Tech Brief:
Orbital Data Centers — Radiators, Watts, & Orbits

Paul Struhsaker
Arrasar Partners
www.struhsaker.com/arrasar-partners

October 27, 2026 9:45 am – 10:00 am

Orbital Data Centers vs Terrestrial Data Centers

Factor	Space-Based Data Centers	Terrestrial Data Centers
Energy Availability	Continuous, uninterrupted solar power (polar orbit) gigawatt-scale potential ($0.01 per kWh)	Dependent on local grids; rising energy constraints ($0.05 to $0.15 per kWh)
Cooling Efficiency	Radiant cooling in vacuum of space - efficient heat dissipation (weight	Requires massive water and/or HVAC systems; costly
Latency	High latency for Earth-bound users; good for orbital edge computing	Low latency; close to users and fiber networks
Scalability	Vast physical space; no land constraints	Limited by land, zoning, and environmental impact. Ocean/Submerged data centers as an alternative
Deployment Cost	**Highly sensitive to Launch Costs**, early-stage technology with resulting R&D costs	Mature, cost-optimized industry
Environmental Impact	Reduces Earth-based energy/water use BUT, launch emissions and **end-of-life disposal issues** must be dealt with	High energy and water consumption; noise and land use issues
Silicon - GPU/CPU	Early reuse of existing silicon (non-rad-hardened) industry maturity will require rad hardening and power savings	Commercial CPU/GPU optimized for terrestrial deployment – not optimized for power savings
POLITICAL RESISTANCE	Currently, there is little political resistance; however, pollution from end-of-life disposal could be an issue	Significant Political Resistance to new data center construction

High-Density Modular "Compute Tiles" Space-based data center architecture [Thales Alenia Space]

Key Stake Holder Constellation Plans

Polar Dusk-Dawn Line Orbit [Starcloud]

Company/ Constellation	FCC Filing	Number of Satellites	Planned Orbital Planes
SpaceX	X	1,000,000	500 – 2,000 km Polar & 30 degree
Starcloud	X	88,0000	600-850 km Polar (multiple orbital planes)
Blue Origin /Amazon "Project Sunrise"	X	51,600	500 and 1,800 km Polar
Orbital Compute	X	100,000	500 km to 850 km (100kW/sat – 10gW/constellation)
Space Cowboy "Stampede"	X	20,000	700 km to 1,000 km
Google "Project Suncatcher" Launched October 1	N/A	N/A	Orbital AI edge commuting tests of Google TPU AI ASIC
Technology Dev. Companies: Sophia Space, Axium Space, Space Cowboy ... and more coming	N/A	N/A	TBD

Space-based data centers filings are 2 orders of magnitude greater (100x) than current Mega- communication constellations

Google “Project Suncatcher”

- **Project Suncatcher Launched October 1 n SpaceX Transporter-18 rideshare mission**
- **Goals Of Project Suncatcher**
  - Achieving data center-scale inter-satellite links (800 MBPS bidirectional using spatial and wave division multiplexing)
  - Controlling large, tightly-clustered satellite formations (80+ satellites in 1Km radius cluster
  - Radiation tolerance of TPUs (and memory) – mission dose at 750 rads tolerance is 2000 rads
  - Economic feasibility and launch costs -at $200/kg by the mid-2030s reach cost parity with terrestrial data center costs
- **Suncatcher Hardware configuration of**
  - **4 Google Trillium Tensor Processing Units (TPUs) -**equivalent to a single **Google Cloud TPU v6e-4 slice**
  - **Packaged to withstand 50-100 g launch vibrations**
  - **1 kilowatt (kW) of solar array - 15-minute operation, then cool-down cycling**
- **Orbital status:**
  - Basic radio contact has been established to track chip durability, thermal behavior, and radiation damage over a planned **one-year testing window**

Suncatcher Prototype at Planet facility

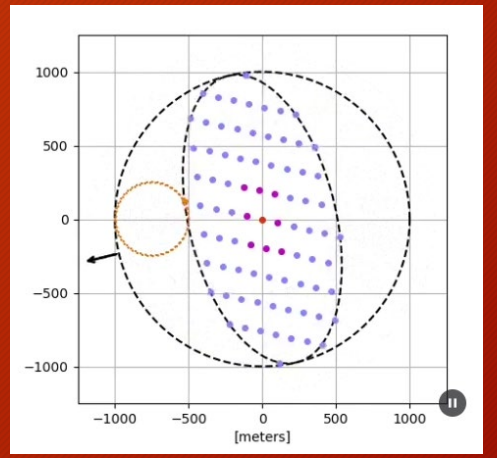

Sun-synchronous Free-Fall (“no thrust”) Orbit Cluster

AI Data Center Compute Life Cycle

- AI Data Center Blade Lifespan ~ **5 years**
- Performance Obsolescence (The "Nvidia" Effect)
  - Each AI chip generation offers 4x to 5x improvement in processing power
  - A 3-year-old AI chip can be made functionally obsolete (e.g., Nvidia H100 to GB200/Blackwell)

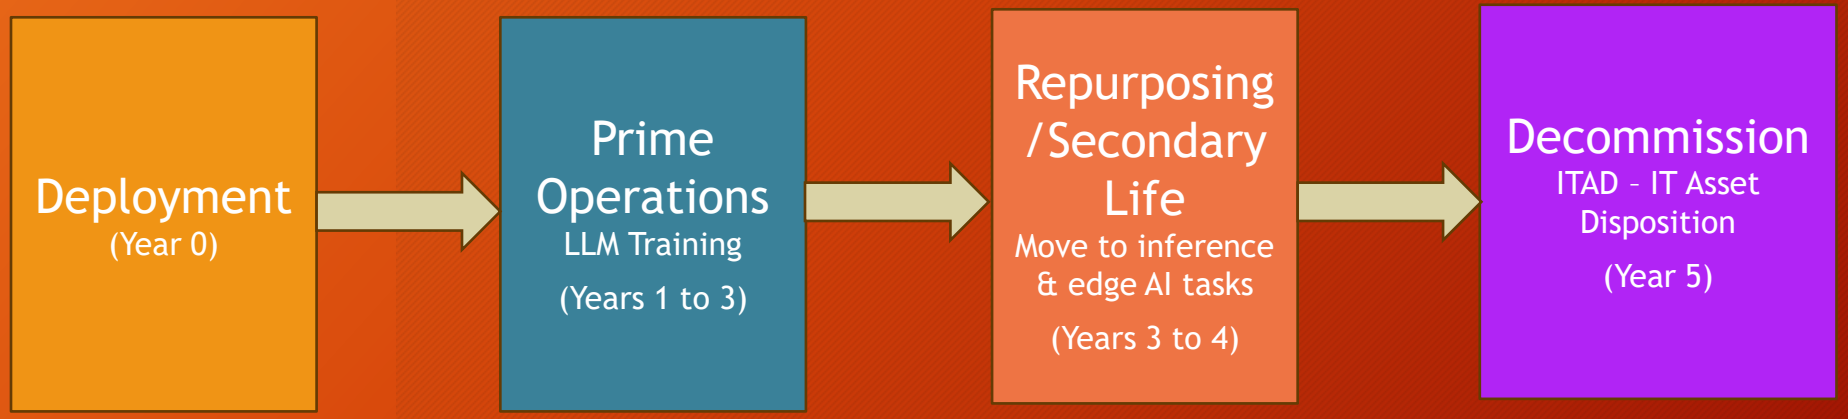

Compute Lifespan and Replacement is a Significant Logistics Issue
What about decommissioning? *Deorbiting pollution or Parking Orbit Junk*

How do changes in Moore's Law change this cycle?

Space Data Center Design

- Three factors drive Space Data Center Design
  - **AI ASIC power consumption:** 1.5kw moving to 3kw per ASIC
  - **Solar Panel Area:** 400 to 600 W/m^2 acts as a sun shield for the radiator
  - **Radiator:** (sunshield @ 27 deg C) 450 W/m^2 one sided to 900 W/m^2 two sided – ammonia fused for heat transfer
- Space-Based Constellation Architectures
  - Distributed "Compute Constellations" (Mesh Networks)
    - Clusters of modular laser-linked AI satellites act as an orbiting node
    - Multiple nodes on multiple orbiting planes make the constellation
    - Proposed by: SpaceX, Google Suncatcher, & Starcloud (Stem & Leaf)
  - High-Density Modular "Compute Tiles"
    - Large-scale satellite node composed of integrated AI ASIC(s), Solar, cells and radiator tiles
    - Multiple nodes on multiple orbiting planes make the constellation
    - Proposed by Sophia Space, and Axiom Space
- Hybrid Earth-Space Architecture for the entire system
  - Edge AI in LEO, non-latency-sensitive workloads in SPACE, and latency-sensitive tasks in EARTH-based data centers
- Sun-synchronous orbits (polar dusk-dawn) from 500km to 2000km
  - Majority of constellations will be above 700km for lower drag, decreased fuel budget, and long orbital life (25+ years at 700km to >1000 years at 2000km)

Distributed Modular Compute Constellation

High density Compute Tiles Architecture

Orbital Data Centers — Radiators, Watts, & Orbits - P Struhsaker, Arrasar Partners

AI ASIC Power Consumption: The Driver of All Space-Base Data Center Design

- Horizontal & future wafer scale integration are driving 40% increase in AI CPU ASIC Platform power per year – **Focus is on Performance NOT Power**
  - Power is increasing despite Moore's Law improvements in silicon process nodes
- Each kW of power drives 2.5 m^2 solar and 2.25 m^2 single sided radiator
- Note: Intel, AMD, SpaceX/Tesla, Google, & Amazon have significant multi-generation AI ASIC programs

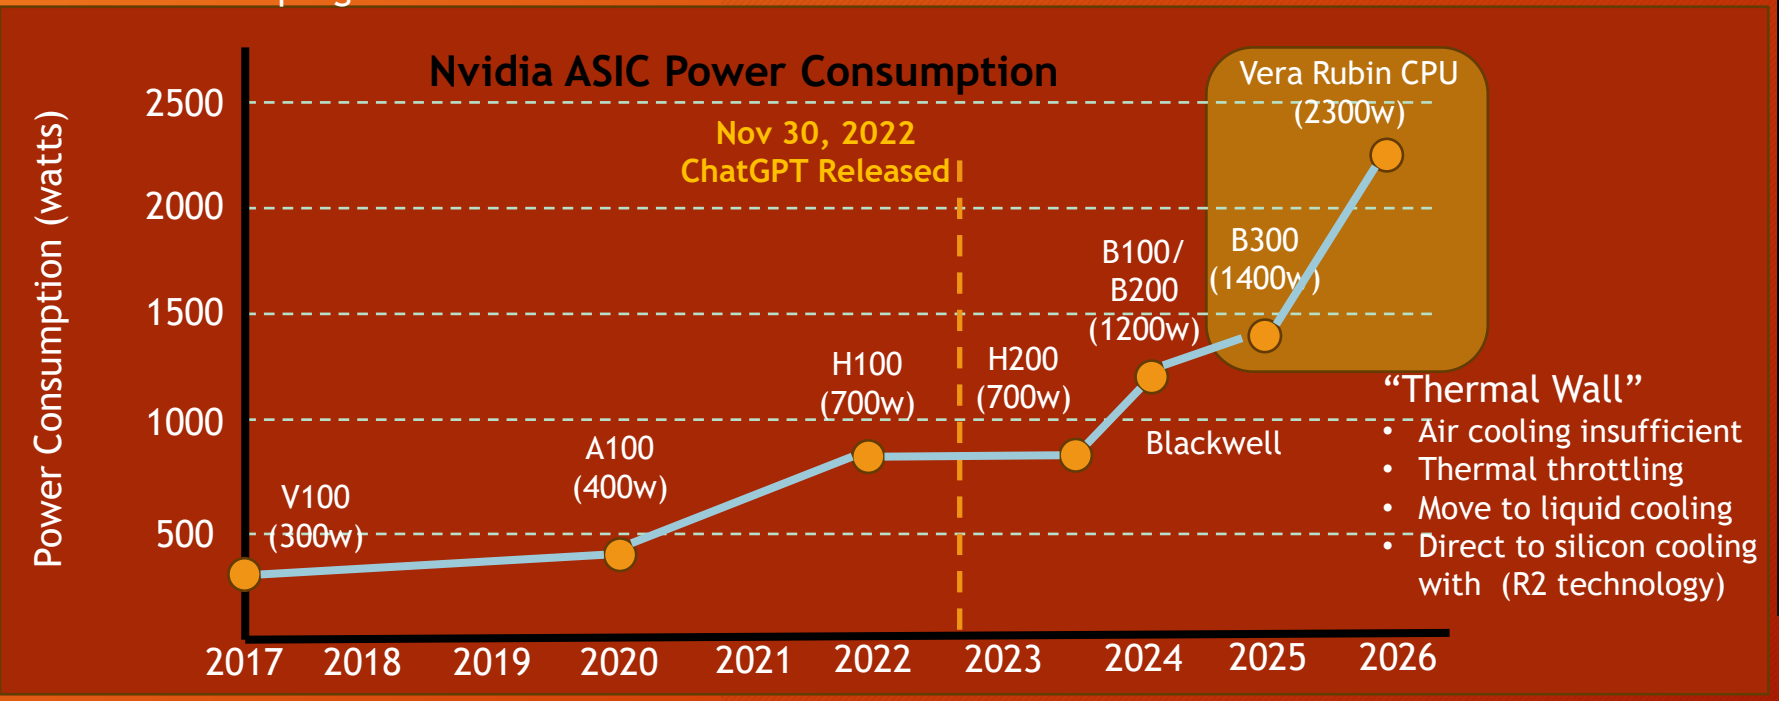

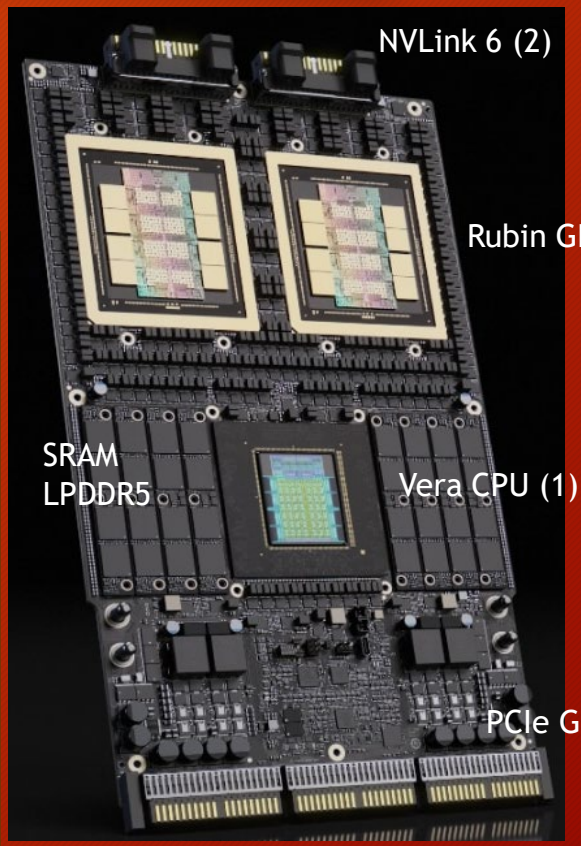

Vera Rubin Super Chip Platform based on TSMC's 3nm (N3P) [Nvidia]

Space-Based Radiators

- Cooling in space is based on infrared radiation
- Solar cells are used as a sun block, providing shade for the radiators (-270 deg C)
- Radiators use Passive or Active heat pipes to distribute heat to the Radiators using Aluminum-ammonia or copper-water working solutions
  - **Passive Heat Pipes** are based on a capillary wick structure to operate in zero gravity
    - **Oscillating Heat Pipes (OHPs):** passive radiators improve weight and performance
    - **Advantage:** Reliable, long life, no active parts, minimum weight, no power load
    - **Disadvantage**: Designed for a fixed thermal load (not variable)
  - **Active Heat Pipes** are based on pumps to move the working fluid
    - **Loop Heat Pipes (LHPs): Capable of transferring higher heat loads over longer distances with high efficiency**
    - **Advantage:** Allows for variable load, more efficient, higher thermal load
    - **Disadvantage**: complex with moving parts, lower reliability, energy use
  - Modern/Lightweight (Carbon-Carbon) radiators 2kg per m^2

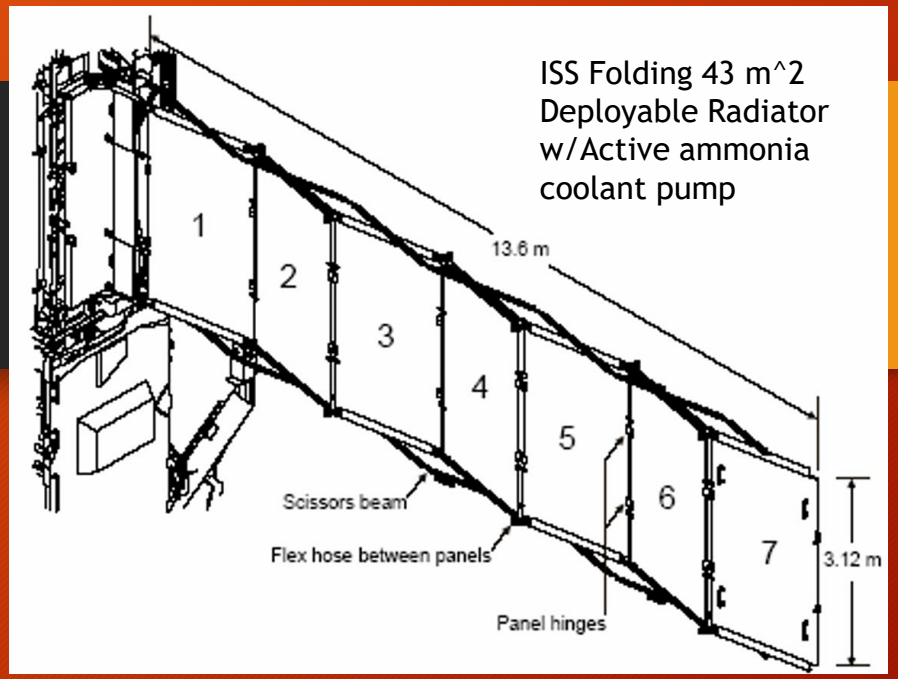

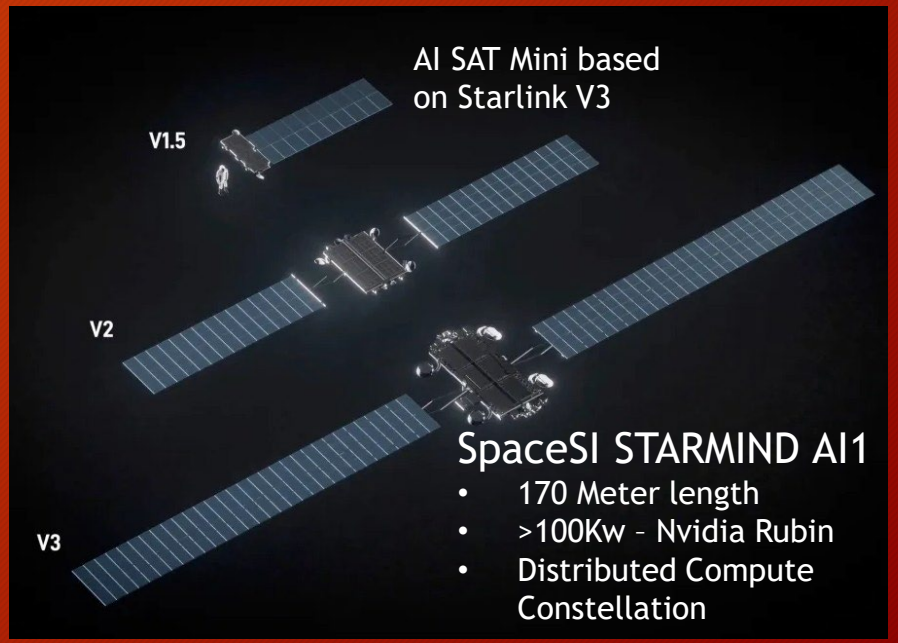

Orbital Life Cycle:
A case for modular design in-orbit servicing

- Majority of space data center satellites will be in orbits between 700 to 2000 Km
  - 25 years till 1000s of years before orbit decay
- Solar Cells, Radiators, have 15 to 25-year lifespan
- AI CPU/Blade, Free Space optics,& station keeping propellant have short lifespans Short Lifespan components (~5 Years):
- Clearly a modular design using solar cells/radiators as a backbone skeleton could cut down decommissioning waste
  - A 5-year replacement cycle using am in -orbit replaceable server farm module
  - The old module would be decommissioned (deorbit or parking orbit) and replaced by a new module
  - The replacement module would upgrade CPU/Free Space Optics recharge propellant and radiator working fluid
- A modular replaceable architecture will reduce the mass to orbit life cycle costs

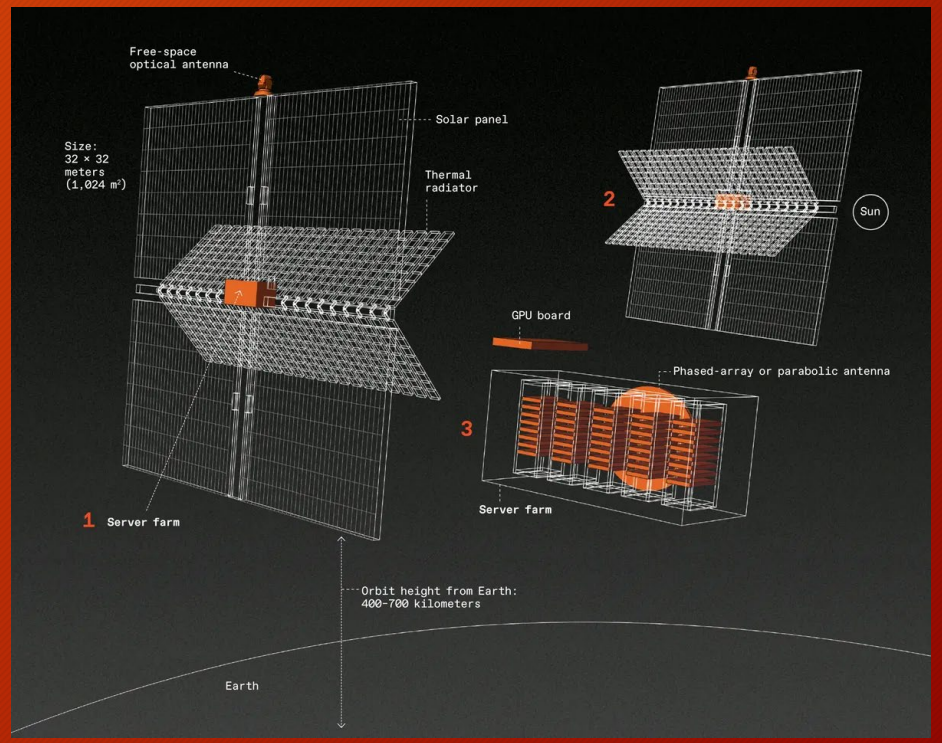

Large-scale Data Center Satellite with Replaceable Server Farm module (IEEE)

The Advantages of Vertical Integration

- Launch costs of <$300/kg are required for an economically viable space data center deployment
  - In the near-term, SpaceX Starship and the next-generation Blue Origin are the only launch vehicles that can meet launch cost goals: an 8x reduction from current launch costs
- SpaceX and Amazon/Blue Origin can leverage their “full stack” capabilities to deliver the vision the space-based data center vision:
  - Launch vehicles and launch cadence build out their respective constellations
  - Exiting satellite communications and ground station network Starlink & Amazon LEO
  - Terrestrial Data Centers and AI business and partnerships
  - Mature, well-funded custom AI chipset development programs
- Despite these leveraged capabilities, large opportunities exist in the ecosystem and niches:
  - **Chipsets and Foundries:** TSMC, Intel, Nvidia, AMD, etc.
  - **Earth Observation edge compute AI:** Nvidia Jetson-Orin & Thor, & alternatives
  - **Full-stack supply chain:** optical terminals, thrusters, etc.
  - **Alternate launch services:** Rocket Labs, Firefly, etc.

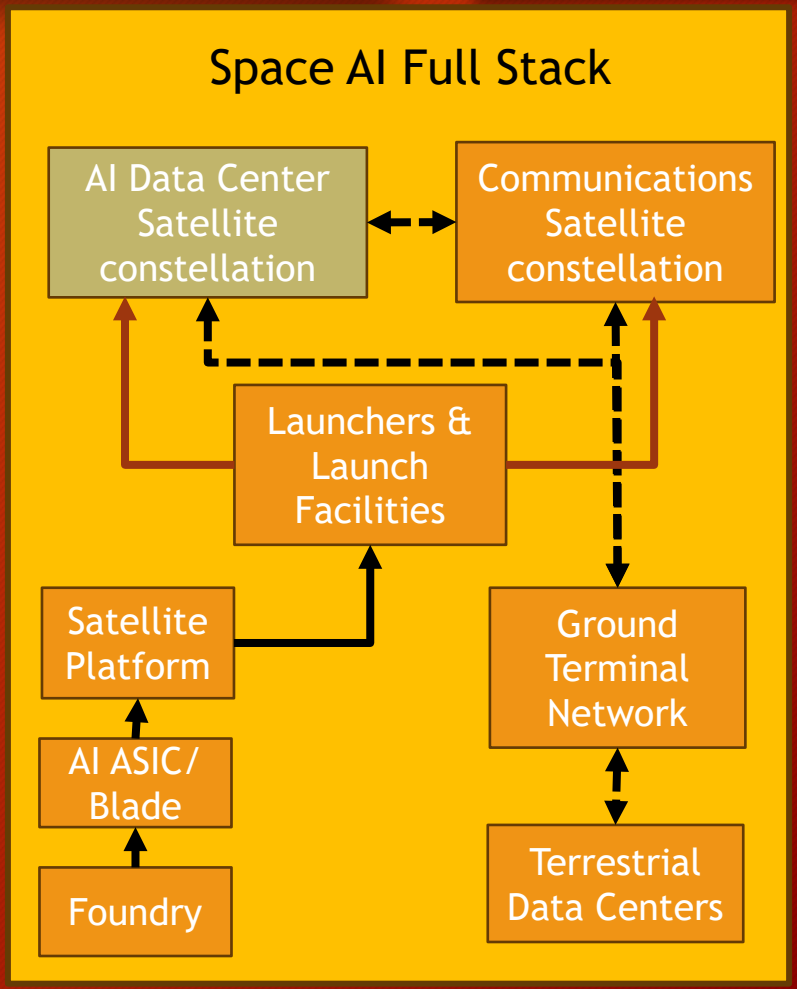

Custom Space-Based AI/CPU

- Early Space Applications are testing existing Core and Edge AI ASICs
  - Design for Performance, Not Power Savings
  - Physical shielding and error detection and correction SW
  - Radiation Effects are only partially mitigated
- Custom AI ASICs will be needed for scale, cost, and longevity
  - Advanced fabrication technology (for example R2)
    - R2 silicon interposer reduces power consumption by 50%
    - R2 tech licensed by Intel & Samsung
  - System-wide Error Correcting Code (bus & memory)
  - Fault Detection and Redundant Hardware (processors & memory)

Radiation Effects on 2n m ASICs

- **Single Event Effects (SEE):** single high-energy particle note strike
- **Multi-Bit Upset (MBU):**single strike with multiple failures and ECC failure
- **Single Event Latch-up (SEL):** Hardware failure/destruction
- **Total Ionizing Dose (TID):** cumulative damage from long-term exposure – causes power leakage

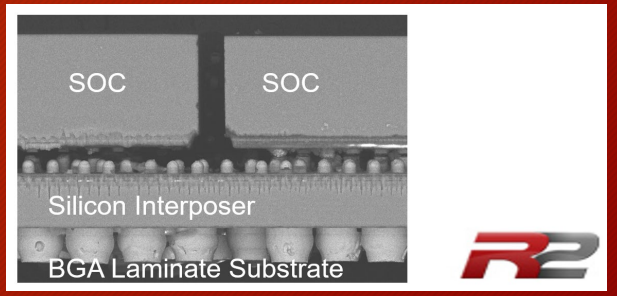

SpaceX Terrafab: The Ultimate in Vertical Integration

- SpaceXSI will use a combination of off-the-shelf AI GPUs while transitioning to SpaceXSI custom silicon over time:
  - IBD reports $40B raise to purchase next gen Nvidia Rubin/Vera chipsets for Colossus and other Data centers
  - Value Add VC reports Starmind AI1test satellite will use Vera Rubin NVL72 Space hardened rack that integrates 72 Rubin GPUs, 36 Vera CPUs, ConnectX-9 SuperNICs and BlueField-4 DPUs
- Longer term, SpaceX has committed over $130M+ to the Terrafab Project and its own AI chips for cars, humanoids, & SpaceXSI
  - Reuters reports Terrafab will use Intel A14 (1.4nm) process technology
    - A14 process uses R2 technology for possible doubling of compute power for a given AI satellite Power configuration and 3D chip stacking
  - SpaceX is also in talks with TSMC for a Terrafab
  - The final outcome is TBD

Artist Rendering SpaceXSI TerraFab Grimes County TX (SpaceX)

The Economics of Data Centers

Achieving cost parity and offering bespoke AI services are the path to profits

Can Orbital Data Centers be cost-effective?
An initial view from the IEEE

- Data Centers orbital or terrestrial, are being designed around gigawatt (Gw) capacity (power and cooling dominate)
  - Terrestrial 1GW are in operation today, with plans of up to 100GW facilities
  - Example: SpaceX SI Colossus 2 at 1 GW, expanding to 1.5 to 2 GW
- Initial estimates put cost at 7x to 10x terrestrial costs
  - Launch cost reduction to <$300/Kg were not factored in
- Glenn Zorpette at IEEE generated a revised model in Feb. 2026
  - 4,300-satellite constellation with 1GW capacity
  - 1,024-square-meter 250 KW solar array (per satellite)
  - 175 NVIDIA Vera Rubin GPUs + companion Vera CPUs and associated switches and interconnect (based on racks like Spacemind AI/NVidia)
  - **Revised Cost Premium: 3.375x**

How do we Close the gap?

Watt-Efficient Rad-Hardened GPUs & low launch costs

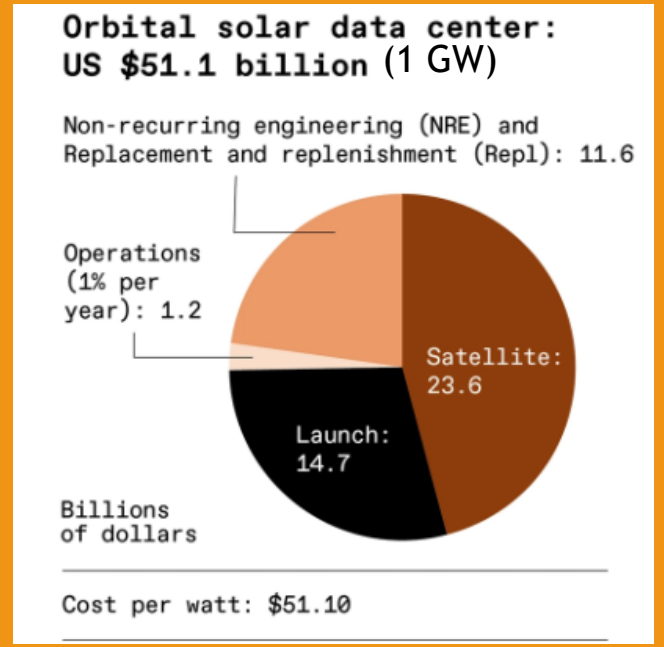

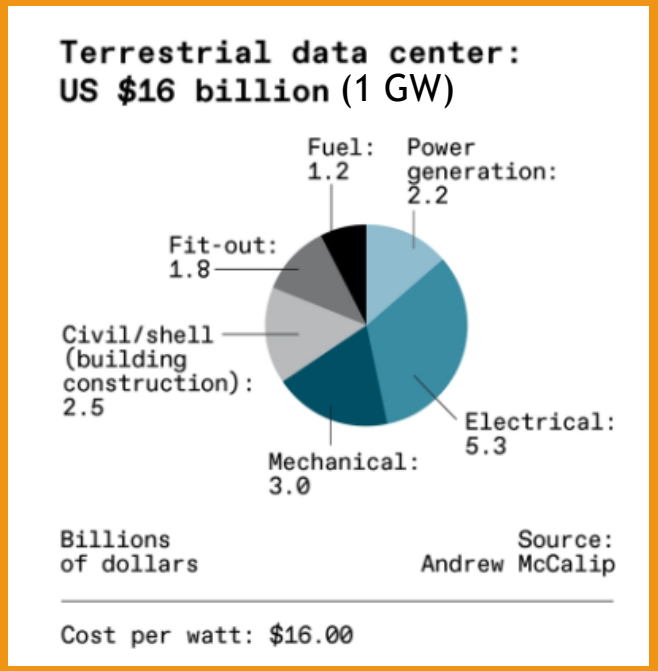

Boston Consulting Group (BCG) cost sensitivity analysis:
Launch Cost & Improved GPU Efficiency: Close the Gap

Launch cost ($/kg)	Satellite weight (tons/100 kW): 2 tons/100 kW (baseline)	1 ton/100 kW (50% improvement)	0.5 ton/100 kW (75% improvement)
$1,500/kg (baseline)	2.3X	1.6X	1.3X
$500/kg (67% improvement)	1.9X	1 1.5X	1.2X
$100/kg (93% improvement)	1.8X	1.4X	2 1.2X

1 Optimistic scenario 2 Aggressive scenario

BCG selected two specific scenarios, Optimistic & Aggressive to

Boston Consulting Group (BCG) economic view for 2030 & 2040

SBDC: Space-Based Data Center

		Conservative scenario	Optimistic scenario	Aggressive scenario
		SBDCs remain limited to niche use cases (such as sovereign AI workloads and space-generated data)	SBDCs become a complementary layer, capturing nonsensitive inferencing and niche use cases	SBDCs become a major market segment for nonsensitive inferencing
2030	Overall SBDC capacity	Negligible; focused on niche use cases	1 GW–2 GW	2 GW–3 GW
2030	Share of total AI data center market	Negligible; focused on niche use cases	1%–2%	2%–3%
2030	Total revenue	Negligible; focused on niche use cases	$8 billion–$16 billion	$16 billion–$24 billion
2040	Overall SBDC capacity	10 GW–20 GW	30 GW–40 GW	50 GW–100 GW
2040	Share of total AI data center market	3%–5%	10%–15%	20%–30%
2040	Total revenue	$80 billion–$160 billion	$240 billion–$320 billion	$400 billion–$800 billion

Ideal SBDC Workloads:

- ***Latency-tolerant:*** batch document, image, and video generation, enterprise back-office AI, (and bulk translation and tagging.
- ***Sovereign AI workloads:*** government, health, tax, secondary backup capability
- ***Space-generated data processing:*** EO edge processing

BCG Link

The End of Moore's Law

CMOS process technology is reaching its fundamental physical unsustainable heat densities that conventional cooling methods cannot dissipated at the 1 nm to 1.4 nm nodes (~2029)

- Geometric scaling ends: 2 to 3 atomic layers deep (0.7nm)
- Quantum Tunneling - catastrophic current leakage and immense power
- Boltzmann Tyranny & Overheating: unsustainable heat densities that conventional cooling methods cannot dissipate

Fabrication Node Label	Physical Gate Length	Physical Phenomenon/Limit
3nm	~10–18 nm	Transition from FinFET to Gate-All-Around (GAA) to maintain control
2nm	~9–12 nm	Gate oxides thin down to 2–3 atomic layers. High leakage begins.
1nm-1.4nm	~4–5 nm	**The Silicon Wall:** Quantum tunneling dominates; silicon loses ability to shut off completely.

Further performance increases will start with 2/2.5D "suburban sprawl," then move to a vertical 3D "skyscraper" integration paradigm

- 2D expansion is limited by long interconnects, latency, and signal integrity
- 3D expansion is limited to ~4 layers due to thermal accumulation (125°C to 150)°C

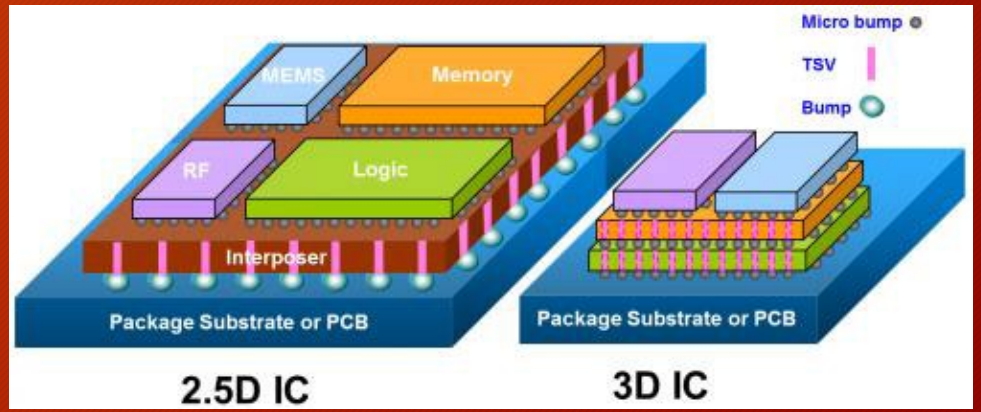

Final Thoughts …
The 2D-3D Wafer Scale - The Last Frontier

- Industry is pursuing both 2D & 3D expansion models to compute solutions
- 2D limit example: Cerebras Wafer Scale Engine WSE-3 Turbo:
  - 46,225 mm^2,4 trillion transistors, 250 petaflops AI compute via 900,000 AI-optimized cores, 56x NVIDIA B300
- The ultimate limit is the 3D wafer-scale solution limit of thermal management

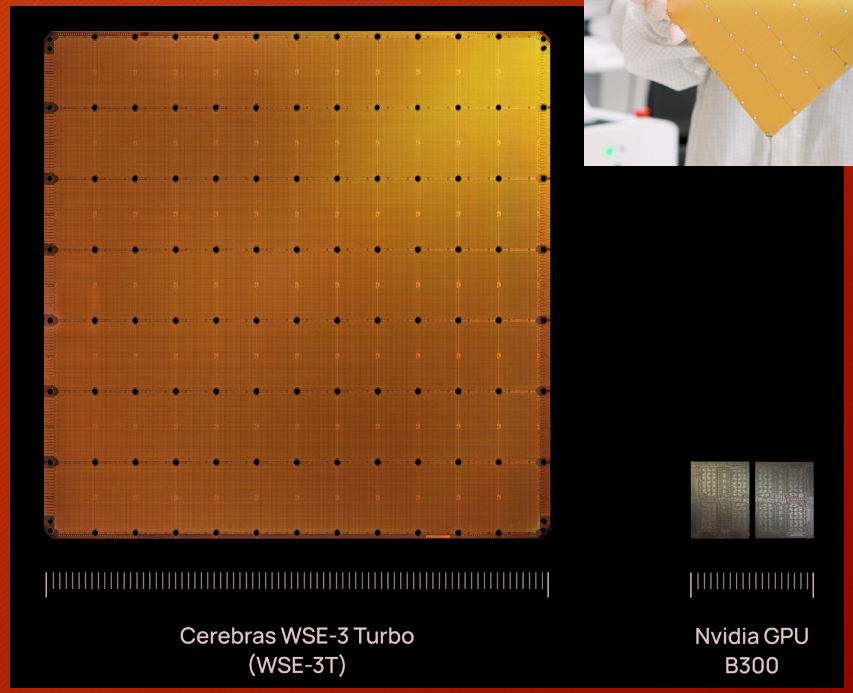

Slowing/End of Moore's Law Extends AI Compute Lifespan:
Improves & Eventually Favors SBDC Economics

Thank You

Space is the new theater of war, and EO Reconnaissance and Edge AI are critical for US and Allied deterrence. Find out more in my new book:

About the Book

About Me

Thank You

Questions/Comments Contact: paul@struhsaker.com
Visit the Arrasar Web Page: www.struhsaker.com/arrasar-partners
PDF download available at: www.struhsaker.com/dystopic-science-stories