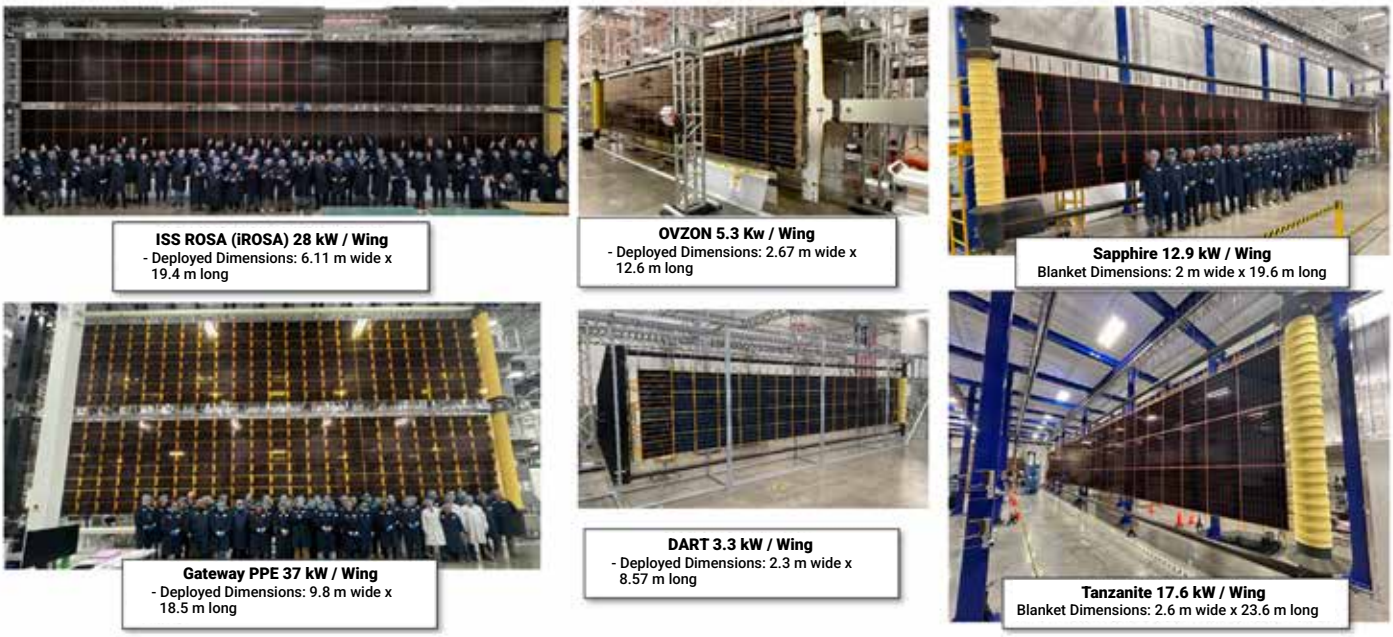

Figure 6. Heritage ROSA solar arrays.

POWER ARCHITECTURE

The solar arrays illustrated in this concept are Redwire's Roll-Out Solar Array (ROSA) architecture. ROSA provides a versatile, flight-proven solution for high-power space applications, with demonstrated scalability across a wide range of missions and environments. Its architecture, based on elastically deployable slit-tube booms and a flexible PV blanket, enables a highly modular system that can be scaled in size without introducing significant mechanical complexity. Unlike traditional rigid-panel or multi-hinged deployable arrays, ROSA maintains a consistent deployment approach as it grows, avoiding additional hinges, motors, and complex kinematic mechanisms. The result is a system that is lighter, more compact in its stowed configuration, and inherently well-suited for large-area power generation.

Meeting the 10 kW power need at a 1,500 km SSO altitude requires a Si solar cell array power generating area of approximately 80 m^2. The scalable ROSA array architecture was baselined for its compact stow volume, high deployed structural performance, high mass efficiency, and heritage. Redwire has existing configurations of 2 m x 20 m blanket arrays, with options to scale as needed.

Scaling the solar array area can be achieved by increasing the deployed width and length dimensions, or by increasing the number of blanket assemblies. Redwire has heritage designs with 2.6 m-wide blanket widths that can be directly utilized at the expense of stowed dimensions. Increasing the length imposes some practical limits on manufacturing as well as spacecraft control stability, with low flexible modes with relatively high mass participation. This factor, which was not investigated in this study, is also applicable to scenarios that increase the number of arrays.

THERMAL ARCHITECTURE

In this architecture, the compute payload is directly coupled to the radiator assembly to improve thermal efficiency, simplify thermal transport, and reduce system complexity. Structural design of the radiator assembly is driven by launch load requirements, necessitating a deployable configuration. The radiator is stowed against the spacecraft body using hold-down and release mechanisms (HDRM) and deployed via damped spring-loaded hinges. This approach minimizes system complexity by routing only electrical harnessing across hinge lines, avoiding the need for the flexible fluid connections required in PFL systems.

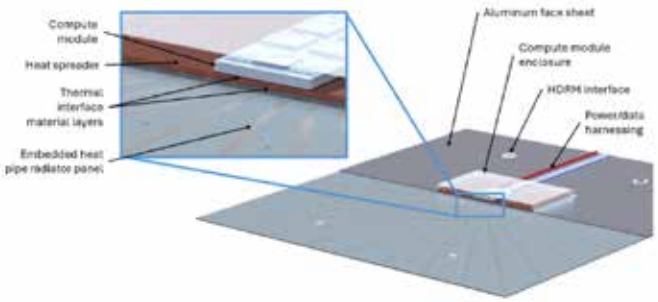

Figure 7. Concept for radiator panel assembly.

A section view of the proposed radiator panel assembly is shown in Figure 7. It is assumed that the compute module consists of the GPU/CPU/on-board memory components, where auxiliary storage, networking, and other support hardware are remotely located within the spacecraft. The assembly design does not rely on this configuration and can be revisited with anticipated impacts to overall size and structural implications. The compute module is coupled to a 1 m x 1 m heat spreader component (e.g., oscillating heat pipe or vapor chamber) that is mounted to a 3.25 m x 3.25 m radiator sandwich panel. The two-sided radiator (21.125 m^2 rejection area) consists of space-grade high-emissivity coated aluminum face sheets (e.g., AZ-93 white paint with $\alpha = 0.15$ and $\varepsilon = 0.85$), with internal radially arranged heat pipes (e.g., axial grooved constant conductance heat pipes or oscillating heat pipes) and core fill such as POCO® foam or aluminum honeycomb. A series of metallic inserts are integrated for HDRM and hinge mechanism interfaces.

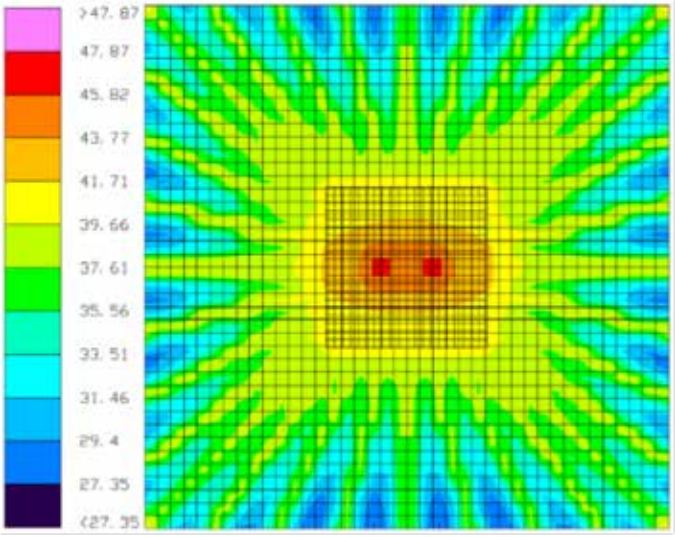

Figure 8. A 3.25 m x 3.25 m two-sided heat-pipe embedded radiator with efficient thermal spreader can manage the 8 kWt heat load.

Thermal analysis was conducted on this architecture. A 1 m x 1 m thermal spreader was modeled to provide the interface between the 8 kWt heat load and radiator. A radial array of heat pipes was attached to the radiator to provide thermal transport/spreading and improve radiator efficiency. This system (with bus and solar arrays) was placed in a 1,500 km SSO (i.e., Beta Angle = 90°). All thermal interfaces were high-performance. Steady-state results (Figure 8) demonstrated that under these conditions, the maximum temperature of the compute module interface was 48°C (a similar analysis with a 10 kWt heat load had a maximum temperature of 66°C). Further optimization would improve radiator efficiency while minimizing mass.

To further scale the heat rejection system, the radiator panel dimensions can be increased up to current manufacturing methods and heat pipe length-driven efficiencies as described earlier. One approach is to structurally combine multiple thermally independent 10 m^2 subpanels into a larger super panel. The structural design and launch constraint system would have to be designed accordingly for launch. Another approach would be to have a multi-panel deployable radiator system. Mass and stiffness of the power and data networking cabling merit further study in this approach. Another alternative replaces the passive radiator system approach with a deployable PFL radiator system where the compute system remains on the spacecraft structure.

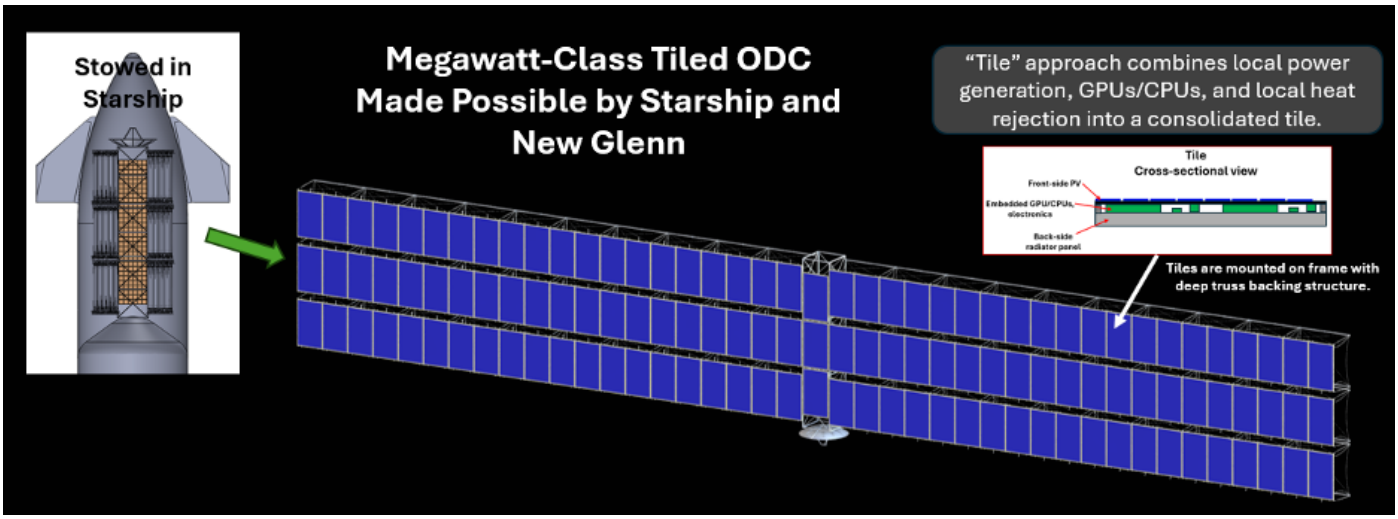

Figure 9. Concept for a megawatt-class tiled ODC.

ALTERNATIVE SCALABLE ARCHITECTURES

An alternative approach moves away from the conventional paradigm of scaling dedicated solar arrays and radiators to support increased compute capacity toward a tiled aperture architecture in which power generation, compute, and thermal rejection are co-located within a single tile. Such architectures are being actively explored by emerging companies such as Sophia Space. In this concept, each tile integrates front-side photovoltaic generation, embedded compute hardware, and backside radiative surfaces for heat rejection, forming a self-contained energy-balanced element.

This architecture directly addresses one of the primary limitations identified in this paper: the need to transport large amounts of electrical and thermal energy across physically separated subsystems. By localizing power generation and heat rejection at the tile level, the tiled approach largely eliminates long-distance power distribution and thermal transport, reducing losses, harness mass, and system complexity.

The result is an inherently modular and scalable system, where overall compute capacity is increased by adding additional tiles. Figure 9 illustrates a representative concept in which tiles are mounted onto a deployable truss structure to form a large aperture exceeding 2,000 m². The truss provides the structural depth and stiffness required to maintain stability and controllability of the deployed system, addressing challenges associated with large flexible structures. The full system can be packaged within next-generation launch vehicles such as Starship and deployed on orbit to form large-scale compute platforms.

However, this approach introduces its own design challenges. Each tile must be carefully optimized to balance compute density, power generation, and heat rejection within a fixed area, requiring tightly integrated layout of components. Sufficient front-side photovoltaic area must be allocated to meet local power demands, while backside radiator area must be sized to reject the corresponding heat load. These competing requirements define the performance envelope of each tile. But once a tile solution has been found, developed, and matured, this approach offers excellent scalability.

This white paper examined the scaling of power and thermal systems for ODCs within a single-spacecraft architecture, where power generation and thermal heat rejection are supported by separate subsystems.

A key takeaway is that while solar energy and radiative cooling are available in space, system-level constraints, and not resource availability, have a big impact on scalability. The fundamental limitation is energy transport. As systems scale, increasing separation between power generation, compute, and heat rejection introduces compounding penalties in electrical losses, harness mass, and system complexity. Thermal management, though well understood, becomes a dominant driver at higher power levels, with radiator sizing, placement, and environmental interactions tightly coupled to overall system design. These effects can create clear points of diminishing return for conventional architectures. More scalable approaches, such as tiled, modular architectures that localize power generation and heat rejection, may offer a path forward by minimizing transport distances and reducing non-linear penalties.

In general, it is recommended that power and thermal systems be treated as first-order architectural drivers. Energy transport (electrical and thermal) should be modeled and evaluated early in design trades when architectural decisions are still fluid. By incorporating these effects early, teams can build a more complete understanding of system behavior, identify practical scaling limits, and uncover opportunities to optimize architecture before constraints become difficult to unwind. This approach enables more informed trade decisions, highlights where non-linear penalties begin to emerge, and supports the development of concepts that remain robust as performance demands grow.