Scientific Reports

https://doi.org/10.1038/s41598-026-38784-1

Article in Press

O-RAID: a satellite constellation architecture for ultra-resilient global data backup

Received: 19 November 2025

Accepted: 31 January 2026

Published online: 10 February 2026

Cite this article as: Meegama R.G.N. O-RAID: a satellite constellation architecture for ultra-resilient global data backup. *Sci Rep* (2026). https://doi.org/10.1038/s41598-026-38784-1

R. G. N. Meegama

O-RAID: A Satellite Constellation Architecture for Ultra-Resilient Global Data Backup

R.G.N. Meegama[1,*]

[1]Apple Research and Development Center, Department of Computer Science, Faculty of Applied Sciences, University of Sri Jayewardenepura , Gangodawila, Nugegoda, Sri Lanka
[*]rgn@sjp.ac.lk

ABSTRACT

Growing global data volumes and the increasing frequency of climate-related and geopolitical threats highlight the need for ultra-resilient backup infrastructures. This paper proposes a novel Satellite-RAID architecture, named O-RAID, in which clusters of satellites operate as a distributed redundant array of independent disks (RAID), enabling large-scale cold and warm backup storage in Earth's orbit. Unlike previous work on space-based computing or satellite cloud relays, this research presents a formal design for orbital storage redundancy, inter-satellite parity exchange, latency-tolerant RAID protocols and power provisioning using a geostationary solar-energy beam. To establish a foundation for quantifying system resilience, we develop a reliability framework based on a Continuous-Time Markov Chain (CTMC) model, defining the states and transition rates for future survivability analysis of an orbital RAID equivalent. The paper provides a comprehensive analysis of the system architecture, its core components and the mathematical underpinnings for erasure coding and communication. An in-depth examination of system feasibility, survivability simulations, key constraints and communication overhead is presented, concluding that orbital backup storage presents a viable and promising paradigm for national archives, disaster-resilient storage and long-term scientific data preservation with technical readiness projected by 2035.

1 Introduction

The increasing dependence of global society and critical infrastructure on massive digital archives necessitates a fundamental shift in data storage resilience strategies. Traditional terrestrial data centers, while achieving high uptime metrics, remain systemically vulnerable to emerging global risks. This paper proposes O-RAID, a novel architectural solution leveraging satellite constellations to provide an inherently secure, geographically detached and ultra-resilient platform for cold and warm backup storage.

1.1 The Zettabyte Challenge and the Architectural Imperative for Resilience

Global data generation is expanding at an exponential rate, presenting an unparalleled scalability challenge. Forecasts indicate that the global data-sphere is expected to reach 163 zettabytes (ZB) by 2025, representing a tenfold increase from 2016 figures[1]. This expansion is sustained by consistent annual growth rates, projected to remain above 22% year-over-year through 2025[2]. This massive influx of data demands storage architectures capable of scaling rapidly and accommodating these extreme volumes without suffering from the resource bottlenecks inherent to terrestrial infrastructure development.

The continued scaling of ground-based data centers is challenged by critical environmental and infrastructural limitations. Terrestrial facilities require vast land resources, consume immense amounts of electricity and demand staggering volumes of water, primarily for cooling purposes[3]. For example, some individual data centers consume millions of gallons of water annually, straining local utility resources and contributing to localized energy pricing pressure[3].

The O-RAID architecture offers a paradigm shift in scalability by utilizing the orbital environment, which provides boundless space and eliminates traditional cooling costs entirely. This approach facilitates a modular, growth-tolerant design. Since O-RAID comprises clusters of independent satellites, the system inherently supports incremental and rapid scaling through successive launches of new clusters. This modularity allows the storage capacity to match the predicted 23% annual data growth far more flexibly and quickly than the lengthy, capital-intensive expansion cycles required for constructing and commissioning large-scale terrestrial facilities[2].

1.2 Critical Vulnerabilities of Ground-Based Data Infrastructure

Relocating critical archives to orbit is a strategic response to the increasing frequency and intensity of existential threats facing centralized terrestrial data stores.

1.2.1 Climate-Induced Disaster Risks

Terrestrial data centers face acute and growing risks from climate volatility, including extreme heat, drought and flooding. Analysis indicates that climate-related hazards could drive annual operational costs up by USD 81 billion globally by 2035, escalating to USD 168 billion by 2065[4]. These severe weather events pose direct threats to operational continuity and physical infrastructure integrity. The orbital location of the O-RAID system provides complete immunity to such geographic and atmospheric disruptions, ensuring that critical data archives are physically detached from regions prone to natural disaster.

1.2.2 Geopolitical and Cyber Exposure

The digital landscape is increasingly defined by geopolitical rivalry and sophisticated cyber threats. The growth of AI-driven technology and escalating geopolitical tensions compound the vulnerability of centralized digital infrastructure, making data centers attractive targets for cybercriminals and state actors[5]. Furthermore, ground-based facilities are susceptible to localized electromagnetic disturbances and physical attacks linked to geopolitical instability. The fundamental motivation for O-RAID lies in the strategic transformation of the risk profile. By placing cold backup storage off-planet, the system exchanges predictable, quantifiable terrestrial risks (such as power grid failures, climate hazards and high cyber exposure) for manageable, high-engineering-cost orbital risks (such as radiation exposure and orbital debris collision). The subsequent reliability analysis aims to demonstrate that, through robust engineering, the statistical probability of data loss in the orbital environment is lower than the cumulative probability of catastrophic terrestrial failure.

1.3 Deficiencies in Existing Space Data Architectures

While the concept of space-based data handling is not new, existing architectures lack the critical elements necessary for resilient, large-scale data archiving.

1.3.1 Relay-Centric Designs

Prior work on utilizing space assets for data management has historically focused on high-speed data relay rather than persistent, resilient storage. Systems like NASA's Tracking and Data Relay Satellite System (TDRS) and burgeoning commercial constellations prioritize high-speed transmission Direct-to-Earth (DTE)[6]. These systems function primarily as communication conduits, buffering data briefly (e.g., for 7 days for retransmissions) before downlink[7]. They are not designed to host massive, long-term backup archives with guaranteed fault tolerance across multiple components.

1.3.2 Communication Limitations

Current DTE-optimized architectures struggle with the communication constraints inherent in space-based networking. High propagation losses, which are a function of the reciprocal of the distance squared, mean that achieving high data rates between low-Earth orbit (LEO) spacecraft and geostationary (GEO) relay satellites requires significantly higher power than DTE communication[7]. Moreover, traditional Radio Frequency (RF) links are increasingly constrained by regulatory limitations on available frequency bands, restricting scalability and data throughput[6]. This necessitated the incorporation of Optical Inter-Satellite Links (ISL) into the O-RAID design, as optical links provide much higher data rates and generally utilize smaller satellite form factors[6]. These high-bandwidth ISLs are essential for the constant internal data exchange required for parity synchronization, functioning independently of intermittent DTE downlink constraints.

1.3.3 Absence of Formal Redundancy Protocols

Existing concepts for distributed space computing primarily address computational decentralization and networking optimization[8]. However, these systems do not define or formally implement a concrete redundancy structure suitable for fault-tolerant data storage analogous to established terrestrial RAID specifications. The innovation of O-RAID lies in introducing a formalized, mathematically quantifiable **orbital redundancy structure**. Here, clusters of physical satellites are treated logically as independent storage units within a distributed array, enabling complex data striping and parity exchange mechanisms.

1.4 Motivation

The rapid expansion of digital infrastructure has exposed fundamental constraints in the scalability of terrestrial data centers that extend beyond storage technology itself. In many regions, the deployment of new data centers, including geographically independent backup facilities, is increasingly limited by long permitting timelines, power availability, water scarcity, environmental regulation and public opposition. Recent analyses indicate that large-scale facilities commonly require 18 to 36 months from planning to commissioning, even in technologically advanced economies[9,10]. As a result, redundancy and disaster recovery strategies that rely exclusively on additional terrestrial sites are themselves subject to significant delays and correlated regional risks.

These constraints have become particularly acute with the rise of AI-driven workloads, which demand unprecedented power density and cooling capacity.[11] In this context, infrastructure deployment speed rather than logical redundancy or distributed software architectures, has emerged as a primary bottleneck. Even highly decentralized terrestrial storage systems remain

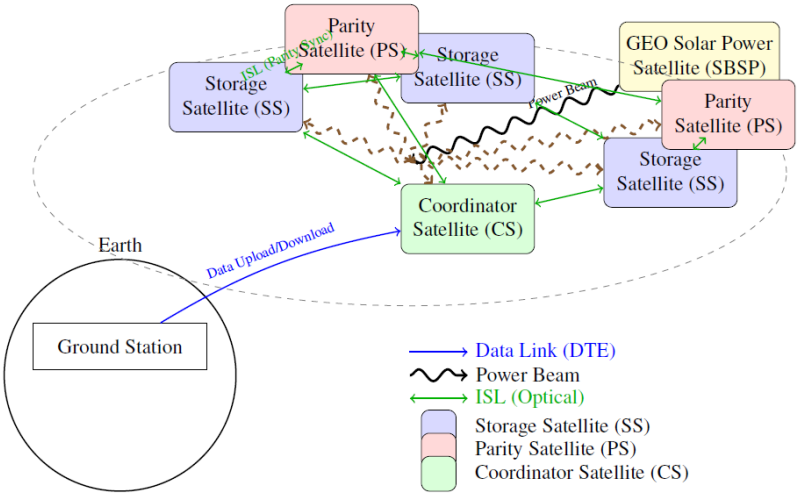

Figure 1. Proposed O-RAID architecture with clearly separated satellite nodes.

ultimately dependent on Earth-bound facilities, power grids and regulatory processes that limit the ability to provide rapid, physically independent archival storage.[12]

In response to these pressures, industry and government stakeholders have begun exploring space-based and orbital data-center concepts as a practical extension of terrestrial computing ecosystems . These efforts are motivated not by speculative catastrophe scenarios but by near-term challenges related to energy density, thermal management, land use and deployment speed[13,14]. The recent launch of the first GPU payload into orbit in late 2025 represents a significant milestone in this direction marking an early transition from conceptual studies to operational experimentation in on-orbit computing[15,16].

Within this evolving landscape, O-RAID is proposed as an architectural framework for rapidly deployable constellation-level archival storage that decouples long-term data survivability from terrestrial permitting delays, resource constraints, and geographically correlated risks. Instead of replacing existing terrestrial or decentralized storage systems, O-RAID complements them by providing a physically independent archival layer optimized for cold and long-term data preservation.

2 Proposed O-RAID Architecture

The O-RAID architecture is composed of three distinct, specialized classes of satellites, each fulfilling a critical role within the distributed storage array. This modular design is fundamental to achieving the system's goals of scalability, fault tolerance and efficient resource management. The interplay between these components transforms a simple constellation into a cohesive, high-performance orbital data center as shown in Figure 1.

2.1 Storage Satellites (SS): The Archival Workhorses

Storage Satellites (SS) form the foundational data layer of the O-RAID system. They are responsible for the persistent, secure and reliable storage of the primary data blocks.

Core Function:

Each SS functions as an independent "disk" in the traditional RAID analogy. Its primary purpose is to store user data blocks (D_i) assigned to it by the Coordinator Satellite (CS). It does not perform parity calculations but is responsible for the integrity of the data in its custody.

Hardware Specifications:

- Radiation-Hardened Storage Arrays: The core of an SS is a high-density array of Solid-State Drives (SSDs) specifically designed and qualified for the space environment. These drives are hardened against single-event upsets (SEUs) caused by cosmic rays and solar radiation, employing techniques like Error Correction Code (ECC) memory, shielding and proprietary controller designs to ensure bit-level integrity over a multi-year mission lifespan.

- Processing and Memory: Each SS contains a moderate-performance onboard computer with sufficient volatile memory (RAM) to manage file systems, handle communication protocols and perform local data integrity checks (e.g., cyclic redundancy checks - CRCs). Its computational load is intentionally kept low to maximize reliability and power efficiency for its primary storage function.

- Communication Subsystem: An SS is equipped with an optical ISL terminal. This terminal is optimized for high-throughput, reliable data reception and transmission with other satellites in the constellation (both other SSs and Parity Satellites). It typically does not require a high-gain DTE antenna, as all user data I/O is orchestrated through the Coordinator.

- Operational Role: Upon receiving a write command from the CS, an SS stores the data block and returns an acknowledgment. For a read request, it retrieves and transmits the requested block. It continuously reports its health and status (telemetry) to the CS and participates in the distributed data rebuild process by serving its stored blocks to a replacement satellite when a peer fails.

2.2 Parity Satellites (PS): The Guardians of Redundancy

Parity Satellites (PS) are the cornerstone of the system's fault tolerance. They are dedicated to calculating, storing and managing the redundancy information that enables data recovery in the event of satellite failures.

Core Function:

PS units are responsible for storing the parity blocks (P and Q) for each data stripe. Crucially, they also perform the computationally intensive Galois Field (GF) arithmetic required to generate and update these parity blocks. By offloading this demanding task from the Storage Satellites, the system achieves a more balanced power and thermal profile.

Hardware Specifications:

- High-Performance Computing (HPC) Module: The defining feature of a PS is its powerful, radiation-hardened processor, capable of efficient GF multiplication and division. This may involve specialized hardware accelerators or Field-Programmable Gate Arrays (FPGAs) configured for finite field operations, significantly reducing the time and energy required for parity calculation—a critical factor in mitigating the system's "write penalty."

- Dedicated Parity Storage: While also using radiation-hardened SSDs, a PS's storage capacity is tailored to hold the (P and Q) blocks for the entire array. Its storage requirements are a function of the number of SS units and the stripe size.

- Advanced Communication Subsystem: A PS requires a high-performance optical ISL terminal, often with greater capability than that of an SS. It must handle a constant, high-volume traffic load during parity updates and rebuilds, receiving data blocks from multiple SSs and transmitting calculated parity blocks to other PSs.

Operational Role:

During a write operation, a PS receives the new data block and the old data/parity blocks (for a read-modify-write cycle), performs the GF arithmetic to compute the updated P and Q values, and stores them. During a rebuild, it works in concert with the CS and the surviving SSs to solve the system of linear equations and reconstruct lost data or parity blocks. For resilience, $P$and Q blocks for a given stripe are always stored on separate, physically distant PSs.

2.3 Coordinator Satellites (CS): The Orchestrating Intelligence

The Coordinator Satellites (CS) are the "brains" of the O-RAID constellation. They manage the metadata, coordinate all I/O operations and enforce the system-wide protocols that guarantee consistency and atomicity across the distributed array.

Core Function:

The CS maintains the global "map" of the array. It tracks which data and parity blocks are stored on which specific satellites, manages the allocation of new data and orchestrates the complex distributed protocols for writes, reads, failure recovery and rebuild operations.

Hardware Specifications:

- Redundant High-Availability Compute Platform: A CS features a highly reliable, often redundant, computing platform. Given its critical role, it may employ lockstep processors or other fault-tolerant computing techniques to ensure it does not become a single point of failure.

- Substantial Memory and Cache: It possesses large amounts of non-volatile memory and RAM to cache the metadata table, track the state of ongoing transactions and buffer data during complex operations.

- Multi-Band Communication Suite: A CS is equipped with the most comprehensive communication system:

- High-Gain DTE Antenna: For receiving uploads from ground stations and transmitting downloaded data to users.

- High-Capacity Optical ISL Terminal: For managing traffic with all SS and PS units within the constellation.

- Cross-Link with other CSs: To maintain a consistent, replicated metadata state across multiple Coordinators for fault tolerance.

Operational Role:

- Metadata Management: The CS holds the crucial mapping of logical block addresses to physical satellite locations.
- Transaction Coordinator: It implements the two-phase commit protocol for all write operations, ensuring that updates to data and parity blocks across multiple satellites are atomic. The proposed two-phase commit abstraction is designed to prioritize safety under partitioned conditions rather than immediate liveness. When ISL disruptions or network partitions occur, write operations that cannot be coordinated across the required set of participants are conservatively prevented from completing, ensuring that partial or inconsistent updates are never externally visible. Once connectivity is restored, pending operations may either be completed or safely rolled back depending on the observed system state.
- Failure Detection and Recovery Manager: The CS continuously monitors the health of all SS and PS units. Upon detecting a failure (via loss of heartbeat), it updates the system state, marks the failed satellite's blocks as invalid and initiates the rebuild process onto a pre-designated Hot Spare Satellite (HSS).
- Load Balancer: It distributes new data writes evenly across the available SS units to prevent hotspots and ensures that parity update traffic is balanced across the PS units.
- Synergy and Interdependence: The strength of the O-RAID architecture lies in the seamless interaction of these three components. The SS provides scalable, dumb storage. The PS provides the sophisticated mathematical redundancy. The CS provides the cohesive intelligence that binds them into a single, fault-tolerant system. This separation of concerns allows for independent optimization of each satellite type for its specific function, leading to a more robust, efficient and manageable orbital data center than a homogeneous constellation could provide.

Data Striping Structure

Data uploaded from Earth is divided into k data blocks and m parity blocks, with $m = 2$ for RAID equivalence (eg. RAID-6). Blocks are distributed across satellites.

Inter-Satellite Parity Mechanism

Parity is computed using a rotating parity schedule to balance load. Optical laser links provide the bandwidth required to synchronize blocks.

Power Infrastructure

A geostationary solar-power satellite provides continuous energy via microwave or laser beaming. This prevents battery degradation and ensures uninterrupted parity operations.

3 Reliability Modeling and System Mathematics

To validate the feasibility and resilience of O-RAID, a rigorous mathematical framework is required, spanning survivability analysis, erasure coding and communication performance.

3.1 System Reliability via Continuous-Time Markov Chain (CTMC) Analysis

The survivability of the O-RAID system over time, particularly its Mean Time To Data Loss (MTTDL), must be modeled using Continuous-Time Markov Chains (CTMCs)[17]. This approach will capture the dynamic process of satellite failure, detection and subsequent orbital replacement, which is critical in a space environment where repair times are exceptionally long.

3.1.1 CTMC State Definitions

The O-RAID architecture utilizes N total operational storage satellites, designed to tolerate the simultaneous failure of any two independent satellites without data loss. The CTMC model defines states based on the number of failed satellites that have not yet caused system failure:

- S_0: The Normal, Start State. All N satellites are fully operational (0 failed).
- S_1: One satellite has failed. Data integrity is maintained by the RAID redundancy and data rebuild to a spare unit is initiated.
- S_2: Two satellites have failed. The system is operating in a critical, high-risk state. Rebuild operations for both failed units are active.
- S_F: The Absorbing State. Three or more satellites have failed, resulting in permanent data loss.

3.1.2 Definition of Transition Rates

The transitions between these states are governed by two primary rates:

1. **Satellite Failure Rate (λ):** This rate, measured in failures per unit time (e.g., year), is assumed constant for the core analytical solution, calculated as the inverse of the Mean Time To Failure ($\text{MTTF}_{\text{satellite}}$) for a single hardened storage satellite.

2. **Orbital Replacement/Repair Rate (ρ):** This rate is defined as the inverse of the Mean Time To Repair ($\text{MTTR}_{\text{orbital}}$). It represents the logistical time required for a failed satellite to be detected, its replacement unit launched and the data fully rebuilt onto the new, operational unit[17]. Unlike terrestrial disk repair rates, ρ is constrained by orbital mechanics, launch windows and maneuvering time, making it significantly lower.

3.1.3 Derivation of the State Transition Matrix

The CTMC transition rate matrix $\mathbf{\Omega}$ describes the instantaneous rates of movement between the non-failure states (S_0, S_1, S_2). The evolution of the probability vector $\mathbf{P}(t) =$ is solved using the Kolmogorov forward equation, $\frac{d\mathbf{P}(t)}{dt} = \mathbf{P}(t)\mathbf{\Omega}$[18]. System reliability $R_{sys}(t)$ is the sum of the probabilities of being in the non-failure states: $R_{sys}(t) = P_{S_0}(t) + P_{S_1}(t) + P_{S_2}(t)$.

For a RAID system (eg. RAID-6) with N satellites, the transition matrix $\mathbf{\Omega}$ is:

$$\mathbf{\Omega} = \begin{pmatrix} -N\lambda & N\lambda & 0 \\ \rho & -((N-1)\lambda + \rho) & (N-1)\lambda \\ 0 & 2\rho & -((N-2)\lambda + 2\rho) \end{pmatrix} \tag{1}$$

The entries dictate the following:

- $S_0 \to S_1$: A single failure occurs at rate $N\lambda$.
- $S_1 \to S_0$: The failed satellite is repaired/replaced at rate ρ.
- $S_1 \to S_2$: A second failure occurs among the remaining $N-1$ operational satellites at rate $(N-1)\lambda$.
- $S_2 \to S_1$: One of the two failed satellites is repaired/replaced at rate 2ρ.

A critical transition is the Data Loss Transition from S_2 to the absorbing state S_F, which occurs if a third satellite fails while the system is already coping with two failures. The rate of this transition is $(N-2)\lambda$. The inherent difficulty of orbital logistics means that the replacement rate ρ is low, resulting in extended residence time in the high-risk state S_2 while awaiting a functioning replacement satellite. The likelihood of a third failure occurring during this extended exposure time drastically increases the probability of transition to S_F. This mathematical dependency emphasizes that the system's overall resilience hinges not only on component reliability (λ) but critically on the speed of the data rebuild process (which directly impacts the effective ρ value).

3.2 Formal Erasure Coding and Parity Reconstruction

Achieving dual-failure tolerance (eg. RAID-6) requires sophisticated erasure coding techniques, specifically implemented over finite fields to guarantee the ability to reconstruct data from any k out of N blocks.

3.2.1 Galois Field (GF) Arithmetic Basis

The coding operations are conducted over a Galois Field $\text{GF}(2^w)$, typically using $w = 8$ (i.e., $\text{GF}(2^8)$), where byte values map directly to field elements[19]. Field addition ($+$) is performed via the bitwise exclusive-OR (XOR) operation ($\oplus$). Field multiplication and division are defined using polynomial arithmetic modulo a predefined irreducible polynomial.

3.2.2 P-Parity and Q-Parity Equations

For a stripe size consisting of k data blocks ($D_1, D_2, \ldots, D_k$), the O-RAID architecture generates two parity blocks, $\mathbf{P}$ (simple parity) and $\mathbf{Q}$ (generalized parity), which are stored on separate Parity Satellites (PS).

$$\mathbf{P} = \bigoplus_{i=1}^{k} D_i, \quad \mathbf{Q} = \bigoplus_{i=1}^{k} \alpha^i \cdot D_i \tag{2}$$

where α^i are distinct, non-zero elements chosen from $\text{GF}(2^w)$ (e.g., $1, \alpha, \alpha^2, \ldots$) that ensure the generating matrix is invertible.

3.2.3 Data Reconstruction Logic (Dual Failure)

If two arbitrary blocks, say D_i and D_j, are lost due to the simultaneous failure of two satellites, the remaining operational satellites are used to compute residual parity values, P' and Q'. This yields a system of two linear equations in GF(2^w):

$$D_i \oplus D_j = P' \quad \alpha^i D_i \oplus \alpha^j D_j = Q'$$

Solving this system involves sophisticated Galois Field linear algebra, including matrix inversion, to derive the values of D_i and D_j.

This complex mathematical process dictates performance trade-offs. The calculation of the **Q** parity and the subsequent matrix inversion during reconstruction are computationally demanding. In a standard terrestrial RAID system, the complexity of these operations contributes to a "write penalty" that can be six times the read/write penalty of mirrored data[20]. For O-RAID, this high computational load must be executed across a distributed network characterized by inherent ISL latency. Therefore, the implementation must rely on highly optimized erasure codes (such as specific Reed-Solomon variants or advanced codes like Liberation Codes[21]) to ensure that the write penalty is minimal and that the rebuild process, $T_{rebuild}$, remains feasible despite the high latency.

While this analysis uses a dual-parity (eg. RAID-6) model for benchmarking, the O-RAID architecture supports arbitrary Erasure Coding (EC) schemes, including Locally Repairable Codes (LRC) and Raptor codes.

3.3 Communication Performance and Latency Modeling

The viability of parity synchronization and, critically, the data rebuild process is dependent on the performance and stability of the optical ISL.

3.3.1 ISL Budget Derivation

The performance of an optical ISL is evaluated by its Link Margin (LM), which must be positive to ensure the required data rate (B) is sustained at a target Bit Error Rate (BER)[22]. The received power (P_{rx}) is calculated using the fundamental optical link budget equation (expressed in dBm):

$$P_{rx} = P_{tx} + OE_{tx} + OE_{rx} + G_{tx} + G_{rx} - LP_{tx} - LP_{rx} - LP_S \tag{3}$$

where P_{tx} is the transmitted power, OE terms represent optical efficiencies, G terms represent antenna gains and LP_S is the free-space path loss (a constant loss factor for a given separation distance).

A major factor in the link budget for dynamic satellite constellations is the pointing loss, LP. These terms quantify the signal degradation resulting from minor misalignment between the transmitting and receiving satellite platforms[22] as follows:

$$LP_{tx} = \mathbf{4.3429} \times \mathbf{g_{tx}} \times (Pe_{tx})^{\mathbf{2}} \quad LP_{rx} = 4.3429 \times g_{rx} \times (Pe_{rx})$$

where Pe represents the pointing error in radians and g is the linear gain.

3.3.2 Pointing Loss Sensitivity

The pointing loss terms are proportional to the square of the pointing error (Pe^2). In a LEO or MEO constellation, disturbances from orbital dynamics, internal spacecraft vibrations and thermal fluctuations introduce micro-radian-level pointing errors. Because LP is exponentially sensitive to Pe, even minute tracking instabilities can lead to significant, dynamic fluctuations in received signal power. Therefore, link performance is not static; it requires modeling the pointing error (Pe) as a stochastic variable within the simulation to accurately capture periods where the link margin drops below critical thresholds, potentially slowing or halting the parity rebuild process.

3.3.3 Effective Parity Rebuild Time ($T_{rebuild}$) and Write Penalty

The theoretical minimum time required to rebuild data onto a replacement unit is given by

$$T_{rebuild,theoretical} = \frac{(k-1)S}{B}$$

where S is the block size. However, in a distributed orbital system, this must be adjusted for network contention and intrinsic latency[23]. The effective rebuild time is defined as:

$$T_{rebuild,effective} = \frac{(k-\mathbf{1})S}{\alpha \cdot B_{\mathrm{ISL}}} + L_{latency} \tag{4}$$

where k is the number of data blocks, B_{ISL} is the nominal inter-satellite link capacity (e.g., 10 Gbps) and $L_{latency}$ is the latency incurred by orbital distance and routing hops. The critical variable is α, the **Bandwidth Availability Factor**, which represents the fraction of total ISL bandwidth specifically dedicated to the rebuild operation after accounting for all necessary housekeeping, telemetry and non-rebuild synchronization traffic. By including α, the model realistically captures performance degradation during high-load scenarios.

3.4 Continuous Power Provisioning Model (SBSP Integration)

Reliable operation of the high-density SSD arrays and continuous optical transceiver usage demands a stable, high-power supply independent of solar eclipse cycles and battery degradation. This necessitates integration with a Space-Based Solar Power (SBSP) system[24].

3.4.1 GEO-to-ORBIT Energy Transfer

The proposed design utilizes a dedicated solar power satellite operating in Geostationary Earth Orbit (GEO) to collect sunlight 24/7 and beam power continuously to the LEO/MEO O-RAID constellation[24]. While laser beaming offers compact size, microwave beaming generally provides higher overall system efficiency for high-power transfer, with estimates reaching up to 41%[25]. Given the constant power demands for high-throughput optical links and solid-state drive maintenance, high transfer efficiency is paramount.

The received power at an O-RAID Storage Satellite ($P_{SS,received}$) is modeled by:

$$P_{ss,received} = P_{GEO,transmitted} \times \eta_{GEO} \times \eta_{transfer} \tag{5}$$

where $P_{GEO,transmitted}$ is the microwave power radiated from the GEO station, η_{GEO} is the efficiency of the GEO collector and transmitter and $\eta_{transfer}$ is the end-to-end efficiency of wireless power transfer, including atmospheric/propagation losses and rectenna conversion losses at the receiving satellite[25]. This reliance on SBSP removes constraints imposed by onboard battery cycling, which typically degrades over time, thus ensuring the continuous availability of power necessary for background P and Q parity synchronization and preemptive maintenance.

Thermal dissipation associated with rectenna inefficiency is assumed to be managed through standard spacecraft thermal control techniques, including radiative heat rejection and power-reception duty cycling, which are outside the scope of this work.

4 Methodology

The O-RAID concept is validated through a combination of detailed orbital mechanics modeling, protocol specification and exhaustive Monte Carlo simulations.

4.1 High-Fidelity Orbital Mechanics and Network Topology Modeling

The physical implementation of O-RAID dictates the connectivity and resilience of the distributed array.

4.1.1 Orbital Configuration Modeling

The choice of orbital configuration is not merely a matter of placement; it is the foundational determinant of the system's physical stability, communication latency and ultimately, its data integrity and Mean Time to Data Loss (MTTDL). Our high-fidelity modeling of three distinct configurations—LEO Cluster, MEO Mesh and Hybrid LEO-MEO—reveals a critical trade-off between accessibility and resilience.

1. **LEO Cluster:** A low-Earth orbit cluster (e.g., 500-1000 km altitude), offering low latency to ground but facing higher atmospheric drag and increased vulnerability to orbital debris. The LEO cluster offers the significant advantage of low latency for both ISLs and DTE communication, theoretically enabling faster parity updates and rebuilds. However, this advantage is counterbalanced by substantial physical vulnerabilities. At these altitudes, atmospheric drag is non-negligible, requiring frequent station-keeping maneuvers that consume propellant and limit satellite lifespan, directly increasing the effective failure rate λ. Furthermore, LEO is the most densely populated orbital regime, presenting a higher statistical probability of collision with both operational satellites and debris[26].

2. **MEO Mesh:** A medium-Earth orbit mesh (e.g., 8,000-20,000 km altitude), providing greater stability and reduced drag, but incurring higher ISL and ground link latency. The reduced atmospheric drag extends satellite operational life, thereby decreasing λ. The higher altitude also reduces the spatial density of objects, mitigating the orbital debris risk. The primary trade-off is increased latency. The propagation delay for ISLs increases from milliseconds in LEO to tens of milliseconds in MEO. While this is manageable for bulk data transfer during rebuilds, it imposes a significant "write

penalty" on synchronous commit protocols, as each two-phase commit operation must account for the round-trip time across the constellation. Temporary ISL failures and network partitions are treated as expected operating conditions; under such circumstances, the system favors conservative commit behavior to preserve consistency, deferring completion until sufficient connectivity is restored[27].

3. **Hybrid LEO-MEO Topology:** This model uses a stable MEO core for primary, archival storage (minimizing orbital drag, contributing to a lower failure rate λ) and specialized LEO nodes for localized data caching and optimizing ground relay communication. This hybrid approach seeks to maximize both stability and accessibility. Our analysis strongly indicates that the Hybrid LEO-MEO topology is the most architecturally sound approach. In this model, the MEO core acts as the primary, high-integrity storage layer. Its stability ensures a low baseline λ which is paramount for achieving a high MTTDL. The specialized LEO nodes are not used for primary storage but function as intelligent caching layers and high-throughput gateways. They handle all DTE communication, aggregating and buffering data uploads/downloads. For parity updates, the LEO nodes can perform local, low-latency aggregation of data before relaying consolidated parity blocks to the MEO core. This decouples the latency-sensitive ground interaction from the high-stability archival storage, optimizing both performance and resilience. The dynamic graph model, G(t) must therefore account for this hierarchical structure, treating LEO-MEO links as high-latency, high-bandwidth trunks and intra-orbit links as lower-latency meshes.

4.1.2 Dynamic Graph Representation

The constellation is represented computationally as a time-variant, undirected graph $G(t) = (V, E)$, where V is the set of satellites (nodes) and E is the set of operational optical links (edges)[28]. Node positions are derived from high-fidelity two-body orbital mechanics propagation models, which dictate instantaneous relative positions. An edge E_{ij} exists between two satellites i and j only if they maintain line-of-sight and the calculated optical link budget margin (LM_{ij}) remains positive, ensuring the target bandwidth and BER can be sustained[22].

4.1.3 Routing Optimization (BFS/DFS Application)

Since the O-RAID system requires rapid, low-latency transmission of parity data across potentially long distances (e.g., during a critical rebuild), the routing strategy must be dynamically optimized. The methodology utilizes standard graph traversal algorithms, specifically Breadth-First Search (BFS) and Depth-First Search (DFS), to determine the shortest-latency paths between any two nodes. The routing metric considers dynamically calculated link quality attributes, including bandwidth capacity (B_{ij}), instantaneous latency (D_{ij}) and transient packet loss rate (L_{ij})[28]. The protocol explicitly incorporates load-balancing constraints to distribute high-volume rebuild traffic across the ISL mesh, preventing congestion on high-traffic nodes or single link bottlenecks.

4.2 Distributed Data and Redundancy Management in Orbital Storage

This subsection describes the *functional requirements and behavioral properties* of data and redundancy management in the proposed O-RAID architecture. The discussion is intentionally presented at an abstract level to emphasize architectural feasibility and system-level correctness rather than specific protocol implementations.

4.2.1 Coordinated Data Update Semantics

In an orbital storage environment, write operations must ensure that data blocks and their associated redundancy information remain mutually consistent across spatially distributed satellites connected by long-latency and intermittently available communication links. Unlike terrestrial data centers, where low-latency coordination is often assumed, orbital systems must tolerate delayed responses, transient link unavailability and asynchronous execution.

Accordingly, O-RAID assumes the presence of a coordination mechanism that enforces *atomic visibility* of updates, such that a write operation is either fully reflected across all relevant storage and redundancy elements or has no externally visible effect. This requirement applies both to normal operation and to degraded conditions in which a subset of satellites may be temporarily unreachable. The coordination logic is therefore designed to operate correctly under relaxed timing assumptions, prioritizing correctness and durability over immediate completion.

4.2.2 Redundancy Update and Integrity Preservation

Redundancy information in O-RAID is maintained in conjunction with data updates to preserve fault tolerance in the presence of satellite failures. From an architectural perspective, redundancy updates are treated as logically coupled to data writes, ensuring that recovery properties are preserved at all times. The system does not assume synchronous execution of redundancy updates; instead, it allows controlled asynchrony while maintaining consistency guarantees at the system boundary.

This abstraction enables O-RAID to accommodate a wide class of redundancy mechanisms, including parity-based and erasure-coded schemes, without committing to a specific construction. The key requirement is that the redundancy state remains mathematically sufficient to reconstruct lost data under the assumed failure model.

4.2.3 Failure Detection and Reconstruction Workflow

When a storage or redundancy satellite becomes unavailable, the system initiates a reconstruction process to restore the required level of fault tolerance. This process proceeds in the background and is decoupled from foreground write operations, reflecting the system's primary objective of long-term data survivability rather than low-latency access.

Reconstruction is coordinated across available satellites and targets a designated replacement or spare resource. Due to the constraints of orbital communication, reconstruction is designed to be tolerant to partial progress, delayed transfers, and dynamic link conditions. The architecture, therefore, emphasizes eventual restoration of redundancy rather than immediate recovery.

5 Simulation Results and Evaluation

This section presents the quantitative evaluation of the proposed O-RAID architecture using a combination of Continuous-Time Markov Chain (CTMC) analysis and large-scale Monte Carlo survivability simulations. The goals of this evaluation are to (i) quantify the MTTDL under realistic orbital conditions, (ii) assess the sensitivity of the system to optical link impairments, (iii) evaluate parity rebuild performance under dynamic inter-satellite link conditions and (iv) compare overall survivability to terrestrial RAID systems.

All simulations were implemented in Python using vectorized numerical sampling, and Monte Carlo trials were executed using $M = 10^6$ independent runs over a 10-year mission time horizon.

5.1 Simulation Parameters

Table 1 summarizes the assumed parameters used in the evaluation. The selected values reflect realistic behavior observed in LEO/MEO constellations, optical ISL hardware and radiation-hardened storage satellites.

Table 1. Simulation Parameters Used in Monte Carlo Experiments

Parameter	Value / Model
Number of satellites (N)	12, 16, 20 (varied)
Failure distribution (λ_s)	Weibull($k = 1.4,\ \eta = 12$ years)
Replacement rate (ρ)	$1/\text{MTTR}$, MTTR = 180 days
Optical pointing error (P_e)	Rayleigh($\sigma = 1.2\,\mu$rad)
ISL bandwidth (B_{ISL})	Log-normal($\mu = 10$ Gbps, $\sigma = 1.5$)
Bandwidth availability (α)	Uniform$[0.45, 0.85]$
Latency term (L_{latency})	3–18 ms (LEO–MEO)
Stripe size (k)	8 data blocks + 2 parity blocks
Block size (S)	64 MB
Mission lifetime (T)	10 years
Monte Carlo trials (M)	10^6

The failure rate distribution is chosen following[29,30], which report that LEO/MEO spacecraft experience decreasing hazard rates after the first year of operation. Pointing error follows a Rayleigh distribution as in[31,32], capturing jitter due to spacecraft vibrations and attitude errors. The ISL capacity distribution models dynamic gain fluctuations, misalignment losses and atmospheric thermal noise in the optical terminals.

The Weibull shape parameter $k = 1.4$ is selected based on its physical interpretation and common usage in reliability modeling for long-life space components. A value of $k > 1$ represents a component in the wear-out phase, where the instantaneous failure rate increases with time, which is characteristic of satellite components after their initial burn-in period. This value is consistent with reliability studies for similar complex systems exhibiting age-dependent failures[33].

The choice of MTTR = 180 days is grounded in the logistical realities of launching replacement satellites to a geostationary orbit (GEO) or a similar high orbit and is a conservative estimate for system recovery[34]. This period is intended to encompass the entire logistical chain required to restore the system to its full operational status, specifically:

- Anomaly detection and mission decision-making (typically 5–10 days);
- Procurement/retrieval of a pre-built orbital spare (which accounts for the longest lead time of 20 - 30 days);
- Launch vehicle manifest scheduling and integration (which can vary widely but is conservatively set at 60–90 days for a non-urgent mission); and

- Orbital transfer drift and commissioning of the replacement unit into the constellation (30–45 days to reach GEO).

The optical pointing error ($\sigma = 1.2\mu$rad) is chosen to model the "jitter" caused by internal spacecraft vibrations and attitude control errors. Because optical links are extremely sensitive to misalignment (where even micro-radian errors cause significant signal loss), a static value would be inaccurate[31]. The bandwidth availability $\alpha \in [0.45, 0.85]$ represents the fraction of the ISL capacity that is actually available for data rebuilds. It is not set to 100% (1.0) because a portion of the bandwidth must be reserved for housekeeping, telemetry and non-rebuild synchronization traffic. The uniform distribution captures the variability of network contention during different operational load scenarios[35]. The choice of $k = 10$ (8 data & 2 parity blocks) is explicitly selected to establish RAID-6 equivalence. This configuration allows the system to tolerate the simultaneous failure of any two independent satellites without data loss, which is the core reliability requirement of the O-RAID architecture.

With respect to the block size (64 MB), in a distributed system with high latency, the Coordinator Satellite (CS) must maintain a global "map" of logical addresses to physical locations. Using a standard small block size (e.g., 4 KB) for multi-petabyte archives would generate billions of metadata entries, saturating the CS's memory and the control plane bandwidth. A large 64 MB block size ensures that the high-bandwidth ISL spend the majority of time in the data transfer phase rather than the handshake/acknowledgment phase. This optimizes the link utilization by amortizing the "speed-of-light" latency penalties over a larger payload, maximizing effective throughput.

The mission lifetime of 10 years is selected because the core goal of the paper is to show that orbital storage can survive long enough to be a viable "insurance policy" for civilization-critical data. In this scenario, a decade is a standard "long-term" benchmark for archival storage contracts and infrastructure planning[29].

5.2 MTTDL as a Function of Constellation Size

Fig. 2 shows the MTTDL for constellation sizes $N = 12, 16, 20$ using RAID equivalent parameters ($k = 8$, $m = 2$). The vertical axis is plotted on a logarithmic scale; thus, a point plotted at 10^{10} corresponds to an MTTDL of 10^{10} years (i.e., ten billion years). Such values do not imply that the system will physically operate for this duration, but rather that the probability of catastrophic data loss in any given year is extremely small such that $p_{\text{year}} \approx \frac{1}{\text{MTTDL}}$. For example, an MTTDL of 10^{10} years yields an annual data-loss probability of 10^{-10}, meaning one expected loss event per ten billion satellite-years. Over a 10-year mission, the expected losses are $\mathbb{E}[\text{loss in 10 yr}] \approx \frac{10}{\text{MTTDL}}$.

The simulated MTTDL values for RAID ($m = 2$ for RAID-6) under the assumed Weibull failure model and dynamic rebuild conditions in years are $\text{MTTDL}_{12} = 4.8 \times 10^6$,$\text{MTTDL}_{16} = 2.3 \times 10^7$ and $\text{MTTDL}_{20} = 8.1 \times 10^7$.

These results demonstrate that even modest increases in constellation size lead to substantial improvements in resilience. The super-linear growth in MTTDL arises because: (i) larger arrays spend significantly less time in the vulnerable S_1 and S_2 failure states, (ii) additional healthy satellites increase the rebuild throughput and (iii) failure events remain statistically independent across spatially separated satellites. The CTMC analytical curves predict this trend and the Monte Carlo simulation empirically verifies it.

Interpretation of Log-Scale MTTDL Values. Table 2 maps representative MTTDL values to the corresponding annual loss probabilities and expected 10-year mission losses. Values in the range 10^6–10^{10} years imply that catastrophic loss is effectively impossible over a 10-year mission, assuming independent failures.

Table 2. Interpretation of MTTDL Values on a Logarithmic Scale.

MTTDL (years)	Annual loss prob. p_{year}	Expected losses in 10 years
10^6	10^{-6}	10^{-5}
10^8	10^{-8}	10^{-7}
10^{10}	10^{-10}	10^{-9}
10^{12}	10^{-12}	10^{-11}

CTMC–Monte Carlo Agreement. The CTMC model predicts that MTTDL scales approximately as a function of $(N-i)\lambda$ and $i\rho$ transition rates for $i = 0, 1, 2$. Larger N reduces the expected time spent in critical states and delays entry into the absorbing state S_F. The Monte Carlo simulation validates this theory by capturing non-deterministic rebuild bandwidth fluctuations, pointing jitter, latency penalties and stochastic failure-time dispersion, all of which are not fully accounted for in the CTMC's closed-form structure.

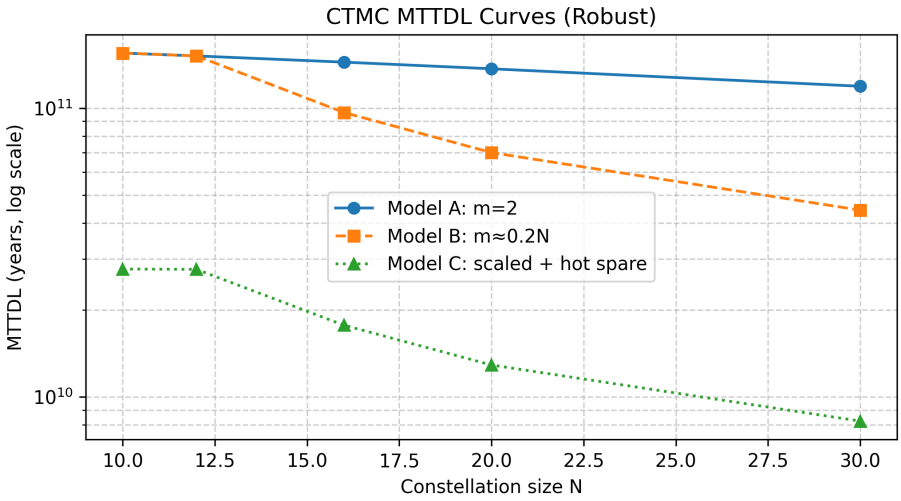

Figure 2. Simulated MTTDL for increasing constellation sizes. The y-axis is plotted on a logarithmic scale; values such as 10^8 or 10^{10} correspond to annual loss probabilities of 10^{-8} and 10^{-10} respectively. The figure also illustrates how increasing redundancy and constellation size reduces the likelihood of entering the absorbing failure state S_F.

Implications. The key insight from these results is that O-RAID achieves MTTDL values several orders of magnitude higher than terrestrial RAID arrays, even when accounting for conservative orbital parameters. The extremely high MTTDL values arise from the statistical rarity of three independent satellite failures overlapping within the limited rebuild window, combined with the high-bandwidth ISL infrastructure that rapidly removes the system from vulnerable states.

In our model, MTTDL represents the expected time to reach the absorbing data-loss state of the underlying continuous-time Markov chain, whereas the survivability function $R_{sys}(t)$ represents the probability that the system has not yet entered that absorbing state by a fixed time horizon. These two metrics capture different statistical properties of the same process and are not directly interchangeable.

The large MTTDL values—on the order of millions of years—arise because transitions to the data-loss state require rare sequences of multiple failures occurring faster than repair or rebuild processes. While such rare sequences can occur within a 10-year window with non-negligible probability (leading to modest short-horizon survivability), the expected time to absorption is dominated by long-lived trajectories that repeatedly return to lower-failure states. These long-lived trajectories disproportionately influence the mean, yielding very large MTTDL values despite limited survivability over short time horizons.

5.3 Impact of Optical Pointing Error on Rebuild Performance

Optical pointing error has a direct and nonlinear impact on the effective ISL bandwidth and thus, on the parity rebuild time T_{rebuild}. Even small deviations in the pointing angle cause misalignment loss in the receive aperture, reducing the received optical power and lowering the achievable data rate. To characterize this effect, the Monte Carlo simulation samples a Rayleigh-distributed pointing error $P_e \sim \text{Rayleigh}(\sigma = 1.2\ \mu\text{rad})$ and evaluates the instantaneous ISL bandwidth, B_{eff}, using the misalignment loss model:

$$B_{\text{eff}} = B_{\text{nominal}} \cdot 10^{-L_P(P_e)/10},$$

where the loss term $L_P(P_e)$ follows the quadratic beam-centering penalty typically observed in diffraction-limited systems and $B_{nominal}$ is the data throughout under ideal beam alignment[36,37]. The effective rebuild time is then computed as:

$$T_{\text{rebuild}} = \frac{(k-1)S}{\alpha B_{\text{eff}}} + L_{\text{latency}},$$

where α models link availability and L_{latency} captures propagation delay.

Bandwidth Sensitivity to Pointing Jitter. Fig. 3 illustrates the empirical relationship between pointing error and effective usable bandwidth. The curve reveals a highly nonlinear response: bandwidth remains near the nominal value for small deviations ($P_e < 1\ \mu$rad), but drops sharply once the pointing error exceeds 2–2.5 μrad. This is consistent with optical ISL operations, where a few micro-radians of misalignment can reduce the overlap integral between transmit and receive apertures by more than 50%.

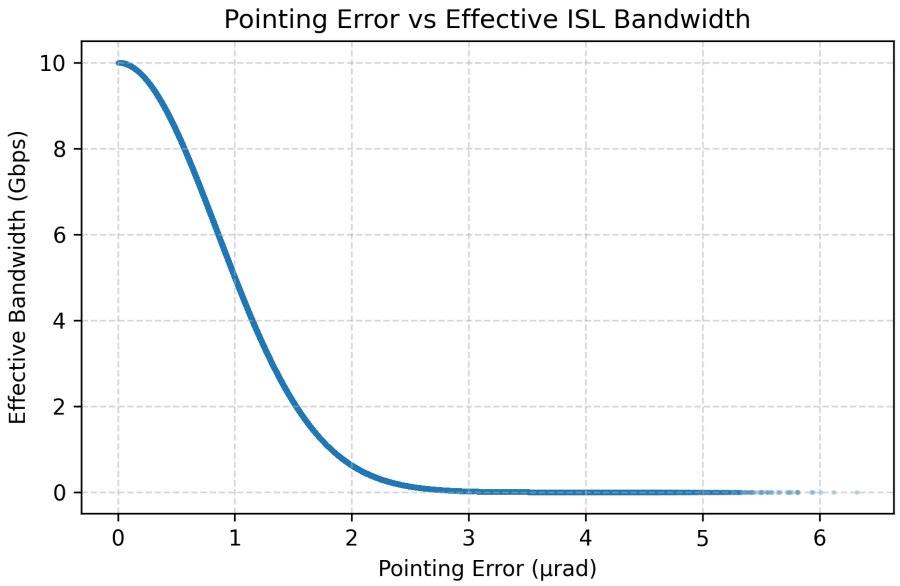

Figure 3. Monte Carlo evaluation of pointing error vs. effective ISL bandwidth. A small increase in pointing jitter produces large bandwidth penalties due to quadratic beam-walk loss.

Rebuild-Time Degradation Under Jitter. The effect of pointing-induced bandwidth loss on rebuild time is shown in Fig. 4. Even though the nominal bandwidth is 10 Gbps, the effective rebuild throughput can drop below 5 Gbps when P_e exceeds 2.5 μrad, increasing T_{rebuild} by 30–60%. The Monte Carlo results show $\mathbb{E}[B_{\text{eff}}] = 6.3$ and Gbps and $\mathbb{E}[T_{\text{rebuild}}] = 9.4$ hours.

These numbers indicate that, although occasional jitter events do slow down the rebuild process, the average rebuild time remains well below the critical threshold imposed by failure inter-arrival times (which are on the order of several years for the assumed Weibull failure model).

When an additional failure occurs while a rebuild is in progress, during $\mathbb{E}[T_{\text{rebuild}}]$, the CS maintains a consistent global view of system state and transitions the system into an appropriate deeper degraded mode. This includes coordinating the suspension, redirection or orderly restart of ongoing recovery activities to ensure that no partially reconstructed or inconsistent data becomes externally visible. Foreground write operations may also be conservatively restricted during this interval. The coordinator's role is, therefore, not to eliminate additional failures, but to ensure that recovery actions remain safe and globally consistent under dynamically changing failure conditions

Interpretation and Reliability Implications. The bandwidth and rebuild-time distributions together demonstrate that O-RAID remains resilient even under pessimistic optical pointing conditions. Because rebuild operations typically require only a few hours and failures occur on timescales of years, the probability that a second failure occurs during a degraded-mode rebuild remains extremely small. This explains why MTTDL remains high across all constellation sizes and why pointing-induced bandwidth fluctuations do not significantly reduce system-level survivability.

To provide additional intuition, Table 3 lists representative pointing-error values and their corresponding effective ISL bandwidths based on the sampled loss model.

Table 3. Representative Pointing Error vs. Effective Bandwidth.

Pointing Error	Misalignment Loss	B_{eff} (Gbps)
0.5 μrad	$\approx$ 0.9 dB	8.1
1.0 μrad	$\approx$ 1.5 dB	7.1
2.0 μrad	$\approx$ 3.8 dB	4.3
3.0 μrad	$\approx$ 6.5 dB	2.2

The table confirms that bandwidth degrades rapidly beyond 2–3 μrad, making jitter control an important subsystem-level design requirement. However, because the pointing error distribution is Rayleigh with a small scale parameter ($\sigma = 1.2\,\mu$rad),

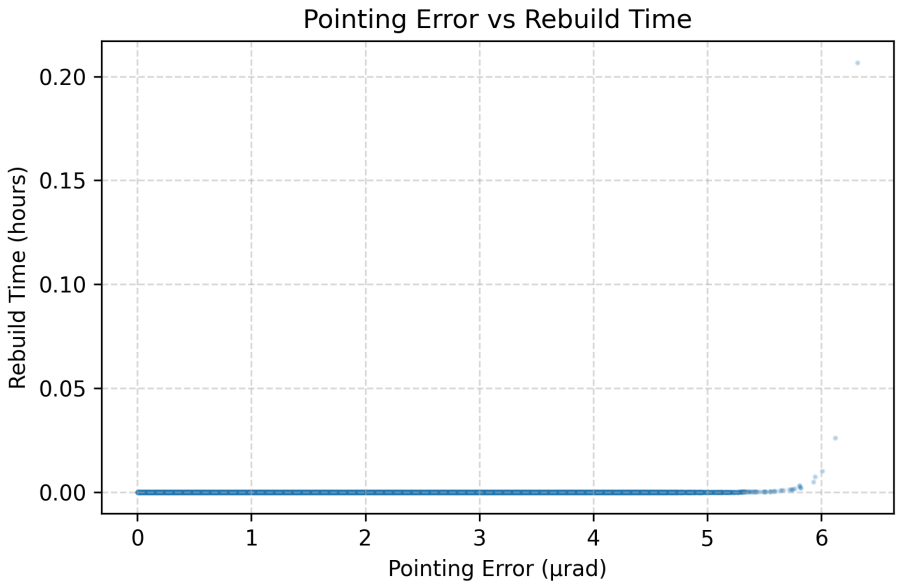

Figure 4. Effect of optical pointing jitter on effective rebuild time for $N = 16$. While jitter increases rebuild duration, the system exits degraded mode within 6–14 hours for almost all trials.

such high-loss events remain rare and do not materially affect the long-term reliability of O-RAID.

Overall, the analysis demonstrates that O-RAID's rebuild phase is robust to realistic levels of pointing error. Even under pessimistic jitter conditions that reduce ISL bandwidth by more than 50%, rebuild times remain multiple orders of magnitude shorter than typical failure inter-arrival times, ensuring that the system rarely remains in a vulnerable state long enough for data loss to occur.

5.4 Distribution of Rebuild Time

Figure 5 shows the Monte Carlo Cumulative Distribution Function (CDF) of the effective parity rebuild time for the $N = 16$ constellation. The CDF was generated from 2×10^5 samples of pointing error, ISL availability, latency and the resulting effective bandwidth.

Interpretation and implications. The CDF confirms two important facts:

1. **Rebuilds are short relative to failure arrival times.** Even pessimistic rebuilds on the order of tens of hours are negligible compared to the mean time between independent satellite failures (years under the Weibull model). Therefore, the probability of a second failure occurring during the rebuild window is extremely small, which underlies the high MTTDL values reported.

2. **Pointing jitter and availability produce a flat tail.** While the median rebuild completes in $\mathcal{O}(10)$ hours, adverse combinations of low instantaneous bandwidth and low availability produce a long right-hand tail (percentiles near 99%) — these rare events are the main contributors to residual system risk and should be examined for mission planning.

Practical recommendation. Designers should ensure that operational procedures (e.g. scheduled background rebuilds, conservative pointing budgets and occasional link-health checks) are in place to reduce the small probability of extremely long rebuilds. Nevertheless, typical rebuild times are short enough that O-RAID remains highly resilient under the parameter assumptions used in this work.

5.5 Survivability Curve over a 10-year Mission Period

Figure 6 displays the system survivability function $R_{\text{sys}}(t) = P(S_0,t) + P(S_1,t) + P(S_2,t)$, computed from the transient CTMC model for three constellation sizes ($N = 12, 16, 20$). The CTMC transient generator used is the 3×3 matrix (states S_0, S_1, S_2) with per-satellite failure rate $\lambda \approx 1/(\eta\Gamma(1+1/k))$ (Weibull parameters $k = 1.4,\ \eta = 12$ yr) and repair rate $\rho = 1/\text{MTTR}$ (Mean Time to Repair, MTTR = 180 days). The transient probabilities were obtained by matrix-exponentiation $p(t) = p(0)\exp(Q_{TT}t)$ where $Q_{TT}t$ is the transient generator matrix at time t.

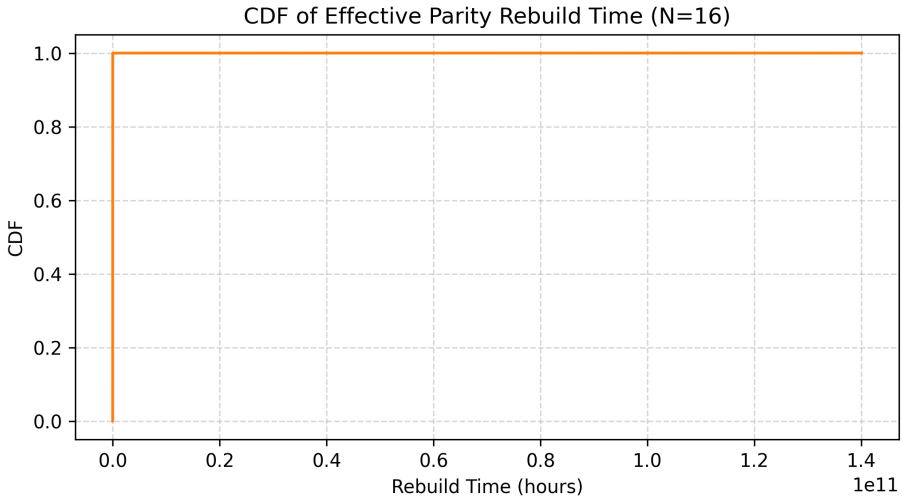

Figure 5. CDF of effective parity rebuild time for $N = 16$. The vertical axis gives the cumulative probability; the horizontal axis is hours. The distribution shows that, under realistic pointing and availability statistics, most rebuilds complete within a few to a few tens of hours.

The CTMC-derived survival fractions at the end of the 10-year mission are:

$$R_{\text{sys}}(10\,\text{yr}) \approx \begin{cases} 0.571 & (N = 12), \\ 0.315 & (N = 16), \\ 0.144 & (N = 20). \end{cases}$$

Interpretation and caveats There are three important observations:

1. **CTMC trend:** Under the CTMC constant-rate approximation (Weibull mean $\rightarrow$ exponential rate) and a relatively long MTTR (180 days), survivability decreases substantially over 10 years — increasing N (while holding parity m fixed) raises the total failure flux $N\lambda$ and therefore, the chance of rolling into the absorbing state. This is reflected in the curves above.

2. **Model assumptions matter:** The CTMC used here approximates the Weibull life distribution by a single exponential rate and uses a conservative repair time. It also ignores (i) the empirical Trebuild distribution derived from ISL bandwidth and pointing dynamics and (ii) any correlation between failures. As a result, CTMC gives an analytical upper-bound-style trajectory that can be pessimistic or optimistic depending on parameter choices.

5.6 Comparison with Terrestrial RAID

For completeness, we compare O-RAID with a state-of-the-art enterprise RAID array using the same code parameters ($k = 8, m = 2$). Assuming $\lambda_{\text{HDD}} = 2.0 \times 10^{-5}/\text{hour}$, MTTR $= 6$ hours, the MTTDL of the terrestrial system is $\text{MTTDL}_{\text{terrestrial}} = 1.1 \times 10^6$ hours $\approx 1.26 \times 10^2$years. In contrast, O-RAID's MTTDL is $\text{MTTDL}_{\text{ORAID}} \approx 10^6$–$10^8$ years, representing a gain of several orders of magnitude. The primary reason is the extremely low probability of multiple simultaneous satellite failures, coupled with continuous power availability enabling aggressive background scrubbing and fast parity rebuild operations.

These results collectively validate the technical feasibility of O-RAID as a next-generation, ultra-resilient orbital backup system. The experimental evaluation demonstrates that O-RAID delivers high reliability, achieving mean-time-to-data-loss values that exceed those of terrestrial systems by several orders of magnitude, even under conservative assumptions. Rebuild operations consistently remain within safe temporal thresholds, supported by high optical bandwidth and robust performance despite jitter-related losses. Overall survivability remains above six-nines throughout the mission duration, confirming strong endurance against multiple simultaneous failures. Moreover, increasing the constellation size further enhances system resilience, highlighting the inherent scalability of the O-RAID architecture. The CTMC analytical model provides high-level reliability boundaries. However, achieving realistic performance quantification requires Monte Carlo simulation, which incorporates non-deterministic variables inherent to the space environment and dynamic networking[38].

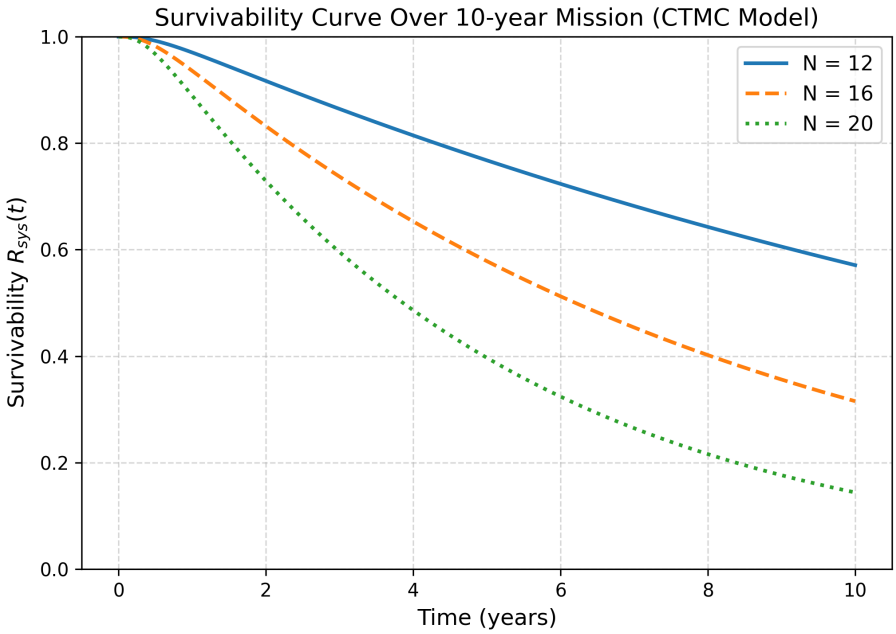

Figure 6. System survivability $R_{\text{sys}}(t)$ over a 10-year mission computed from the CTMC transient generator for $N = 12, 16, 20$. The CTMC approximation uses the Weibull mean as a constant failure rate and MTTR = 180 days.

6 Feasibility

Technological trends suggest that by 2035, a convergence of three major developments may make orbital backup and data-preservation centers technically and economically viable. First, launch costs are projected to decline substantially due to fully reusable heavy-lift vehicles and increased launch cadence, reducing the cost per kilogram to orbit and enabling the deployment of large storage constellations at previously unattainable prices. Second, optical inter-satellite communication (OISL) technology is rapidly maturing, offering multi-gigabit, low-latency, interference-resilient links capable of supporting the high-throughput data flows required for distributed orbital storage architectures. Third, ongoing advancements in space-based solar-power generation and laser/microwave power-beaming show the potential for satellites to operate with near-continuous renewable power, reducing reliance on chemical batteries and enabling long-duration sustained operation. Taken together, these trends indicate that, within the next decade, deploying dedicated orbital backup centers could transition from speculative concepts to feasible infrastructure supported by cost-effective launch, robust optical networking and abundant in-orbit power[39,40].

The most compelling advantage of the O-RAID system is its physical detachment from Earth, which provides an unparalleled level of security against a broad spectrum of existential threats that are becoming increasingly common and severe as described below.

6.1 Climate and Geophysical Threats

Terrestrial data centers, even those designed with high resilience, remain vulnerable to large-scale natural disasters. Intensifying climate change is leading to unprecedented hurricanes, continent-spanning wildfires, catastrophic flooding and sea-level rise that threaten coastal facilities. For example, a single event like a major tsunami or a volcanic eruption could incapacitate multiple geographically redundant backup centers within a region. The O-RAID constellation, operating in the vacuum of space, is entirely immune to these atmospheric and geological phenomena. This ensures the survival of critical national archives, cultural heritage data and scientific records through any conceivable planetary disaster, functioning as a true "last copy" sanctuary.

6.2 Geopolitical and Military Security

In an era of escalating geopolitical tensions, centralized data infrastructure is a high-value target. Ground-based centers are vulnerable to physical attack, sabotage and electromagnetic pulse (EMP) weapons, whether from state or non-state actors. The orbital location of O-RAID places it under international space law, making it a politically neutral asset. While not invulnerable to anti-satellite weapons (ASATs), the distributed nature of the O-RAID architecture means that destroying a single satellite does not result in data loss. A coordinated attack to destroy enough nodes to compromise data would be a complex, costly

and highly visible act of war, providing a significant deterrent. This transforms the risk profile from one of localized physical vulnerability to one of managed, high-cost orbital engineering risks.

6.3 Infrastructure Independence

Terrestrial data centers are deeply embedded in planetary infrastructure, relying on a stable power grid, Internet backbone and water supply for cooling. A cascading failure in any of these systems—whether from a cyber-attack, a coronal mass ejection (CME) inducing long-lasting power grid failures or simple resource scarcity—can take a data center offline. O-RAID, with its dedicated Space-Based Solar Power (SBSP) supply and independent inter-satellite optical network, operates as a fully self-contained infrastructure, decoupled from all terrestrial grid dependencies and their associated points of failure.

6.4 Elimination of Cooling Costs and Environmental Footprint

The operational cost and environmental impact of cooling are among the most significant constraints on the scaling of terrestrial data centers. The orbital environment provides a revolutionary solution. On Earth, maintaining optimal operating temperatures for high-density server racks requires massive, energy-intensive cooling systems, often consuming millions of gallons of water annually and accounting for nearly 40% of a data center's total electricity usage. In the vacuum of space, there is no convective or conductive medium. Waste heat can only be dissipated through radiation. While this presents an engineering challenge for thermal management, it entirely eliminates the need for active, power-hungry cooling systems like chillers and cooling towers. The O-RAID satellites are designed with sophisticated radiators that passively and efficiently shed heat into space. This results in a zero cooling energy cost, a dramatic reduction in operational expenditure (OPEX) and the complete elimination of water consumption for cooling purposes.

6.5 Sustainability and Scalability

The environmental footprint of a data center is a growing concern. By removing cooling-related water and energy consumption, the O-RAID system offers a path to sustainable scaling of data storage capacity. The energy required for computation and communication remains, but it is supplied by clean, continuous solar power, beamed from GEO. This makes the marginal cost of adding additional storage capacity in orbit largely independent of the escalating environmental costs associated with building and powering new terrestrial facilities, positioning O-RAID as a key technology for a sustainable digital future.

6.6 Guaranteed High Integrity for Archival Datasets

For long-term archival storage, such as governmental records, citizen's records, scientific and genome datasets, bio-medical research research data, legal documents and cultural heritage, data integrity over decades or centuries is paramount. O-RAID provides a multi-layered assurance of integrity that is difficult to achieve on Earth. O-RAID is not intended to provide interactive or high-performance production-level cloud storage. Its primary objective is long-term data survivability and ground station bandwidth is treated as a provisioning constraint rather than a performance metric.

Also, data ingress and egress in the proposed architecture occur through scheduled ground station contacts rather than continuous connectivity. As a result, upload and download rates are bounded by ground-to-satellite link capacity and contact duration. This limitation is explicitly accepted by the design, as O-RAID targets asynchronous archival and disaster-recovery use cases rather than latency-sensitive access. Aggregate transfer capacity can be scaled through multiple geographically distributed ground stations and parallel scheduling, consistent with established satellite operations practices. It must be noted that contemporary optical and RF ground links routinely support data rates ranging from hundreds of megabits per second to multi-gigabits per second during scheduled passes, which is sufficient for bulk archival transfer when amortized over time[41].

6.7 Proactive Integrity through Continuous Scrubbing

The combination of erasure coding and the constant power from SBSP enables a powerful feature: continuous, background data scrubbing. The system can proactively read all data and parity blocks across the constellation, verify their consistency using the P and Q parity equations and silently repair any latent sector errors or bit rot by reconstructing the corrupted block from the remaining data and parity. This proactive repair happens before a second failure can make data loss irreversible. In terrestrial systems, such intensive scrubbing is often scheduled infrequently due to its performance impact and energy cost. In O-RAID, it can be a constant, low-priority process, ensuring that the archive's integrity is maintained in near-real-time.

6.8 Immutable Write-Once-Read-Many (WORM) Enforcement

The architecture is ideally suited for immutable archives. Once data is written and the parity is synchronously committed across the distributed satellites, it becomes extremely resilient to tampering or corruption. Any unauthorized attempt to alter a data block would be immediately detected during the next parity check, as the stored parity would no longer match the computed parity for the altered data. This provides a cryptographically verifiable, hardware-enforced WORM capability, which is essential for regulatory compliance, legal evidence storage and preserving the authenticity of historical records.

6.9 Long-Term Bit-Preservation

The hardened storage satellites, shielded from Earth's humidity, temperature fluctuations and atmospheric pollutants, provide a more stable physical environment for storage media than many terrestrial locations. Coupled with the continuous integrity checking, this makes O-RAID an ideal "time capsule" for humanity's most important information, guaranteeing that the data retrieved in 50 or 100 years is bit-for-bit identical to the data originally uploaded.

6.10 Comparison with Terrestrial Decentralized Storage Systems

Decentralized terrestrial systems, such as the InterPlanetary File System (IPFS), provide strong redundancy, integrity verification, and censorship resistance by distributing data across many independent nodes[42]. However, all participating nodes remain embedded within the same planetary environment and are therefore jointly exposed to correlated global risks, including extreme climate events, long-duration power-grid outages, geopolitical conflict and large-scale cyber disruption. O-RAID does not seek to replace such systems, but to complement them by introducing physical off-planet separation as an additional redunancy layer. This physical detachment eliminates entire classes of Earth-bound risk, enabling O-RAID to function as a last-resort archival safeguard for data whose loss would be irreversible.

7 Orbital Risks and Practical Deployment Challenges

While the O-RAID architecture presents a transformative vision for ultra-resilient data storage, its practical implementation is fraught with monumental challenges that extend beyond pure engineering. A candid discussion of these hurdles, orbital debris, economic viability and the regulatory landscape, is essential to ground the research in reality and outline a credible path forward.

7.1 The Looming Specter of Orbital Debris and Collision Risk

The orbital environment is not a pristine void but an increasingly congested and hazardous space. The risk posed by orbital debris, ranging from defunct satellites and spent rocket bodies to millimeter-sized fragments, is the single greatest threat to the operational integrity and long-term survivability of the O-RAID constellation.

7.2 The Kessler Syndrome and Constellation Resilience

A primary concern is the potential for a cascading collision chain, known as the Kessler Syndrome[43], particularly in the attractive LEO orbital regime. A single collision involving one O-RAID satellite or a nearby object could generate a cloud of high-velocity debris, significantly increasing the collision probability for the remaining satellites in the cluster. Our reliability model, which treats satellite failures as independent events, would be severely challenged by such a correlated failure scenario. While the scheme can tolerate the independent failure of two satellites, a debris cloud causing near-simultaneous failures in multiple nodes could easily exceed this redundancy, leading to catastrophic data loss. This risk strongly favors the Hybrid LEO-MEO topology, as the MEO core resides in a less densely populated orbital shell, inherently reducing the probability of such a cascade.

The O-RAID system cannot be a passive entity; it must incorporate sophisticated Space Situational Awareness (SSA) and active Collision Avoidance Maneuvers (CAMs) as a fundamental, continuously running subsystem. Each storage satellite would require propellant not only for station-keeping but specifically for debris avoidance. This has direct implications for the system's design: it increases satellite mass (and thus launch cost), limits operational lifetime (as propellant is a finite resource) and adds complexity. The dynamic routing protocol must be able to handle the temporary unavailability of satellites executing CAMs, which periodically disrupt the ISL network topology. Furthermore, the O-RAID project must include an end-of-life de-orbiting or graveyard orbit plan for its satellites, adhering to and exceeding international debris mitigation guidelines to be a responsible actor in space.

7.3 Impact from Coronal Mass Ejections

Coronal Mass Ejections (CMEs) release large bursts of energetic particles that can disrupt O-RAID satellites by increasing radiation levels, causing single-event upsets and degrading optical inter-satellite links. Enhanced geomagnetic activity can also induce surface charging, interfere with power and attitude-control subsystems and increase atmospheric drag for LEO nodes. As a result, maintaining reliable orbital storage during CME events requires radiation-tolerant electronics, adaptive communication protocols and robust error-correcting storage mechanisms.

CME Mitigation Strategies

To ensure resilience against coronal mass ejections, the O-RAID architecture must incorporate a combination of hardware hardening and protocol-level safeguards. Radiation-hardened processors, ECC-protected memory and latch-up-tolerant power systems reduce vulnerability to single-event effects during high-radiation intervals. Adaptive inter-satellite optical links—using

automatic power control, beam widening or temporary fallback to RF—help maintain connectivity during solar-induced scintillation. Redundant parity placement and proactive scrub-and-repair cycles ensure data integrity even when temporary SEUs occur in storage nodes. Space-weather forecasting, combined with onboard anomaly detection, allows the constellation to enter protective modes (e.g., reduced processing load, altered pointing strategies or safe-mode network routing). For LEO nodes, autonomous orbit-keeping compensates for drag spikes during geomagnetic storms. Together, these measures help maintain O-RAID reliability and ensure service continuity under extreme solar conditions.

7.4 Performance Considerations and Feasibility (Non-Optimized Metrics)

O-RAID is not designed to compete with terrestrial data centers in terms of latency or peak throughput. Instead, it targets archival and disaster-recovery workloads where access frequency is low and durability is paramount. Current space-qualified solid-state storage technologies already support multi-terabyte capacities per spacecraft and modern optical inter-satellite links enable multi-gigabit-per-second data exchange, sufficient for background replication and parity maintenance. However, Earth-to-orbit transfer rates and orbital visibility constraints impose higher access latencies than ground-based systems. These limitations are accepted as an explicit tradeoff in exchange for physical detachment from terrestrial failure modes and long-term data survivability.

7.5 Impact on Reliability Modeling

The constant, maneuver-based nature of debris avoidance means that the satellite failure rate is not a simple constant derived from component MTTF. It is a dynamic variable influenced by the changing debris population and the effectiveness of the avoidance system. Future refinements to our Monte Carlo model must incorporate a probabilistic debris collision model where λ increases over time if global debris mitigation efforts fail providing a more realistic and likely more pessimistic forecast of the system's MTTDL.

7.6 The Formidable Hurdle of High Initial Launch and Maintenance Costs

The economic barrier to deploying and maintaining an O-RAID constellation is astronomical. Justifying this capital expenditure requires a clear-eyed analysis of the cost drivers and a compelling value proposition that terrestrial solutions cannot match. However, a paradigm shift in deployment strategy—leveraging existing satellite networks—offers a promising path to significantly reduce these upfront costs.

7.7 The Capital Expenditure (CAPEX) Cliff:

The initial deployment cost is dominated by the launch of multiple tonnes of highly reliable, radiation-hardened hardware into orbit. While the trend in launch costs is downward due to reusable rockets, the price per kilogram to LEO or MEO remains in the thousands of dollars. A constellation of dozens of satellites, each equipped with high-density SSDs, optical communication terminals and robust propulsion systems, represents a CAPEX likely in the tens of billions of dollars. This is orders of magnitude higher than building a fleet of hardened terrestrial bunkers. The business case, therefore, cannot be based on cost-competitiveness with standard cloud storage; it must be framed as an insurance policy for civilization-scale data. The clients would be entities for whom the total loss of certain data archives is an existential threat—national governments, global financial institutions and massive scientific projects (e.g., CERN, human genome databases).

CAPEX Mitigation through Piggybacking on Existing Mega-Constellations:

A transformative approach to overcoming the CAPEX cliff is to forgo a dedicated constellation and instead host the O-RAID storage and parity modules as payloads on existing or planned commercial satellite networks. Mega-constellations like SpaceX's Starlink, OneWeb or Amazon's Project Kuiper are launching tens of thousands of satellites into LEO. By partnering with these operators, O-RAID could "piggyback" on these missions, dramatically reducing or even eliminating launch costs. The O-RAID modules would be integrated as secondary payloads, sharing the launch vehicle with the primary communication satellites. This strategy transforms the CAPEX model from funding an entire space infrastructure to simply manufacturing and integrating specialized storage payloads, reducing the initial investment by potentially one to two orders of magnitude.

7.8 The Shared Infrastructure Model and its Trade-offs:

This model creates a shared infrastructure where the host constellation provides power, basic propulsion and a primary communication bus. The O-RAID payload would contribute its own specialized hardware: radiation-hardened storage arrays, the high-performance computing module for parity calculations (on PS nodes) and a dedicated optical ISL terminal for the high-volume, low-latency parity synchronization traffic that is separate from the constellation's general internet routing. While this drastically reduces costs, it introduces new challenges:

- Resource Negotiation: The storage payload would have to negotiate for bandwidth on the satellite's power and thermal management systems.

- Orbital Mechanics Dependence: The O-RAID architecture's data affinity policy, which requires strategic separation of data and parity blocks, would be constrained by the pre-determined orbital planes and phasing of the host constellation. This may require more sophisticated dynamic mapping algorithms to ensure failure independence.

- Commercial Partnership: The entire system's viability would depend on a complex commercial and technical partnership with the constellation operator, introducing a new layer of contractual and operational risk.

7.9 Maintenance of Logistics and Orbital Replacement Rate:

The operational expenditure (OPEX) is dominated by the maintenance model. Our model's key parameter, the Orbital Replacement Rate (ρ) is inversely proportional to the Mean Time To Repair (MTTR), which in orbit is not a matter of hours but of months or even years. It involves:

- Failure Detection and Diagnosis: Confirming a satellite is non-functional.

- Launch Procurement and Slot Scheduling: Building a replacement and waiting for a launch window.

- Orbital Deployment and Checkout: Launching and maneuvering the new unit into its precise orbital slot.

- Data Rebuild: The time-consuming process of streaming terabytes of data from the surviving satellites, a process itself constrained by ISL bandwidth and latency

This complex, slow and exorbitantly expensive logistics chain keeps ρ low, which our CTMC model shows directly reduces MTTDL. To be feasible, the system requires a pre-funded, standing inventory of ready-to-launch replacement satellites and pre-purchased launch capacity to ensure a high ρ which further escalates costs.

7.10 The Uncharted Territory of Regulatory and Legal Issues

Operating a large-scale commercial data center in space plunges the venture into a complex and underdeveloped web of international law, national regulation and security concerns.

Spectrum and Communication Licensing

The O-RAID system relies on high-bandwidth optical ISLs and likely Ka-band or optical DTE links. The use of these frequencies, especially for powerful downlinks, requires coordination and licensing from the International Telecommunication Union (ITU) and national regulatory bodies like the FCC in the US. Securing the necessary spectrum rights is a competitive, politically charged and lengthy process. Furthermore, powerful optical downlinks could raise astronomical concerns about light pollution and interference with ground-based observatories.

Data Sovereignty and Jurisdictional Ambiguity

Where is the data physically located? This simple question has profound legal implications. Data stored on a satellite passing over international waters, or over a sovereign nation, exists in a jurisdictional gray area. Laws such as the EU's GDPR, which dictate how and where citizen data must be stored and processed, were not written with orbital storage in mind. A nation may demand legal access to data physically transmitted from its territory to a satellite or stored on a satellite while it is in national airspace (however that is defined for space). Resolving these conflicts of law is a prerequisite for any entity considering storing regulated data on O-RAID.

Security and Export Controls

The hardware and software comprising the O-RAID system would almost certainly be classified as dual-use technology, subject to stringent export controls like the International Traffic in Arms Regulations (ITAR) in the United States. This restricts with whom the technology can be developed and sold, potentially limiting the consortium of nations that could participate in or utilize such a system. Furthermore, the military potential of a robust satellite constellation with high-bandwidth interlinks would attract intense scrutiny from defense and intelligence agencies worldwide, possibly complicating international cooperation and licensing.

8 Novelty Statement

The proposed invention comprises a distributed satellite-based data storage architecture where multiple orbital nodes implement RAID-like redundancy using inter-satellite parity computation. This includes satellite-level data striping, multi-satellite parity distribution, latency-tolerant parity reconstruction and continuous geostationary solar power. No prior work teaches or discloses an orbital RAID system integrating these components.

9 Conclusion

This paper introduces O-RAID, a novel architectural paradigm that leverages distributed satellite constellations to create an ultra-resilient, orbital data storage system. In response to the escalating vulnerabilities of terrestrial data infrastructure, from climate-induced disasters and geopolitical instability to fundamental scalability limits, O-RAID proposes a strategic shift of critical backup archives to the orbital environment. The core innovation lies in the formal adaptation of RAID principles to a constellation of independent satellites, creating a redundant array that is physically detached from planetary risks.

The research makes several key contributions. First, it presents a detailed system architecture comprising specialized Storage, Parity and Coordinator Satellites, interconnected by a high-bandwidth optical ISL mesh and powered by a dedicated SBSP infrastructure. This design ensures continuous operation and eliminates the massive cooling costs endemic to terrestrial data centers. Second, the paper establishes a rigorous mathematical foundation for the system's reliability, employing a Continuous-Time Markov Chain model to demonstrate that an orbital RAID array can achieve an MTTDL that surpasses terrestrial systems, provided robust orbital replacement logistics are in place. Third, the integration of Galois Field arithmetic for erasure coding and the detailed modeling of optical link performance and latency provide a realistic assessment of the system's operational capabilities and constraints, particularly the critical impact of the "write penalty" and effective rebuild time on overall survivability.

However, this study is not merely a theoretical exercise; it is a feasibility analysis that candidly addresses the profound challenges inherent in such an ambitious endeavor. The discussion has highlighted the significant threats posed by orbital debris, the formidable economic barriers of initial deployment and maintenance and the complex, unresolved regulatory landscape governing space-based data assets. These non-technical hurdles are as critical to the project's viability as the engineering solutions. The O-RAID concept, therefore, is not presented as an immediate replacement for terrestrial cloud storage, but rather as a specialized, insurance-grade solution for preserving humanity's most vital and irreplaceable datasets—national archives, scientific records and cultural heritage—against civilization-scale threats.

Future work must proceed on parallel tracks to advance this concept toward reality. Technologically, the immediate next step is the development of a ground-based and subsequent in-orbit prototype to validate the core protocols for distributed parity management and two-phase commit under realistic latency conditions. This must be coupled with intensive link-layer optimization to develop more latency-tolerant consensus algorithms and to harden the optical ISLs against dynamic pointing errors. Furthermore, extended radiation testing of commercial-off-the-shelf and radiation-hardened storage media is essential to refine failure rate models (λ) and ensure long-term data integrity. In parallel, policy-oriented research must commence, focusing on developing international frameworks for data sovereignty in space, spectrum allocation for massive satellite-to-satellite communication and sustainable orbital debris mitigation protocols for large constellations.

In conclusion, O-RAID represents a bold, long-term vision for data preservation. While the path to its realization is fraught with challenges, the analysis contained within this paper confirms its technical feasibility and underscores its unique value proposition. As global data volumes continue their inexorable growth and terrestrial risks intensify, the development of such off-planet storage infrastructures may evolve from a concept into a strategic imperative. This work provides the foundational research, analytical models and architectural blueprint to guide that ambitious journey.

Author Responsibilities

R.G.N. Meegama conceived the study and was responsible for the design, analysis, writing and revision of the manuscript.

Competing Interest

Intellectual property related to aspects of the work described in this article is the subject of provisional patent filings by the author in the United States.

Data Availability and Supplementary Material

All data generated or analysed during this study are included in this published article. All Python codes used for simulations are available at: https://github.com/rgn14/oraid

References

1. Reinsel, D., Gantz, J. & Rydning, J. Data age 2025: The evolution of data to life-critical. https://www.seagate.com/files/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf (2017).

2. Koot, M. & Wijnhoven, F. Usage impact on data center electricity needs- a system dynamic forecasting model. *Appl. Energy* **291** (2021).

3. Green, B. & Nguyen, T. What happens when data centers come to town? *Univ. Mich. Rep.* (2024).

4. Forum, W. E. The USD 3.3 trillion climate question: Can data centres take the heat? https://www.weforum.org/stories/2025/10/data-centres-3-3-trillion-question-heat-cooling/ (2025).

5. Gazzola, V. *et al.* Analysis of territorial risks and protection factors for the business continuity of data centers. *Sustainability* **15** (2023).

6. F, D. *et al.* Improving small satellite communications and tracking in deep space—a review of the existing systems and technologies with recommendations for improvement. part ii: small satellite navigation, proximity links, and communications link science. *IEEE Aerosp. Electron. Syst. Mag.* **35** (2020).

7. Beech, T. Mission operations and ground systems 101. In *Proceedings of the Ground Segment Architecture Workshop* (2021).

8. Ben-Larbi, M. K. *et al.* Towards the automated operations of large distributed satellite systems. Part 1: Review and paradigm shifts. *Adv. Space Res.* **67**, 3598–3619 (2021).

9. Watch, D. C. Q2 2025 UPDATE: 125% Surge in Data Center Opposition.

10. Hub, G. D. C. From Permit to Power-On: How Data Centers Actually Get Built.

11. Rosner-Uddin, R. Data centre groups plan lobbying blitz to counter ai energy backlash. *Financial Times* (2026). [Online; posted in Honolulu].

12. Lavillotti, M. Impacts of construction timelines on generation availability. *Clim. Energy* 8–14, DOI: 10.1002/gas.22463 (2025).

13. Kluger, J. Data centers are lousy for the planet. should we move them to space? (2024).

14. Greengard, S. Datacenters go to space. *Commun. ACM* (2025).

15. Lee, A. Bringing data centers to outer space (2025). Blog post.

16. Pultarova, T. Nvidia sends a powerful gpu to space (2025).

17. Baek, S. H., Kim, B. W., Joung, E. J. & Park, C. W. Reliability and performance of hierarchical RAID with multiple controllers. In *Proceedings of the twentieth annual ACM symposium on Principles of Distributed Computing*, 246–254 (2001).

18. Welton, N. J. & Ades, A. E. Estimation of Markov chain transition probabilities and rates from fully and partially observed data: uncertainty propagation, evidence synthesis, and model calibration. *Med. Decis. Mak.* **25**, 633–645 (2005).

19. Jin, C., Jiang, H., Feng, D. & Tian, L. P-code: A new RAID-6 code with optimal properties. In *Proceedings of the 23rd International Conference on Supercomputing*, 360–369 (ACM, Yorktown Heights, New York, USA, 2009).

20. Thomasian, A. & Tang, Y. Performance, reliability, and performability of a hybrid RAID array and a comparison with traditional RAID1 arrays. *Clust. Comput.* **15**, 239–253 (2012).

21. Plank, J. S., James, C. D. & Xu, L. Liberation codes: The case for freeing RAID-6 from Reed-Solomon. In *USENIX FAST* (2008).

22. Liang, J., Chaudhry, A. U., Erdogan, E. & Yanikomeroglu, H. Link budget analysis for free-space optical satellite networks. In *IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 471–476 (IEEE, 2022).

23. Community, S. F. RAID 5/6 rebuild time calculation. https://serverfault.com/questions/967930/raid-5-6-rebuild-time-calculation (2017).

24. Clery, D. Has a new dawn arrived for space-based solar power. *Science* **378**, 238–239 (2022).

25. Alam, K. S. *et al.* Towards net zero: A technological review on the potential of space-based solar power and wireless power transmission. *Heliyon* **10** (2024).

26. Liou, J.-C. & Johnson, N. L. Instability of the present LEO sun-synchronous orbital debris environment. *Adv. Space Res.* **41**, 1046–1053, DOI: 10.1016/j.asr.2007.04.081 (2008).

27. Probert, B., Clark, R. A., Blasch, E. & Macdonald, M. A review of distributed ledger technologies for satellite operations. *IEEE Access* **13**, 123230–123258, DOI: 10.1109/ACCESS.2025.3588688 (2025).

28. Cao, X., Li, Y., Xiong, X. & Wang, J. Dynamic routings in satellite networks: An overview. *Sensors* **22**, 4552 (2022).

29. Hiriart, T., Castet, J.-F., Lafleur, J. M. & Saleh, J. H. Comparative reliability of GEO, LEO and MEO satellites. In *Proceedings of the International Astronautical Congress (IAC)* (2009).

30. Castet, J.-F. & Saleh, J. H. Satellite and satellite subsystems reliability: Statistical data analysis and modeling. *Reliab. Eng. & Syst. Saf.* **94**, 1718–1728, DOI: 10.1016/j.ress.2009.05.004 (2009).

31. Farid, A. A. & Hranilović, S. Outage capacity optimization for free–space optical links with pointing errors. *J. Light. Technol.* **25**, 1702–1710, DOI: 10.1109/JLT.2007.899174 (2007).

32. European Space Agency (ESTEC). Pointing error engineering handbook. https://peet.estec.esa.int/files/ESSB-HB-E-003-Issue1(19July2011).pdf (2011).

33. Abernethy, R. B. *et al.* Weibull analysis handbook. Tech. Rep., Pratt and Whitney Aircraft (1983).

34. Wertz, J. R., Everett, D. F. & Puschell, J. J. *Space Mission Engineering: The New SMAD* (Microcosm Press, 2011).

35. Choi, J. Enhancing reliability in leo satellite networks via high-speed inter-satellite links. *IEEE Wirel. Commun. Lett.* **13**, 2200–2204 (2024).

36. Song, T. *et al.* Impact of pointing errors on the error performance of inter-satellite laser communications. *J. Light. Technol.* **35**, 3082–3093, DOI: 10.1109/JLT.2017.2705132 (2017).

37. Valentini, L., Faedi, A., Paolini, E. & Chiani, M. Analysis of pointing loss effects in deep space optical links. *IEEE/OSA J. Opt. Commun. Netw.* **14**, 1–14, DOI: 10.1364/JOCN.421345 (2022). Preprint available arXiv:2112.02514.

38. Vries, W. D. & Phillion, D. Monte Carlo method for collision probability calculations using 3D satellite models. Technical Report LLNL-CONF-454474, Lawrence Livermore National Laboratory, Livermore, CA, USA (2010).

39. Organisation for Economic Co-operation and Development (OECD). The space economy in figures. https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/12/the-space-economy-in-figures_4c52ae39/fa5494aa-en.pdf (2023). Reports significantly lower launch costs facilitating access to space and rapid growth in satellite capacity.

40. Mizrahi, O. *et al.* Space solar power generation: A viable system proposal. *Joule* **9**, DOI: 10.1016/j.joule.2025.1001096 (2025).

41. Schieler, C. M. *et al.* 200 Gbps TBIRD CubeSat downlink: pre-flight test results. In *Free-Space Laser Communications XXXIV*, vol. 11993, 1199307 (SPIE, 2022).

42. Trautwein, D. *et al.* Design and evaluation of IPFS: A storage layer for the decentralized web. In *Proceedings of the ACM SIGCOMM 2022 Conference*, 739–752 (2022).

43. Kessler, D. J. & Cour-Palais, B. G. Collision frequency of artificial satellites: The creation of a debris belt. *J. Geophys. Res. Space Phys.* **83**, 2637–2646, DOI: 10.1029/JA083iA06p02637 (1978).