

SPACE-BASED DATA CENTERS ROADMAP: THE CASE FOR CAUTION

November 2026

CHAPTER 1.3:

Cooling

Alexis Abramson



Innovation for Cool Earth Forum

WORKING PAPER

TABLE OF CONTENTS

CHAPTER 1.3: 1.3-1
COOLING – *Alexis Abramson*

A. Introduction: Cooling in Space 1.3-1
B. Engineering Space-Based Thermal Management 1.3-3
i. Lessons from spacecraft thermal management 1.3-3
ii. Thermal system design 1.3-4
iii. Thermal performance metrics 1.3-8
C. Risks and Open Questions 1.3-12
D. Recommendations 1.3-13
E. References 1.3-14

CHAPTER 1.3:

COOLING

Alexis Abramson

MAIN MESSAGES

- Space-based data centers (SBDCs) can shed their computers' waste heat only by radiating it into space. This makes cooling the central limit on how much computing a satellite can do—and a major driver of cost.
- Counterintuitively, the best system design runs chips as hot as they can safely tolerate: a hotter radiator sheds far more heat, which keeps radiator panels small and launch mass low.
- Real world thermal performance may not be easily achieved over the system's lifetime since coatings degrade and debris-avoidance maneuvers may periodically require reorientation of satellites toward sources of heat.
- SBDC developers should consider cooling as a critical design constraint given thermal management dictates key system characteristics. Standard thermal performance metrics and disclosures should be developed to independently verify claims.

A. Introduction: Cooling in Space

Data centers produce significant heat. Terrestrial data centers are cooled with fluids, such as air and water, that carry heat away from the hardware by convection. However, no fluids are present in the vacuum of space. As a result, an orbiting spacecraft can shed its heat only by emitting infrared radiation—a process known as radiative cooling (see Figure 1).

Radiative heat rejection is governed by a simple principle: the hotter an object is, the more heat it radiates, and this heat output rises very rapidly as its temperature increases. This principle is captured in the Stefan-Boltzmann law, which states that power radiated per unit area is $P = \varepsilon\sigma T^4 A$, where ε is thermal emittance, σ is the Stefan-Boltzmann constant, T is absolute temperature and A is the radiating area. Because radiated power rises so steeply with temperature, a modest increase in radiator temperature yields a large increase in heat rejected. For example, taking the same radiator from 60 °C up to 90 °C adds roughly 40% to the radiated power per unit area it produces. So, the hotter radiator at 90 °C can reject the same heat load with a panel roughly 30% smaller and lighter than a cooler radiator at 60 °C. What this buys is capacity rather than efficiency: nothing about the thermodynamics improves, but the hardware shrinks, and hardware size carries real cost in an industry that must pay to lift every kilogram off the ground.

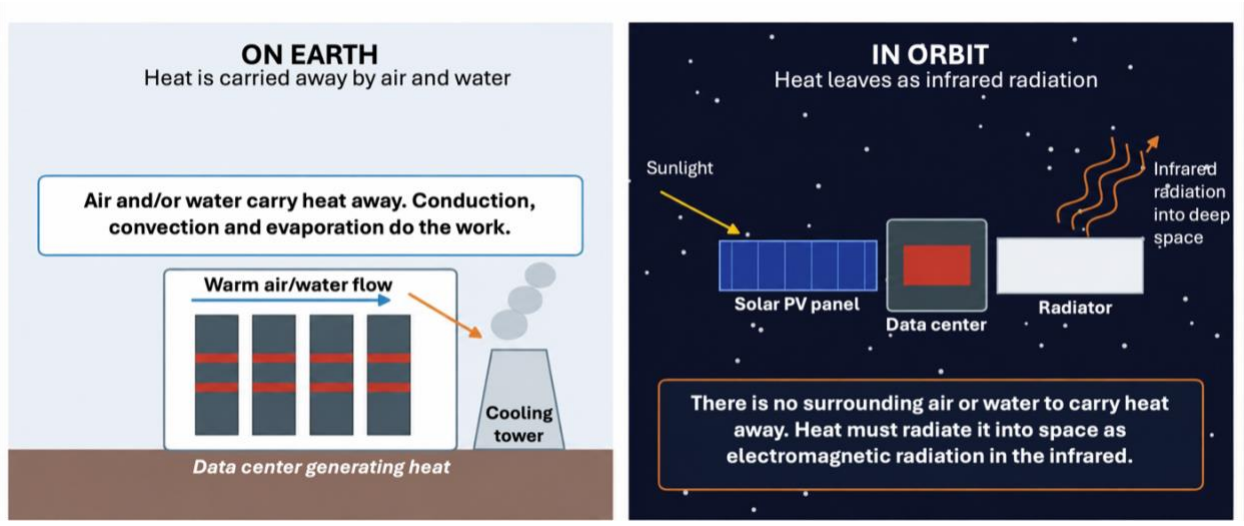


Figure 1. *Data center cooling in space works much differently than on Earth. On Earth a moving fluid carries waste heat from the servers and releases it to the atmosphere; in orbit there is no surrounding medium, so heat can leave only as infrared radiation. The diagram is a schematic and not to scale; it omits internal coolant loops, environmental heat loads and the orientation constraints examined later in this section (PV = photovoltaic). Source: AI generated using Gemini and ChatGPT*

The Stefan-Boltzmann law leads to a design conclusion that runs opposite to terrestrial practice. On Earth, running the hardware cool is always preferable, and operators keep chips well beneath their rated limits. But in space, the temperature of the radiators is typically a little bit cooler than the temperature of the chips. Any margin between the chip's operating temperature and its maximum rating lowers the radiator's operating temperature by the same amount. The design objective in space is thus to run hardware as hot as the system can tolerate, which for high-performance computing chips means operating near 85–95 °C.[1] The chip's thermal limit thus becomes a parameter designers deliberately target instead of a boundary to avoid. The challenge is that operating at the maximum temperature limit accelerates wear, shortening chip lifetime. This tradeoff must be carefully engineered into the system: lifetime specifications, redundancy provisions and replacement schedules are all set with the knowledge that, to avoid having to launch large amounts of extra radiator mass, the chips will operate for their entire service life pressed up against their rating.

A radiator operating near room temperature needs roughly 2000 m² of surface to shed a megawatt of waste heat.[2] And while space is an attractive cold sink, solar irradiance (~1361 W/m²), Earth's thermal emission (~240 W/m²) and reflected sunlight (albedo, ~30% of solar irradiance) all deposit additional heat onto exposed radiator surfaces. Further, depending on the satellite's configuration, heat radiating from the spacecraft's own warm structures can add to the load. Together these loads reduce the net heat the radiator can shed. What complicates matters is that these loads are not constant. A satellite in typical low Earth orbit (LEO) crosses between sunlight and Earth's shadow every 90 minutes—roughly 55 minutes of sunlight and 35 minutes of eclipse—which means the loads from direct sun and albedo are switched on and then off again 16 times in a single day. Two levers serve to minimize this unwanted incoming radiation: orientation, which sets what falls within the

radiator's field of view, and coatings, which determine how much radiation the radiator surface absorbs. Managing the environmental loads with these two levers is the essence of radiator engineering.

Following this introduction on how cooling in space differs, this section discusses spacecraft thermal management as a mature engineering discipline and examines the radiator architectures, heat transport systems and orientation strategies available to designers of space-based data centers (SBDCs). The section then explores significant challenges created by the cooling needs of SBDCs, including the gap between demonstrated and proposed scale, coating degradation, tradeoffs between efficient cooling and brightness mitigation, risks related to orbital debris and congestion, the unsolved problem of how best to dispose of SBDC radiator panels at the end of their useful lives, and the lack of any disclosure standard. It closes with recommendations.

B. Engineering Space-Based Thermal Management

At steady state, a space-based system must radiate away essentially every watt its solar panels generate. This constraint propagates through the entire design: it constrains the choice of orbit, fixes the satellite's attitude and much of its geometry, ties directly into the power architecture and, together with power, caps the compute a platform can economically carry. One recent industry analysis reports that radiators and solar arrays together can account for 65–70% of a satellite's total mass and that offsetting five years of coating degradation calls for ~40% additional radiator area.[3] Cooling does not sit alongside the other subsystems; it is what determines whether the data center is feasible at all.[4]

As discussed throughout this roadmap, the sector is advancing from concept to flight testing very rapidly, and thermal management design is keeping pace. In late 2025, Starcloud became the first company to operate an NVIDIA H100 in orbit,[5] cooling the chip's 700 W by means of direct-to-chip liquid cooling that carried heat out to body-mounted radiator panels, which then rejected it passively to space. Its successor, Starcloud-2, planned for launch in late 2026, will carry what may be the largest deployable radiator yet flown on a commercial mission.[6] In June 2026, SpaceX unveiled AI1: a 70-meter satellite carrying 110 m² of radiators whose panels are held at a knife-edge attitude to the sun, meaning the radiator presents almost no cross-section to sunlight.[7,8]

i. Lessons from spacecraft thermal management

Spacecraft thermal management is a well-established discipline. Satellites carry batteries, propellant, optics and electronics that function within a narrow temperature range and must react accordingly to the large thermal swings experienced in space. Over the past six decades, rigorous systems and tools have been developed to manage thermal performance, including flight-proven radiators, coatings and hardware for dissipating heat; standardized thermal vacuum testing apparatuses; and analysis software to validate performance with flight data. An accepted figure of merit used by the industry, the α/ε ratio (solar absorptivity over thermal emissivity), assesses the thermal performance of radiator coatings.[9,10] Absorptivity is the fraction of incident sunlight a surface absorbs rather than reflects, and emissivity is how efficiently it radiates its own heat as infrared. A good radiator coating therefore minimizes the first and maximizes the second, giving a low α/ε ratio that keeps the surface

cool under sunlight while still rejecting heat. Specific mass, the kilograms of thermal system hardware required per kilowatt of heat rejected, is the basis on which competing radiator concepts are conventionally compared.[11]

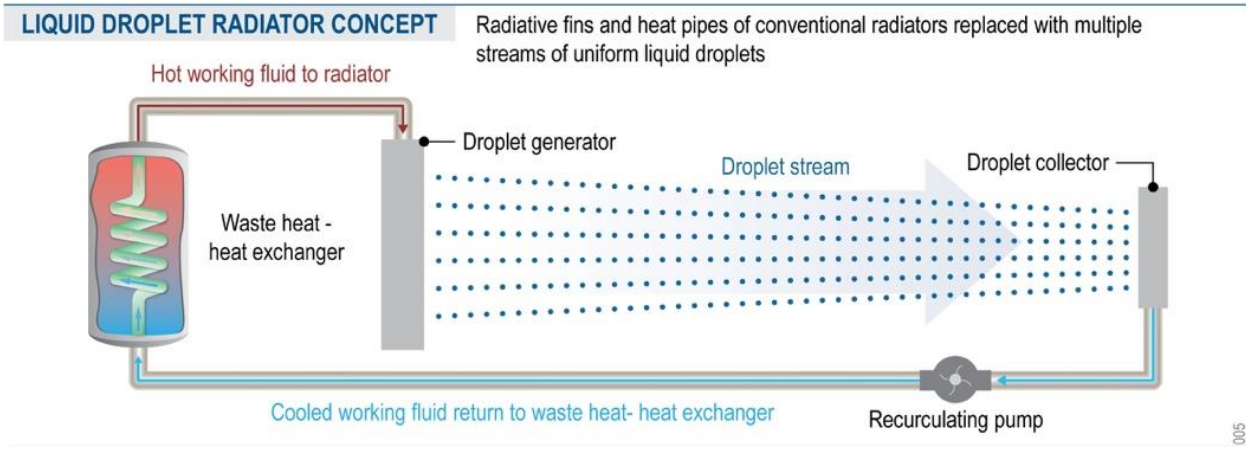
Most of this prior knowledge transfers directly to the data center application, although the problem demands additional insights the spacecraft industry has never needed. In particular, space-relevant thermal strategies approach their limits at larger scales,[12] and the inherited analyses offer no link between thermal performance and compute a system delivers. No analog to power usage effectiveness (PUE), as applied to terrestrial data centers, exists for orbiting data centers. In fact, radiative cooling consumes essentially no power, and so the thermal subsystem contributes negligibly to PUE, rendering the metric insensitive to the quality of the thermal design. Closing that gap will require new measures of thermal performance, reported alongside the figures of merit already inherited from spacecraft engineering.

ii. Thermal system design

Radiators selected for SBDCs are likely to resemble those flown on spacecraft for decades: a flat panel of aluminum or carbon-fiber composite embedded with internal heat pipes and surface-treated to radiate infrared at high efficiency ($\varepsilon \approx 0.85–0.90$) while reflecting most incident sunlight ($\alpha \approx 0.1–0.2$).[10] A conventional spacecraft panel maintained at 20 °C emits ~355 W/m² from one radiator face (calculated from the Stefan-Boltzmann law, $P = \varepsilon\sigma T^4$, with $\varepsilon \approx 0.85$). After accounting for typical orbit-averaged environmental absorption, the net heat rejection can fall to ~250 W/m² (for a panel with Earth in roughly half its field of view). However, for an SBDC, the same panel running at a chip-tolerable temperature of 85 °C sees its emissive power rise to roughly 790 W/m² (both figures follow from the Stefan-Boltzmann law; the call-out box works through a full AI1 example). And a panel at a knife-edge attitude to the sun, kept facing deep space almost continuously, sees little of Earth (less than a tenth of its field of view), and so absorbs little environmental heat (Earth's infrared and reflected sunlight), with its net rejection falling to ~780 W/m². The threefold spread between these two estimates in this example clearly demonstrates that temperature and orientation are what principally determine how a radiator performs.

One advanced concept is the liquid droplet radiator, which dispenses with panels entirely. The working fluid itself is sprayed across open space as a mist of microscopic droplets, which radiate their heat directly to the void before being recollected and recirculated. Because the droplet sheet needs no structure, analyses dating to the 1980s found it can be many times lighter than the lightest panel radiators, with reduced vulnerability to micrometeoroid strikes.[13]

A recent peer-reviewed analysis of a megawatt-class space nuclear system projects the advantage at roughly an order of magnitude, finding the droplet radiator's mass to be about 10% of an equivalent heat pipe radiator's mass.[14] This technology is not examined in depth here: these systems have not been fully tested in real-world scenarios, and the first generation of SBDCs likely will be built on flight-proven panel radiators.

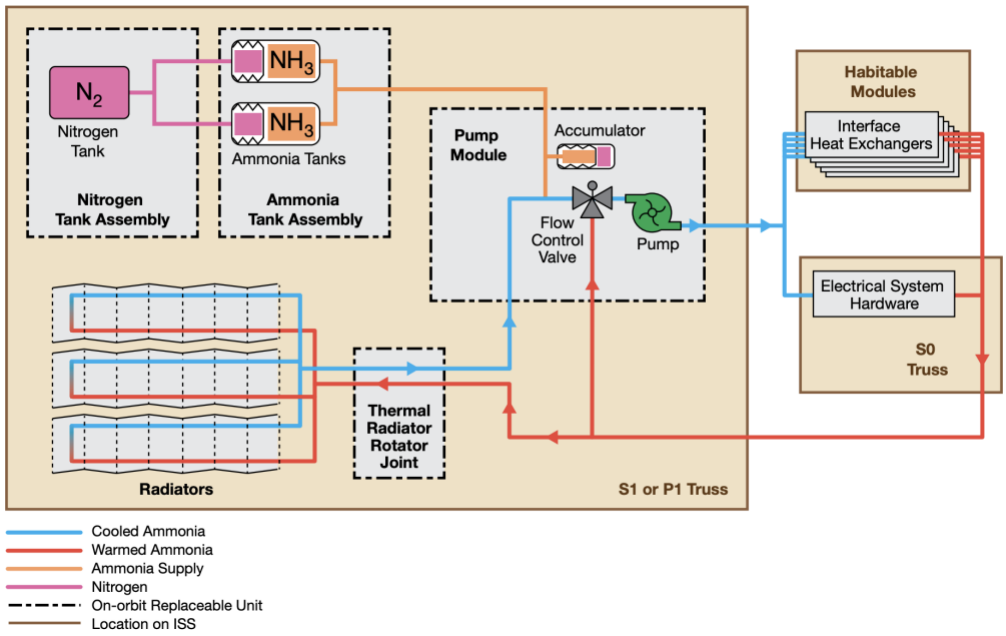


Liquid Droplet Radiator Concept.

Within the system, heat transfers across a series of interfaces from chip to radiator before finally being emitted as infrared. At each interface, thermal transport incurs a penalty: a temperature decrease determined by that interface's thermal resistance. Each temperature drop then sums to the total temperature gap between chip and radiator temperature, and because silicon physics sets the junction temperature of the electronics, each degree lost along the transport chain leaves the radiator that much cooler, carrying the T^4 penalty with it. Consequently, a poorly engineered interface or thermal transport loop reduces rejection capacity, which must then be offset by added area and mass.

The workhorse thermal transport device of spacecraft thermal control is typically the heat pipe: a sealed tube in which working fluid evaporates at the hot end, condenses at the cold end and returns through a capillary wick.[9] Typically, no pump is required—the wick does the work—so a heat pipe consumes no power and contains no moving parts. Further, these devices have flown reliably on spacecraft for decades. However, the capillary force that returns liquid through the wick is weak, which limits both how much heat a heat pipe can carry and how far it can carry it. Once a data center scales from tens to hundreds of kilowatts, pumps become unavoidable and the heat must travel in a mechanically driven fluid loop. For a real-world example, consider the International Space Station (ISS), which uses a ~70-kW thermal control system that relies on mechanically pumped ammonia loops rather than passive wicks.[15] SpaceX's AI1 proposes a similar configuration, with redundant pumps moving fluid across its radiators.[7,8]

These features come at a price since pumps must be designed to run unmaintained for years while being exposed to ionizing radiation that degrades electronics. A pump is a single point of failure, and so redundancy is also essential. Inside the radiator panel, embedded heat pipes take heat from the point where it enters and spread it across the whole radiating surface. This helps keep the panel isothermal, which enables the system to reject heat at full capacity. A solid plate without heat pipes could achieve this spreading by conduction alone, but the radiator would then need to be thick enough to carry the heat laterally without experiencing a large temperature drop, which would add

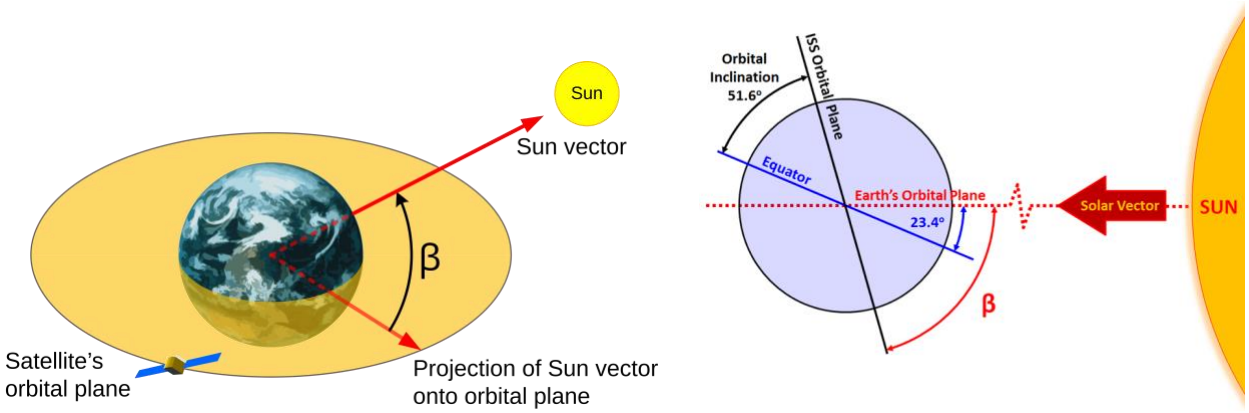


Schematic of the ISS External Thermal Control System (ETCS) with pumped ammonia loops. Source: NASA

weight. Embedded heat pipes accomplish the same spreading without the additional mass, keeping a thinner, lighter weight panel nearly isothermal.

Radiator performance is fundamentally governed by orientation relative to environmental flux vectors. While some SBDC designs (including the SpaceX design described below) align satellites to be “knife-edge” to the Sun, others use different orientations. For example, in a standard nadir-pointing attitude, a satellite keeps one face permanently pointed at Earth (nadir) while the opposite face looks out at deep space (zenith). The angle between the solar vector and the orbital plane, called the beta angle, drifts continuously throughout the year as Earth orbits the sun and the satellite’s orbital plane rotates in space. As the beta angle varies, solar incidence shifts, making a geometry that succeeds at one beta angle insufficient at another.

While a zenith-facing radiator looks out at 3 Kelvin of cold, deep space, deployable double-sided panels inevitably expose one face to Earth’s emission and albedo. Even in a knife-edge attitude, the radiator cannot escape Earth; each face of a deployed panel still sees a large fraction of Earth's disk at grazing angles, intercepting infrared and reflected sunlight. Spacecraft attitude can reduce this Earth-viewing load but, unlike direct sunlight, cannot eliminate it. Optical coatings can be engineered to reflect the sun's short wavelengths that peak in the visible range, but this approach does not work against Earth's infrared since rejecting Earth's long wavelength infrared would mean sacrificing the radiator's own emissive power. (Emissivity and absorptivity at the same wavelength must be equal by Kirchoff’s Law of thermal radiation.) Earth's infrared load must be accounted for in the design budget since no coating or geometry can fully eliminate it. Direct solar loads can be largely engineered away using coatings with lower solar absorptance and through orbit selection, radiator articulation and



The beta angle between the solar vector and satellite orbital plane, which determines how much direct sunlight radiators receive and how much of each orbit is spent in eclipse, both of which affect radiator sizing and cooling efficiency. Source: Wikimedia Commons. CC A-SA 3.0 Unported

The beta angle of the ISS, which cycles between roughly ±75° over the year due to its 51.6° orbital inclination and Earth's 23.4° axial tilt. Source: CERN

spacecraft attitude control. Which of these approaches a design uses, and how completely it limits unwanted heat load, is a critical thermal design decision.

Power generation and thermal management work against each other since solar arrays must face the sun to capture power, while radiators must avoid the sun to stay cold. Designers approach the orientation problem three ways (Figure 2). Fixed radiators are bolted to the body; they have no mechanisms or failure modes, but their performance depends on seasonal geometry, unless the orbit itself eliminates the seasonal dependence. A sun-synchronous terminator orbit keeps the sun's direction nearly constant relative to the orbital plane year-round, allowing even a fixed radiator to maintain edge-on solar geometry permanently while also providing near-continuous sunlight for the solar arrays.[16] In contrast, articulated radiators connect to gimbals that continuously point them away from the sun (the approach taken on the ISS), which ensures more consistent performance but at the price of mass and joints that must survive tens of thousands of tracking cycles over a five-year mission. The SpaceX AI1 illustrates a third path: the radiator is fixed and the satellite is steered instead. Its panels hold a knife-edge attitude, aligning their thin edge toward the sun, the way a blade held edge-on to a light source nearly vanishes, presenting almost no cross-section to sunlight while leaving both of its faces free to radiate to deep space. This is the key advantage over simply shadowing the radiator behind the solar array, which would spare it the sun but leave one radiator surface facing warm hardware instead of cold space. The mechanism involves attitude control of the entire satellite rather than the use of gimbals, which costs almost nothing for a platform carrying few, if any, Earth-pointed instruments. This attitude fixes the panels' orientation relative to the sun only; their alignment toward Earth, and the residual infrared absorbed, is a separate design parameter to consider.

iii. Thermal performance metrics

There is no metric that currently ties thermal performance to delivered compute throughput. The most familiar metric, PUE, collapses in comparison because radiative cooling consumes essentially no power. Others transfer only in part. The thermal resistance from chip-to-coolant, a standard measure of a terrestrial cooling loop, describes only the first link of a chain that in orbit continues through the radiator to deep space. However, for SBDCs, we must consider the total temperature drop across the hardware and into deep space. The cost per kilowatt is still central, but its leading term moves away from real estate and electricity and toward launch mass, which thermal specific mass (a metric used previously in spacecraft thermal design) captures directly.

Useful orbital metrics must capture how well a thermal system converts a fixed heat load into sustained compute and at what cost in mass. These metrics should consider the following:

1. The fraction of rated compute throughput that can be delivered given the thermal constraints of the system.
2. The magnitude of total interface resistance and the resulting temperature drop from the hardware to the radiator.
3. The total mass of the thermal hardware required to deliver the expected compute throughput.

Measures of the first two do not yet exist and still need to be developed; devising them and agreeing on how to report them should be a priority for the field. The third concern does have an established figure of merit, thermal specific mass, inherited from spacecraft engineering. Thermal specific mass (kg/kW) divides the thermal management system's mass by the heat load the system rejects, and it carries over unchanged from the spacecraft literature, in which competing radiator concepts are customarily ranked on exactly this basis.[11] Thermal specific mass also translates directly into cost: multiply it by the cost of putting a kilogram into orbit and the result is the dollar cost of each kilowatt of heat rejection capacity. The ISS benchmark is roughly 100–150 kg/kW when accounting for the radiators, pumps, radiator rotary joint, and ammonia and nitrogen tank assemblies.[15] NASA's Advanced Radiator Concepts program demonstrated reductions in radiator specific mass by a factor of three to four over then-current heat pipe technology,[11] and liquid droplet concepts are projected to cut radiator mass by roughly an order of magnitude.[14]

Note that mass figures of merit appear throughout this roadmap in related forms; the thermal specific mass presented here is specific to heat rejection (for each kilowatt rejected, rather than generated) and should be read in conjunction with these treatments.

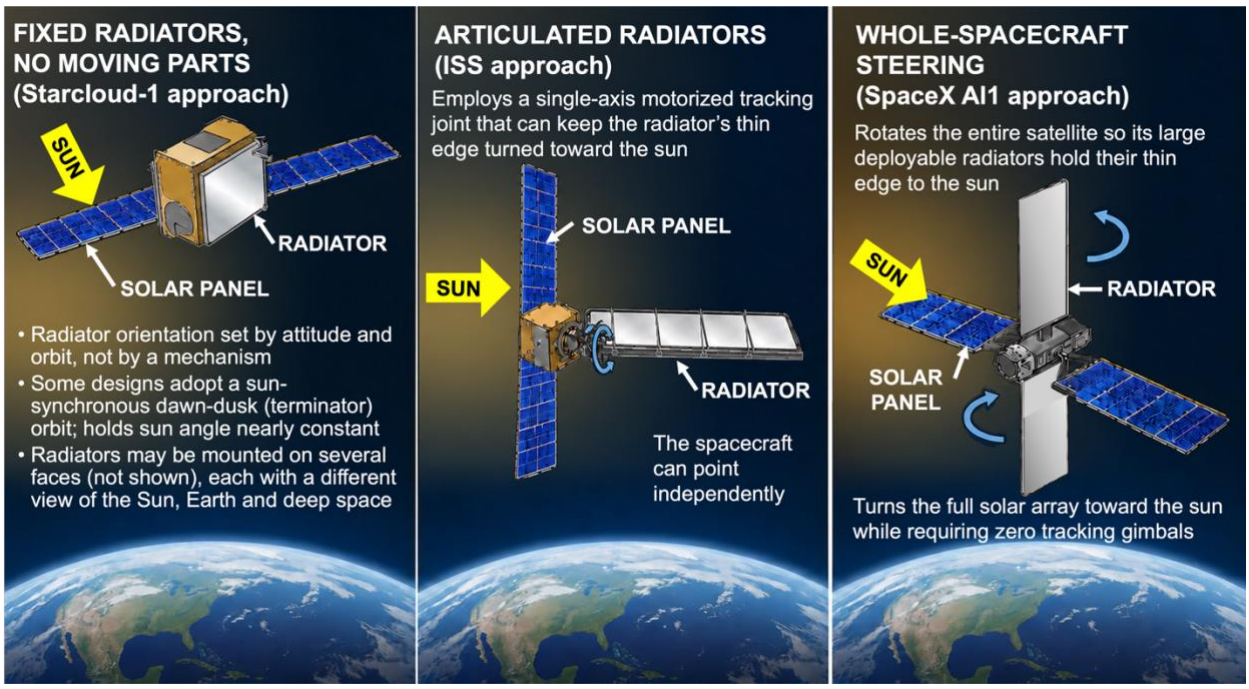


Figure 2. *Comparative thermal orientation architectures for orbital platforms. The panels show three approaches to reconciling the conflicting requirements of solar arrays and radiators. Schematic and not to scale; environmental heat loads, internal coolant loops and residual Earth infrared are omitted. Source: AI generated using Gemini and ChatGPT*

THERMAL PERFORMANCE OF THE SPACEX AI1

The SpaceX AI1, announced in June 2026, has been described publicly in enough detail to support a thermal performance review, although several inputs must be inferred.[7,8] As shown in the table below, the ~1400 W/m² rejection flux SpaceX claims (we calculate 1360 W/m^2) corresponds to a radiator temperature of roughly 76–82 °C, calculated using the Stefan-Boltzmann Law. The range reflects the panel's undisclosed orientation relative to Earth as it travels through its orbit. SpaceX does not disclose the maximum allowable chip temperature, but an industry source reports that Rubin-class GPU (graphics processing unit) die temperatures (the temperature of the silicon itself) can easily reach 95 °C.[17] Assuming a temperature drop of 8 °C across the thermal transport train,[18] a 95 °C junction limit corresponds to a maximum allowable radiator temperature of 87 °C, 5 °C above the top of the calculated range. The 8 °C drop assumed across the transport chain between chip and radiator forfeits roughly 8–9% of the rejection capacity the radiator would otherwise deliver, an illustration of how costly small temperature losses become once the Stefan-Boltzmann law is applied.

Thermal specific mass can only be roughly estimated given that available data are limited. SpaceX has not disclosed the AI1 thermal system's mass, but bounds can be constructed using areal mass, the kilograms of hardware per square meter of radiator panel, and metrics from available ISS data.[15] The ISS provides the only flight benchmark: its active thermal control system rejects up to 70 kW across 6 radiator panels totaling 475 m². Radiator areal mass is about 14 kg/m²,[15] and total radiator mass is roughly 6700 kg. In contrast, modern composite panels exhibit areal mass of roughly 5–10 kg/m²,[3] and next-generation designs are targeting even lower.[19] Including the rest of the rejection system's documented hardware—pump modules, thermal radiator rotary joints, and ammonia and nitrogen tank assemblies—brings the full system to about 9700 kg, yielding a radiator-only specific mass of roughly 100 kg/kW and a full thermal specific mass of ~140 kg/kW. Applying the ISS's 14-kg/m² areal mass to AI1's disclosed 110 m² radiator area gives an estimated radiator mass of roughly 1500 kg; scaling by the ISS's ancillary-to-radiator mass ratio (a further ~44%, covering pumps, rotary joints and tankage) yields a full-system estimate of 2200 kg, or a thermal specific mass of ~10 kg/kW (radiator-only) and ~15 kg/kW (full system) against the 150-kW load. On a full-system basis, this makes the AI1 estimate nearly 10x better than the ISS benchmark. This comparison assumes AI1's ancillary hardware scales with radiator mass the same way it does in the ISS's—a real assumption, not a documented fact, since AI1's pumping and tankage architecture is unknown. Whether AI1's thermal system actually achieves mass in this range and whether the claimed rejection density is realizable in practice both remain to be determined.

Using the ISS estimate of the full thermal specific mass of 140 kg/kW, the thermal system of the Reference SBDC (see Text Box 1-1) would have a total mass of 14 tonnes. By contrast, using the estimate of the claimed SpaceX AI1 full thermal specific mass (15 kg/kW) implies that the thermal system of the Reference SBDC would have a total mass of 1.5 tonnes.

As the table's status column shows, every input is disclosed, claimed, assumed, estimated or derived. Nonetheless, the inputs to these calculations will sharpen and the achievable performance of orbital thermal systems will come into clearer view as systems are flown, operators publish data and independent analyses mature. In any case, the exercise demonstrates that a workable assessment of thermal performance can be assembled from a modest set of design parameters and that such an assessment offers a basis for comparing one SBDC against another.

Parameter	Value	Status
Radiator area	110 m^2	Disclosed (Federal Communications Commission (FCC) filing)
Orientation	Knife-edge to sun	Disclosed
Heat rejection density	~1400 W/m^2	Claimed (SpaceX announcement); analytically verified to order of magnitude
Peak compute/Thermal Design Power (TDP)	150 kW	Disclosed
Coating emissivity (epsilon)	0.85	Assumed, not disclosed
Maximum chip temperature	~95 °C	Assumed using industry source,[17] not disclosed
Radiator operating temperature	8 °C below maximum chip temperature	Assumed,[20] not disclosed
Thermal specific mass	~10 kg/kW (radiator only) ~15 kg/kW (full system)	Estimated from documented panel mass-per-area benchmarks[3,15]

C. Risks and Open Questions

The scaling gap. The largest systems actually flown today are single-GPU demonstrators on the order of a kilowatt. The largest disclosed SBDC design, SpaceX's AI1, is planning for 150 kW.[7,8] Bridging this gap in power requires redesign, not just extrapolation from existing designs: because heat must travel farther, structural support for radiators becomes disproportionately heavy, pumping power becomes non-negligible, and large flexible panels affect attitude control. Ambitions for aggregate SBDCs run to megawatts (or even gigawatts), but this will be achieved chiefly through constellations of many such satellites. Starcloud-2's deployable radiator, launching late 2026, will be the first significant test of scaling at this individual-satellite level.[21]

Degradation. Coatings degrade under exposure to ultraviolet, atomic oxygen and particle radiation. This means that solar absorptivity of radiators typically climbs and capacity erodes. Current estimates suggest that maintaining performance over a five year lifetime may require ~40% more radiator area than a fresh coating would require.[3] But these estimates are based on studies of conventional spacecraft, not from SBDCs, since no coating has yet built up flight history under the high operating temperatures and continuous full-power duty cycles this application demands. Any beginning-of-life evaluation of thermal performance is therefore an upper bound, and real performance will decline over the course of the mission.

Optical and infrared brightness. A radiator's low solar absorptivity, needed to stay cool, is also high reflectivity, so the best coating for heat rejection is close to the worst for brightness. Darkening a radiator to dim its reflectivity raises its absorptivity and forces added area and mass. Holding the radiator edge-on to the Sun helps, reducing absorbed solar load and the sunlit area seen from the ground. Specifically, in the dawn-dusk terminator orbits these systems favor, the geometry directs much of the specularly reflected sunlight away from Earth's surface rather than down toward ground observers.

Debris and congestion. Radiator panels are among SBDCs' most vulnerable systems. Radiators, like solar arrays, have large, thin profiles, and a single debris strike can puncture embedded heat pipes or coolant channels. If the SBDC has the capability to maneuver to avoid the debris, doing so breaks its attitude such that it may no longer be edge-on to the sun. Until its orientation is restored, sunlight may impinge on the radiators, leading to a drop in heat rejection capacity, and the processors must throttle. A beginning-of-life thermal analysis would not capture this effect since the frequency of necessary attitude adjustments cannot be easily predicted. With growing orbital congestion, actual thermal performance likely will fall below the predicted value.

End-of-life disposal. Retiring a large aluminum radiator has no good solution. On re-entry, the panel ablates into nanoparticles of aluminum oxide that reach the stratosphere and catalyze the destruction of ozone.[22] The alternatives to re-entry each fail in their own way. Moving retired SBDCs to higher disposal orbits only defers the problem. Sending them entirely outside of Earth's orbit would eliminate the risk but would require more propellant than the satellites could realistically carry. Each option carries a different cost, and there is not yet an established methodology for comparing them.

The standards vacuum. No standards body has yet taken up comprehensive thermal metrics for this application. It is possible that the first generation of commercial SBDCs will be financed, licensed and sold on the strength of thermal claims that no outside party is able to check. What such claims conceal is thermal throttling: when environmental loads drive a radiator past its rejection capacity, the processors must slow until the load eases, so a platform advertised at one compute capability may deliver materially less of it in practice. Measures of how much of its rated compute a platform sustains over a year in orbit and of how much thermal headroom the chip loses in the transport chain, reported alongside the thermal specific mass at which both are achieved, would render competing systems comparable. As the SpaceX AI1 exercise suggests, assessments of this sort could be assembled today from a modest set of design parameters that operators already possess.

D. Recommendations

1. Developers of SBDCs should treat thermal management as a primary design driver from the outset—sizing radiators and selecting operating temperatures before fixing compute density—rather than adding cooling as a downstream constraint.

2. Operators should design electronics to run at the highest junction temperature their reliability limits safely allow since higher radiator temperatures reject the same heat from substantially less surface area and mass. Because operating near the thermal limit accelerates wear and shortens component life, this trade should be made deliberately, with lifetime specifications and redundancy considered.

3. Developers and operators should disclose the key thermal parameters of their systems, including radiator area, heat-rejection density, coating emissivity and absorptivity, operating temperatures and thermal hardware mass. Given the linkage between thermal management design and both SBDC lifetime and compute throughput, investors and regulators have a stake in thermal performance claims they cannot independently verify.

4. The community should develop shared, quantitative cooling metrics and standards that cover sustained thermal performance over a full year in orbit, the temperature penalty imposed in transit from chip to radiating surface, and the thermal hardware mass required to achieve both. This will allow designs to be compared consistently with one another and over a system's operating life.

5. Funders and research institutions should prioritize development of high-emissivity coatings that resist ultraviolet and atomic-oxygen degradation, so that end-of-life thermal performance, not just beginning-of-life performance, meets mission requirements.

6. Funders and industry should invest in lightweight, reliably deployable radiator structures and in high-temperature-tolerant electronics, as advances in either directly relax the central mass and area constraints on orbital compute.

7. Operators and orbital-traffic authorities should account for the coupling between congestion and cooling. As orbital bands grow more crowded, more frequent debris-avoidance

maneuvers could break radiators' sun-avoiding attitude and force throttling. The odds of a strike puncturing a thin radiator also rise. Both degrade thermal performance in ways a beginning-of-life analysis does not capture.

E. References

1 Sohan Salahuddin Mugdho, Md. Shahedul Hasan & Cheng Wang. Space-CIM: Enabling Compute-In-Memory Accelerators for Thermally-Constrained Space Platforms. arXiv:2606.05741 (2026). https://doi.org/10.48550/arXiv.2606.05741.

2 Satnews. The “Physics Wall”: Orbiting Data Centers Face a Massive Cooling Challenge; Satnews, Sonoma, California, https://satnews.com/2026/03/17/the-physics-wall-orbiting-data-centers-face-a-massive-cooling-challenge/ (2026).

3 Andrew Cavalier. Why Orbital Data Centers Are Harder Than Silicon Valley Thinks; EEE Spectrum, New York, New York, https://spectrum.ieee.org/orbital-data-centers-heat (2026).

4 Tony Pan. The hype is real for space-based data centres. So are the challenges; World Economic Forum (WEF), Geneva, Switzerland, https://www.weforum.org/stories/artificial-intelligence/space-data-centres-cooling/ (2026).

5 Tereza Pultarova. Nvidia Sends a Powerful GPU to Space; IEEE Spectrum, New York, New York, https://spectrum.ieee.org/nvidia-h100-space (2025).

6 Conor McNevin. Starcloud raises US$ 170 million for space-based data centers; w.media, Suntec Tower 1, Singapore, https://w.media/starcloud-raises-us-170-million-for-space-based-data-centers/ (2026).

7 Sebastian Moss. SpaceX details AI1 satellite 'data center,' claims 150kW peak compute; Data Centre Dynamics Ltd. (DCD), London, United Kingdom, https://www.datacenterdynamics.com/en/news/spacex-details-ai1-satellite-data-center-claims-150kw-peak-compute/ (2026).

8 Space Exploration Technologies Corporation (SpaceX). SpaceX - Response to Federal Communications Commission (FCC) (05-29-2026) Re: ICFS File No. SAT-LOA-20260108-00016; Scribd, Inc., San Francisco, California, https://www.scribd.com/document/1045612919/SpaceX-Response-05-29-2026 (2026).

9 Robert C. Consolo & Sandra K. S. Boetcher. Chapter One - Advances in spacecraft thermal control in Advances in Heat Transfer Vol. 56(eds John Patrick Abraham, John M. Gorman, & W. J. Minkowycz) 1-50 (Elsevier, 2023, https://doi.org/10.1016/bs.aiht.2023.04.001).

10 David G Gilmore & Martin Donabedian. Spacecraft thermal control handbook. Vol. 1 (Aerospace Press, El Segundo, CA, 2002, https://www.academia.edu/934756/Thermal_Control_Handbook).

11 Albert J. Juhasz & George P. Peterson. Review of advanced radiator technologies for spacecraft power systems and space thermal control (NASA Technical Memorandum 4555); National Aeronautics and Space Administration (NASA) Technical Reports Server (NTRS), Hampton, Virginia, https://ntrs.nasa.gov/citations/19940032314 (1994).

12 Jiahao Shi, Xudong Zhang, Min Yang & Jing Liu. Thermal Management Technologies for Space Data Centers: Current Status and Prospects. Journal of Refrigeration 47, 1-19 (2026). https://www.sciopen.com/article/10.12465/issn.0253-4339.20251030004.

13 A. T. Mattick & A. Hertzberg. The liquid droplet radiator—An ultralightweight heat rejection system for efficient energy conversion in space. Acta Astronautica 9, 165-172 (1982). https://doi.org/10.1016/0094-5765(82)90084-4.

14 Hao Qin, Chenglong Wang, Wenxi Tian, Suizheng Qiu & Guanghui Su. Energy allocation optimization of the gas-cooled space nuclear reactor. Applied Thermal Engineering 196, 117289 (2021). https://doi.org/10.1016/j.applthermaleng.2021.117289.

15 Boeing IDS Business Support. Active Thermal Control System (ATCS) Overview; National Aeronautics and Space Administration (NASA), Washington, D.C., https://www.nasa.gov/wp-content/uploads/2021/02/473486main_iss_atcs_overview.pdf (2021).

16 Travis Beals. Exploring a space-based, scalable AI infrastructure system design; Google Research, Mountain View, California, https://research.google/blog/exploring-a-space-based-scalable-ai-infrastructure-system-design/ (2025).

17 Anton Shilov. Frore claims its LiquidJet can drop Nvidia Rubin GPU temperatures by 10°C — can also boost performance by 15% as hyperscalers eye using delidded GPUs in production environments; Tom's Hardware (Future US Inc.), New York, New York, https://www.tomshardware.com/pc-components/liquid-cooling/frore-claims-its-liquidjet-can-drop-nvidia-rubin-gpu-temperatures-by-10-c-can-also-boost-performance-by-15-percent-as-hyperscalers-eye-using-delidded-gpus-in-production-environments (2026).

18 Graeme Smecher, Peter Cameron, Jean-François Cliche, Matt Dobbs & Joshua Montgomery. “Development of TRL5 space qualified hardware for tuning, biasing, and readout of kilopixel TES bolometer arrays” in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, (eds Jonas Zmuidzinas & Jian-Rong Gao) 121900X, https://doi.org/10.48550/arXiv.2207.11374 (2022).

19 Mohamed S. El-Genk. Lightweight Heat Pipe Radiator Panel for Lunar Surface Fission Power Systems in AIAA AVIATION FORUM AND ASCEND 2024 (https://arc.aiaa.org/doi/abs/10.2514/6.2024-4937).

20 Pradeep Bhandari, Gajanana C. Birur, Gani Ganapathi, Anthony D. Paris, Mike Pauken, Keith Novak & Glenn Tsuyuki. “Mechanically Pumped Fluid Loops for Spacecraft Thermal Control: Past, Present & FuturePast, Present & Future” in 15th Annual Thermal & Fluid Analysis Workshop, Pasadena, California. https://tfaws.nasa.gov/TFAWS04/Website/program/Speakers/TFAWS2004PresentationonMechLoops-v1.pdf,(2004).

21 Blake Crosley. The Orbital Data Center Race: Every Major Player, Timeline, and Economic Reality in 2026; Introl Solutions, LLC., Chicago, Illinois, https://introl.com/blog/orbital-data-centers-space-computing-race-2026 (2026).

22 José P. Ferreira, Ziyu Huang, Ken-ichi Nomura & Joseph Wang. Potential Ozone Depletion From Satellite Demise During Atmospheric Reentry in the Era of Mega-Constellations. Geophysical Research Letters 51, e2024GL109280 (2024). https://doi.org/10.1029/2024GL109280.