Terrestrial (land-based) data centers remain essential for the vast majority of computing needs today—and likely for decades—because they offer unmatched practicality, cost-effectiveness, and performance for most real-world workloads. While alternatives like underwater pods (e.g., Microsoft's now-shelved Project Natick) or orbital/space-based concepts get attention for addressing power and cooling challenges, they introduce trade-offs that make them unsuitable as full replacements. Here's why land-based data centers are still fundamentally necessary:1. Easy Access for Maintenance and Upgrades Data centers are full of hardware that fails, needs replacement, or requires rapid upgrades—especially with fast-evolving AI chips (new generations often every 1–2 years). On land: Technicians can physically access servers in minutes or hours, swap parts, and restore service quickly.

This allows high utilization rates with minimal over-provisioning (extra backup hardware).

In contrast, underwater systems require surfacing entire sealed pods for repairs (logistically complex and expensive), and space-based ones have no practical on-site maintenance at all—you'd need to launch replacements or design everything with extreme redundancy. Experiments like Natick showed lower failure rates in sealed environments, but scaling and servicing proved impractical for ongoing operations.

itpro.com

2. Low Latency and High-Bandwidth Connectivity - Most data processing serves terrestrial users and applications—web browsing, streaming, cloud services, financial transactions, AI inference for apps, enterprise software, etc. These demand near-instant response times. Fiber-optic networks on land provide massive bandwidth with milliseconds of latency.

Data centers are often sited near population centers, internet exchange points, or users (e.g., dense clusters in Northern Virginia).

Underwater or space options add delays: satellite links introduce hundreds of milliseconds of latency (unacceptable for many interactive or real-time uses), and bandwidth is far more constrained and expensive than terrestrial fiber. Orbital data centers might suit space-specific tasks (processing satellite data in orbit), but not general Earth-based needs.

gartner.com

3. Practical Cooling and Heat Dissipation Servers generate enormous heat (modern AI clusters can consume hundreds of MW). On land: Efficient convective cooling uses air, water, or evaporative systems, often leveraging ambient air, rivers, or lakes as "infinite" heat sinks.

Techniques like free cooling in cold climates further reduce energy use.

Underwater pods benefit from seawater cooling (one reason Natick worked technically), but deployment is limited. In space, there's no air or water—cooling relies on radiative panels in a vacuum, which is far less efficient per unit area and requires massive, complex radiators plus active fluid loops.

Heat rejection in orbit is a major engineering hurdle that drives up costs and size dramatically.

andrewmccalip.com

4. Power Availability and Economics Land-based facilities can connect directly to the electrical grid, tap renewables, natural gas, or nuclear sources, and scale with existing infrastructure. While power grids face strain from AI growth, solutions like new transmission lines, on-site generation, or siting in energy-rich areas (e.g., near hydro or stranded gas) are feasible and far cheaper than alternatives. Building and operating on land avoids the extreme expenses of launch mass, radiation hardening, or specialized marine/space engineering.

Cost analyses show orbital data centers could be 3–10x (or more) more expensive per unit of compute, even with falling launch costs—due to power delivery, hardware lifespan, and redundancy needs.

spectrum.ieee.org

Underwater power delivery is tricky (cables from shore), and space relies on solar arrays (abundant but limited by panel size/mass launched).5. Reliability, Security, and Scalability Radiation protection: Earth's atmosphere and magnetic field shield hardware from cosmic rays and solar radiation, which can cause bit flips or failures in space. Orbital systems need heavy shielding or triple redundancy, inflating costs.

Hardware refresh: AI workloads chew through chips quickly; land allows easy swaps. In space or underwater, you're often locked into older hardware for years.

Regulatory and environmental realities: Land sites involve permitting and community trade-offs (land use, water for cooling, grid impact), but these are manageable compared to international treaties for ocean floors or orbital debris risks (Kessler syndrome from thousands of satellites).

Terrestrial data centers do face real pressures—rising energy demand, land/water use, and local opposition in some areas—but they benefit from mature supply chains, existing infrastructure, and the ability to innovate incrementally (better efficiency, advanced cooling, nuclear co-location, etc.).When Alternatives Might Play a Niche Role Underwater: Great for passive cooling in tests, but abandoned at scale due to maintenance and upgrade limitations.

Space/orbital: Promising for edge cases like processing data from satellites in orbit (reducing downlink needs) or ultra-secure backup, especially with abundant solar power. However, they won't handle mainstream terrestrial computing anytime soon due to physics, economics, and latency.

In short, we need terrestrial data centers because computing is a physical, human-serviced, Earth-connected activity. The planet provides cheap cooling sinks, instant access, robust networks, and protection from space's harsh environment—at a fraction of the complexity and cost of off-world alternatives. As AI and data demands grow, land-based infrastructure will continue to dominate, with innovations focused on making it more efficient and sustainable rather than abandoning it entirely.