



TECHNICAL FRAMEWORK

Technical Framework for the Resilient Evolution of the Global Telecommunications System

Project This Station – Earth

Giuseppe Carilli


I'n'I / Lion of Judah

Final consolidated editorial edition FR1

Completed editorial version, before manual corrections and feedback

Research architecture · Reproducible examples · Defined limits

English · 6 October 2026 · 25 pages



Abstract

The Global Telecommunications System is understood here as a federation of networks, energy systems, computing, timing, maintenance and governance. Its resilience is measured by continuity of defined services during disturbance and by evidence of recovery, rather than by nominal peak speed or duplicated equipment alone.

This paper develops the systems discipline of Project This Station – Earth into a reference framework for terrestrial and non-terrestrial communications. The proposed architecture joins source-aware observability, complete resource and energy accounts, qualified substitution, a traceable digital twin and bounded orchestration. It uses the 2005 procurement-with-substitution model as a planning foundation while adding capacity, timing, security and physically executable recovery constraints.

Worked examples explain free-space attenuation, link margin, ideal channel capacity, latency, usable stored energy and mixed-route allocation. A two-route case distinguishes normal 10 Gbit/s delivery from a guaranteed 6 Gbit/s essential service after either route fails; it does not claim full-rate survivability. A practical testbed, failure scenarios, metrics and a conditional thirty-six-month roadmap define what evidence would justify scaling.

The framework complements established ITU, 3GPP, IETF and NIST work; it is not a new telecommunications standard or a report of a deployed global system. White Flag contributes an example of reproducible signal auditing, not an RF shielding technology. This is the completed editorial version before manual author corrections and feedback.

Keywords: service continuity; network resilience; multi-carrier infrastructure; NTN; digital twin; resource substitution; energy autonomy; evidence provenance.

Contents

01 Mission: specify the service that must survive	4
02 Reference architecture: a federation of accountable layers	5
03 Interfaces and service contracts	6
04 Observability, environment and space weather	7
05 Link budget: received power and margin are different	8
06 Capacity, latency and the limits of substitution	9
07 Energy autonomy: usable reserves and complete loads	10
08 Time, synchronisation and holdover	11
09 Qualified substitution: beyond a commercial label	12
10 Worked routing case: cost and N–1 continuity	13
11 Digital twin: reconstruct decisions, expose uncertainty	14
12 Cyber resilience and bounded authority	15
13 Failures, safe states and recovered quality	16
14 RF compatibility and reproducible signal audits	17
15 A practical testbed and experimental sequence	18
16 Metrics that measure delivered continuity	19
17 Economics, sustainability and maintenance	20
18 Roadmap, governance and progression gates	21
19 Glossary and formula interpretation	22
20 Conclusion, shared annex and author attribution	23

Abstract	2
References	24–25

01 Mission: specify the service that must survive

Telecommunications resilience begins with a service contract. Public warning, emergency coordination, local industrial control, scientific-data access and ordinary entertainment traffic have different tolerances for delay, loss and interruption. A nominal network capacity does not express those differences.

Define the users or functions to protect, minimum quality, duration, geography, failure scenarios and conditions for returning to normal operation. Local control may require low latency that an orbital route cannot provide; scientific file transfer may tolerate delay but require intact data. A qualified substitute must preserve the relevant contract. [T01, T15]

The project's continuity-of-life principle becomes continuity of critical functions and knowledge. Configuration backups, keys, repair procedures and calibration records are part of the infrastructure: losing them can prevent recovery even when replacement hardware exists.

Service property	Requirement to declare
Capacity and reach	Useful throughput, coverage and protected users.
Time	Delay, jitter, outage duration and recovery deadline.
Integrity and access	Security, identity and acceptable data loss.
Dependency	Power, clocks, cooling, control and supply chain.
Degradation	What can be reduced while essential service persists.

The architecture should increase continuity without assuming that every generation of mobile technology solves the same problem. ITU's IMT-2030 framework sets future development objectives; it does not certify this proposal or remove deployment-specific engineering. [T03]

02 Reference architecture: a federation of accountable layers

A global system need not be one homogeneous network. Federated domains can share verifiable interfaces while retaining bounded local operation during isolation. The reference architecture has six interacting layers.

Layer	Responsibility
Environmental observation	Spectrum, propagation, weather, space weather and data quality.
Energy, thermal and maintenance	Supply, conversion, reserves, cooling and repair capacity.
Multi-domain transport	Fibre, terrestrial radio, satellites/NTN and applicable optical links.
Services and distributed computing	Core, edge, identity, emergency functions and local storage.
Twin and orchestration	State estimates, scenarios, qualified options and explanations.
Governance and execution policy	Responsibility, authority, operating limits and review.

Reserves and provenance cross the layers. A new radio should not destroy state traceability; changing power supply should not silently change reserve policy. The optimiser should not acquire command authority merely because it can recommend a route.

This architecture transfers engineering discipline from the companion project. It does not require Mercury industry, planetary migration or an audio-induced physiological effect. [T01]

3GPP's NTN work provides an established standardisation context for integrating certain non-terrestrial access modes. That does not make every satellite, high-altitude platform or deep-space link interchangeable. Terminal capabilities, visibility, spectrum, latency and service interfaces must still be qualified. The federation should expose those limits so domains can evolve without hiding incompatible dependencies. [T06]

03 Interfaces and service contracts

A useful interface carries meaning as well as a number. “Available capacity” should identify the service boundary, time window, uncertainty and whether it depends on a common power or transport route. An inventory should distinguish installed, operational, usable and protected amounts.

Object	Minimum fields
Telemetry	Variable, unit, clock, location, quality, uncertainty and version.
Service	Priority, demand, minimum capacity, delay/loss tolerance and duration.
Resource	Usable amount, quality, location, reserve and replacement options.
Link	Medium, band, margin, throughput, latency and common dependencies.
Decision	Feasible alternatives, active constraints, authority and expected consequences.

Define how fresh telemetry must be and what happens when it is absent. A stale battery estimate or clock reading should not be presented as a current measurement. Use bounded uncertainty or a conservative operating state until observation recovers.

Service admission also needs a degradation policy. Traffic above the protected contract may be rejected or delayed when a reserve threshold is crossed. That decision should be understandable to operators and users without exposing implementation details that do not help them act.

Interoperability tests must include actual transitions, not only document compatibility. Equipment or route replacement may change timing, packet size, encryption, authentication, addressing or available management paths. Successful restoration means the contracted service works again, not merely that an interface reports “up”.

04 Observability, environment and space weather

The observation layer preserves raw telemetry, screening and derived estimates separately. Packet loss, bit errors, latency, RF spectrum, energy, temperature and clock quality are different variables. Missing records must not be converted into zero loss or perfect availability.

For every measurement record calibration, bandwidth or aggregation, instrument floor, time source and firmware/software version. Correlate network anomalies with independent weather or space-weather records only after clocks, coverage and processing changes have been checked. NOAA documents several communication and navigation impacts of space weather; those mechanisms are context, not proof of cause for a particular outage. [T09]

Investigation stage	Question
Integrity	Is the anomaly real data, a gap, a clock error or a version change?
Corroboration	Do independent sensors or endpoints observe it?
Mechanism	Is propagation, interference, software, power or traffic consistent with the effect?
Counterfactual	What would have happened without the proposed intervention?
Decision	Are the action, stopping rule and authority defined?

A fibre cut and an ionospheric disturbance require different responses. Optical links, terrestrial microwave and satellite paths also have different atmospheric and geometry dependencies. Earth-space designs should use applicable propagation models, including rain and other losses beyond free-space attenuation. [T05]

The first observatory deliverable is a frozen, reproducible baseline with coverage and uncertainty. It is more valuable than an unexplained dashboard correlation because another analyst can challenge both data and interpretation.

05 Link budget: received power and margin are different

A link budget adds gains and subtracts physical losses on a logarithmic scale. Transmit power and received power may be expressed in dBm; antenna gains and losses are in dB. Free-space attenuation is a geometric model, not a complete description of weather, interference or hardware. [T04, T05]

$$P_r = P_t + G_t + G_r - L_{fs} - L_{atm} - L_{sys}.$$

$$L_{fs} = 20 \log_{10}(4\pi d/\lambda); \lambda = c/f.$$

$$\text{Available margin} = P_r - P_{\text{required}}.$$

P_t and P_r are transmitted and received power; G_t/G_r are antenna gains. L_{fs} , L_{atm} and L_{sys} are free-space, atmospheric and system losses. Distance is d , wavelength λ , frequency f and vacuum light speed c . P_{required} is the receiver's specified minimum power.

A planning reserve is compared with the available margin. It is not an extra physical loss unless a specific loss mechanism is being modelled. This distinction corrects the former report's expression that subtracted reserve directly from received power.

Illustration: at 10 GHz over 10 km, L_{fs} is 132.45 dB. With 20 dBm transmit power, 30 dBi antennas at both ends and 4 dB combined additional losses, $P_r = -56.45$ dBm. If the receiver requires -70 dBm for the declared mode, available margin is 13.55 dB. A 10 dB planning reserve leaves 3.55 dB beyond that target.

The receiver requirement must specify bandwidth, modulation, coding, error target and implementation. Rain, pointing, polarisation, ageing and uncertainty determine whether the chosen reserve is adequate. These artificial values illustrate bookkeeping; they do not qualify a real site, transmitter or licensed deployment.

06 Capacity, latency and the limits of substitution

For an ideal band-limited additive-white-Gaussian-noise channel, Shannon's relation is:

$$C = B \log_2(1 + \text{SNR}).$$

C is ideal capacity in bit/s; B is bandwidth in Hz; SNR is a linear power ratio, not a dB value. At 20 MHz and 10 dB SNR, the ratio is 10 and the ideal bound is approximately 69.19 Mbit/s. It is a separate illustrative example from the link budget, not a derived performance claim for that link. [T07]

Delivered throughput also depends on framing, signalling, scheduling, coding choices, retransmissions, traffic and outage time. If overhead fraction is assumed to be 15%, $69.19 \times 0.85 \approx 58.81$ Mbit/s is only an overhead-adjusted illustration; it is still not a measured service rate.

Propagation delay sets a physical lower bound. A nominal 384,400 km Earth-Moon distance corresponds to about 1.282 seconds one way in vacuum before processing or routing. Remote closed-loop control must therefore be redesigned for local autonomy rather than promise terrestrial real-time response.

When links are intermittent or long-delay, store-carry-forward mechanisms such as Bundle Protocol Version 7 offer a relevant architecture. They do not make delayed data suitable for every low-latency function. [T12]

A route can substitute for another only if it meets that service's quality and timing. A high-throughput satellite path may preserve file delivery while failing a local control deadline. Aggregate capacity is necessary but insufficient evidence of continuity.

07 Energy autonomy: usable reserves and complete loads

A network node consumes energy through radio, transport, computing, storage, security, timing and thermal management. Batteries and flywheels move previously supplied energy through time; they do not create it. Autonomy must include conversion losses, auxiliary demand, ageing and protected reserve. [T01, T14]

$T_{aut} \approx E_{available} / P_{essential}$, under a constant-load approximation.

Eavailable is usable energy and Pessential includes all essential loads. Expressing them in kWh and kW gives Taut in hours.

Take a 120 kWh rated energy quantity. Holding 15% as unavailable/protected leaves 102 kWh. If essential electronics consume 6 kW and cooling 1 kW, ideal autonomy is $102/7 = 14.57$ hours. Further conversion losses or ageing shorten it. If traffic raises electronics demand to 9 kW with the same cooling, the result is 10.2 hours.

Boundary	What to record
Supply	Grid/local generation, fuel, variability and resupply.
Conversion	Efficiency by load and temperature.
Storage	Usable energy, power rating, state, standby loss and reserve.
Loads	Communication, computing, timing, security and auxiliaries.
Thermal	Ambient conditions, cooling power and derating.

An autonomy estimate becomes operationally useful only with a traffic profile and degradation policy. Decide which nonessential functions are reduced before reserves are exhausted. Shared cooling or power can defeat otherwise redundant links.

Orbital or lunar equipment must conduct heat to radiators because vacuum does not provide convective cooling. Terrestrial equipment faces different ambient constraints but the same obligation to account for thermal dependencies. [T13]

08 Time, synchronisation and holdover

A route can carry packets while failing a required timing service. Synchronisation quality should be a resource with a measurable error budget, source diversity and declared holdover duration. ITU-T G.8271 provides the relevant time/phase framework; the exact profile and equipment limits must be chosen for the intended service. [T08]

A simplified holdover estimate illustrates the distinction between frequency error and time error:

Accumulated time error $\approx$ fractional frequency error $\times$ elapsed time.

At 0.1 parts per million, fractional error is 1×10^{-7} . After one hour without correction, the accumulated component is about 360 microseconds. Initial offset, temperature drift, ageing and noise add further error; this is not a guaranteed oscillator specification.

Record source, timestamp offset, uncertainty, disciplining mode, temperature and source-loss alarms. Two receivers do not offer true diversity if both depend on the same vulnerable external timing source. Packet-path asymmetry may also bias a time-transfer estimate.

Local operation during isolation must state which services tolerate degraded time and when they should stop or change mode. Alarm ordering, data provenance and cryptographic validity may be affected before ordinary throughput appears degraded.

A recovery test should remove the preferred timing source, measure error growth, activate a qualified alternative and verify restored quality. The outcome is service-specific. A recovered clock icon or a successful link handshake is not sufficient evidence that the network again meets its time and phase limits.

09 Qualified substitution: beyond a commercial label

The 2005 model combines offers and transformations to satisfy requirements at minimum cost. In networking, candidates include alternative transport, compatible equipment, local edge functions, power sources and repair processes. The extension must preserve both the resource account and the service contract. [T02]

Candidate substitution	Qualification needed
Fibre to radio/NTN	Coverage, delay, throughput, visibility, spectrum and security.
Equipment replacement	Interfaces, timing, capacity, software and maintenance.
Cloud to local edge	Data consistency, identity, storage and restricted operation.
Grid to local reserve	Starting conditions, usable energy, cooling and fuel.
Repair over replacement	Time, tools, spares and recovered quality.

Initial inventories, delivery bundles and transformation costs form a planning basis. Capacity, location, energy, deadline, quality and protected reserves are added constraints. Discrete equipment choices and continuous traffic flows must not be confused.

Final balances are not always an executable sequence. A battery recharge plan may depend on a generator that cannot start without that battery; a remote recovery command may require credentials or a route lost in the same failure. Time-indexed states and start conditions expose these dependencies.

A declared substitute is unavailable if it fails even one binding service requirement, regardless of its attractive cost. Feasibility is established through qualification and executable recovery tests, not by a solver's output alone. Procurement history and actual transformation evidence remain linked to each candidate.

10 Worked routing case: cost and N–1 continuity

Let route A provide 8 Gbit/s and route B 6 Gbit/s, with separate physical, power and control dependencies assumed and later tested. Normal demand is 10 Gbit/s, including a 6 Gbit/s essential subset. Artificial accounting prices are 8 CU and 5.25 CU per allocated Gbit/s for the same planning interval; they are not market prices.

Minimise $8x_A + 5.25x_B$.
 $x_A + x_B \geq 10$; $0 \leq x_A \leq 8$; $0 \leq x_B \leq 6$.

The nominal optimum is $x_A = 4$ and $x_B = 6$, costing 63.5 CU. Four Gbit/s remains unallocated on A, of which at least 2 Gbit/s is marked as protected headroom under the example policy. Headroom does not by itself demonstrate redundancy.

Scenario	Surviving available capacity	Contract outcome
Nominal	14 Gbit/s total	10 Gbit/s demand can be served.
A fails	B: 6 Gbit/s	Essential 6 survives; nonessential 4 is dropped.
B fails	A: 8 Gbit/s	Essential 6 survives after reallocation.
Shared power/control fails	Possibly neither	Independence assumption is invalid.

This is an N–1 capacity argument for the 6 Gbit/s essential service only. Full 10 Gbit/s after either route fails is infeasible with these capacities. Switching time, latency, packet loss, power and equipment configuration still require verification. If the service cannot tolerate the transition, capacity feasibility alone is insufficient. The example improves the former report by separating cheap nominal allocation, reserve policy and scenario-based survivability.

11 Digital twin: reconstruct decisions, expose uncertainty

A telecommunications twin represents services and dependencies as well as topology. Its state includes links, nodes, resources, clocks, software versions, configuration and maintenance. It should reconstruct a past recommendation from archived inputs and distinguish observed state from assumed scenario conditions.

Model layer	Required behaviour
Observation	Preserve provenance, gaps and freshness.
State estimation	Propagate uncertainty and identify unobservable quantities.
Scenario	Introduce failures, traffic, timing and supply windows explicitly.
Optimisation	Respect hard constraints and explain infeasibility.
Recommendation	Record binding limits, alternatives and expected impact.

Validation begins with dimensional checks, conservation and known-answer examples, including the routing case. Then replay historical events and compare predicted versus observed service outcomes. An attractive topology display cannot substitute for these tests.

A model should say when it cannot identify a cause or produce a feasible plan. Missing measurements and unqualified substitutes must not be replaced by silently favourable assumptions. Robust scenario analysis can protect the essential contract across specified disturbances, but it cannot guarantee safety against every unmodelled event.

A proposed architecture is initially a simulation and decision-support system. It becomes an operational twin only after its live interfaces, data quality and predictive adequacy are demonstrated. Versioned runs and decision explanations should remain portable if implementation products change.

12 Cyber resilience and bounded authority

Network resilience includes anticipation, resistance, recovery and adaptation under compromise. NIST’s cyber-resilient systems guidance supports treating those capabilities throughout the lifecycle. Zero-trust architecture provides a complementary principle: network location alone does not confer trust in access to a resource. [T10, T11]

Separate observation, modelling, recommendation and command execution. An optimiser may produce a feasible plan without being entitled to apply it. Execution follows approved identities, policies, bounds and rollback or safe-state procedures. Low-impact reversible adjustments can use preauthorised automation; actions with wider effects need the appropriate operator authority.

Control requirement	Evidence to retain
Identity and access	Approved role, authentication and least necessary capability.
Configuration integrity	Version, signed/verified provenance where applicable and change record.
Isolation	Tested domain boundaries and minimum local operation.
Recovery authority	Credentials and procedures available during external isolation.
Decision record	Actor, recommendation, active limits, action and outcome.

A central controller, software image or key service can be a common failure mode. Local continuity requires a realistic way to recover without the normal remote management path. Offline restoration knowledge and protected backups therefore belong in the service architecture.

These are design requirements informed by primary guidance, not a claim that this system has passed a security assessment. Tests should demonstrate recovery under bounded compromise scenarios while preserving the essential service and evidence needed to understand what occurred.

13 Failures, safe states and recovered quality

Failover moves service to an alternative; recovery restores the contracted quality. They are different milestones. A path may become reachable quickly while remaining too slow, poorly timed or unable to sustain demand.

Fault scenario	Decisive recovery evidence
Fibre cut	Physical route independence, switching time and restored service.
Regional power loss	Measured reserve drawdown and useful autonomy.
Common software defect	Limited deployment scope and independent restoration.
Timing-source loss	Error growth, alternative source and recovered limits.
Interference/congestion	Qualified reallocation with measured margin and loss.
Compromised control	Isolation, trusted recovery authority and essential operation.
Missing spare	Qualified repair/substitution and credible restoration time.

For independent parallel components, $A_{\text{parallel}} = 1 - \prod(1 - A_i)$. If both links have availability 0.99, the ideal independent result is 0.9999. That figure is overstated if a shared dependency can disable both. Report common causes and maintenance alongside any availability estimate.

A safe state limits further harm: freeze changes, isolate a domain, reduce optional loads or keep emergency communication. Rollback restores an earlier state, which may be impossible after data loss, time discontinuity or persistent physical changes.

Introduce representative faults in a controlled testbed and repeat the recovery sequence. Record detection, switching, useful-service restoration and return to normal configuration separately. The test boundary and failure severity must accompany the result, so a local success is not presented as proof of global survivability.

WHITE FLAG · REPRODUCIBILITY


14 RF compatibility and reproducible signal audits

White Flag

This Station Version

3:56.5 · Mono AAC-LC · 44.1 kHz · Original recording

 **Open the embedded song**
Use the paperclip or the PDF Attachments panel.



Electromagnetic compatibility concerns emissions, susceptibility and the behaviour of equipment in a specified environment. An RF field measurement needs frequency, bandwidth, antenna/probe geometry, calibration, time averaging, instrument floor and uncertainty. Receiver performance and human exposure are different assessment tasks.

The White Flag audit provides a transferable method: preserve the asset, identify decoding and transformations, define the metric and denominator, test processing sensitivity and limit interpretation to the measured domain. Its 8.05% envelope-band share is a property of a recording, not RF attenuation or EEG activity. [T01]

Asset	Reproducibility record
IQ capture	Sample format, rate, centre frequency, gains, clock and calibration.
Packet/protocol trace	Capture point, timestamp source, filters and anonymisation rules.
Spectrum record	Resolution bandwidth, detector, averaging and measurement geometry.
Firmware/configuration	Exact version, checksum, dependency and test condition.

Freeze endpoints before comparison. Otherwise a reported improvement may depend on chosen filters, windows, favourable intervals or discarded observations. Hardware controls help distinguish pickup or instrument changes from a physical effect.

No result in this framework establishes that audio neutralises 5G/6G emissions. RF mitigation must be evaluated through its own specified electromagnetic mechanism and instruments. The shared annex stores the audio as a reproducibility example and the 2005 paper as a planning source, with their original attribution.

15 A practical testbed and experimental sequence

A credible first testbed uses at least two sites or independently instrumented domains, a primary route, a qualified alternative, local energy/timing records and a controllable traffic generator. Begin in a contained environment with realistic service contracts; add complexity after known-answer cases pass.

Stage	Output
Baseline	Fixed traffic, configuration, calibration and source versions.
Component checks	Actual usable capacity, delay, power, clocks and thermal envelope.
Fault injection	Declared fibre, power, timing, software and management-path interruptions.
Recovery	Switching sequence, essential service and restored quality.
Independent replay	Archived inputs reproduce the conclusions.

Use repeated matched runs for the baseline and candidate policy. Keep traffic, service priorities and environmental windows comparable; record unavoidable differences. Include runs where the preferred alternative is unavailable, so the system’s infeasibility report and degradation policy can be assessed.

Prespecify useful effect sizes and acceptance limits. A small apparent gain below measurement uncertainty does not justify a new operational policy. Likewise, a few minutes of successful service cannot establish a month of autonomy.

A readiness review should confirm instrumentation, command limits, fault containment and recovery authority. The programme delivers methods and test evidence; it does not promise a deployed planetary network within three years. Scaling requires a model and tests of the new dependencies introduced at each stage.

16 Metrics that measure delivered continuity

The best metric describes a service delivered across a defined boundary and interval. Peak throughput, component uptime and reserve nameplate ratings are useful but incomplete. Report distributions and adverse scenarios where averages hide unacceptable interruption.

Dimension	Required metric
Essential continuity	Time meeting the protected contract, users/functions served and degradation.
Link quality	Useful throughput, delay, jitter, packet loss and margin by condition.
Energy	kWh per delivered service, usable reserve and autonomy under fault.
Recovery	Detection, switching, restored quality and repeatability.
Resources	Qualified spare stock, delivery time, substitute availability and consumption.
Evidence	Coverage, freshness, uncertainty and reproducible decisions.

One availability percentage should not merge maintenance, unobserved intervals and different services silently. Report denominator, service window and what counts as failure. A loss of telemetry may make availability unknown rather than excellent.

Acceptance requires an evidence package: physical topology, independence assumptions, link budgets, energy states, timing limits, configuration versions, raw/derived records and decision history. Explain negative results and binding constraints, not only favourable runs.

Capacity should be counted at the useful service interface. Overhead-adjusted calculations remain illustrations until measured against traffic and endpoints. When a failover restores reachability but not the specified delay or timing, it has not restored the complete service contract. That distinction keeps resilient design focused on users and functions.

17 Economics, sustainability and maintenance

Evaluate delivered service over its lifecycle. Capital, energy, maintenance, spare inventory, supply delays, replacement, downtime and retirement are part of a resilient-network cost boundary. A cheap route that fails the essential contract is outside the feasible comparison, not an inexpensive solution.

Use scenario costs or quotations dated for a specific study. The routing example’s CU prices are deliberately artificial; the framework contains no tariff forecast or universal technology ranking. Generation and storage costs should also remain distinct. [T14]

Trade-off	Question for the model
Spare capacity	Does reserve protect a defined scenario or remain idle nominal headroom?
Local energy	What service duration is gained after full auxiliary accounting?
Supplier diversity	Are compatibility and independent recovery actually qualified?
Higher redundancy	Does it remove a common cause or add correlated maintenance?
Repairable design	Are tools, skills, parts and quality checks locally available?

Sustainability includes energy consumed during operation, embodied materials, consumables, repairability and end-of-life burdens. Avoid a single per-bit energy metric when it rewards dropping low-rate but essential users or excludes cooling and standby operation.

Compare baseline and candidate under the same service contract and fault set. Report uncertainty and sensitivity to utilisation, energy price, component life and lead time. The decision output is a transparent trade-off within the feasible set, not a promised financial return. Evidence that a simpler design restores service reliably may outweigh a nominal efficiency gain.

18 Roadmap, governance and progression gates

A conditional thirty-six-month programme develops models and bounded demonstrations. Timing begins after requirements, staffing and resources are committed; dependent stages may be revised when evidence changes.

Period	Deliverable and gate
Months 0-6	Service contracts, topology/dependency map, data dictionary and reproducible baseline.
Months 6-12	Observatory, twin and known-answer optimisation; independent reconstruction.
Months 12-18	Qualified node and multi-route pilot design; readiness review.
Months 18-30	Multi-site fault tests, energy/timing loss and repeated recovery.
Months 30-36	Independent comparison and decisions to continue, redesign or close.

In the first ninety days, freeze essential services and failure criteria, qualify instrumentation, test approved failover/energy/timing scenarios and release quality/deviation reports. A schedule is a plan; no testbed construction or review is claimed completed by this document.

Assign separate responsibility for scientific/technical evidence, operational policy and affected-service governance. Record who may authorise changes and which reversible actions are already bounded by policy. Preserve authority and recovery procedures during isolation.

Stop or redesign if balances violate conservation, claimed gains fail replication, common causes defeat the essential contract, recovery cannot restore quality or consequences cannot be bounded. Deployment must use the standards, spectrum conditions and applicable requirements of its actual service and location. Referencing ITU or NIST is not a certificate of compliance. [T03-T11]

19 Glossary and formula interpretation

Term	Meaning in this framework
Resilience	Retain or recover a defined service under specified disturbances.
NTN	Non-terrestrial network; compatibility is service and implementation specific.
N–1	A declared service remains feasible after one specified component/route loss.
Headroom	Unallocated capacity; it may share the same failure as active capacity.
Link margin	Received power minus the receiver requirement for a declared mode.
SNR	Signal/noise power ratio; use linear form in Shannon's equation.
Holdover	Clock operation after loss of its disciplining reference.
Twin	Versioned observation/state/scenario model with declared uncertainty.
Failover / recovery	Alternative-path activation / restoration of contracted quality.
Safe state / rollback	Limit further effects / restore an earlier state.

Gbit/s describes a rate, while a stored byte count describes an amount. kW is power; kWh is energy. dBm references power to 1 mW, dB expresses a ratio, and dBi antenna gain uses an isotropic reference. Mixing these quantities can produce plausible-looking but invalid budgets.

A Shannon bound needs a specified channel and linear SNR. A free-space model excludes environmental losses unless they are added. An availability formula requires defined repair behaviour and dependency assumptions. A reserve estimate requires usable rather than nominal stock.

The worked examples make these distinctions visible. They should be replaced with qualified site and component evidence when used for a real design, while retaining their units, boundaries and decision logic.

20 Conclusion, shared annex and author attribution

Project This Station – Earth provides a useful engineering discipline for telecommunications when continuity is expressed as a defined service with measurable dependencies. The framework combines observability, energy/resource accounts, qualified substitution, independent recovery and traceable authority. Its contribution is this integration, not a unique device or an already built global network.

The next defensible step is a reproducible observatory and twin, followed by bounded multi-route tests with useful-service metrics. The known-answer routing case shows why nominal cost, spare capacity and $N-1$ essential continuity require separate checks. Link budget, timing and energy examples give bachelor-level readers a way to inspect specialist claims.

The common Annex contains `White_Flag_This_Station_Version.m4a` and `Multi-unit_Combinatorial_Reverse_Auctions_2005.pdf`. Methods include the reproduced audio script, JSON, CSV and calculation notes. Each language README explains reading order, file roles and scientific status. The 2005 original remains in English and retains its four authors; the audio is unchanged and retains attribution to White Flag (Dido), with this version supplied by Giuseppe Carilli.

No network performance, RF attenuation or security certification was measured for this paper. Its figures are explicit engineering examples and its experiments remain proposals. AI assistance supported synthesis, numerical implementation, translation and typesetting. This completed editorial edition precedes manual corrections and feedback.

Giuseppe Carilli Author · Project This Station – Earth · I'n'I / Lion of Judah LinkedIn: <https://es.linkedin.com/in/giuseppe-carilli-a03805b1> 6 October 2026 · Final technical edition FR1 · Typeset author attribution.

References · 1/2

Canonical source titles are retained in their original language. Online sources checked on 6 October 2026; versions used are identified below.

[T01] Carilli, G. (2026). Project This Station – Earth: A Clean Sail for Earth. Final consolidated edition FR1; companion paper in this package.

[T02] Giovannucci, A.; Rodríguez-Aguilar, J. A.; Cerquides, J.; Carilli, G. (11 July 2005). Multi-unit Combinatorial Reverse Auctions with Substitutability Relationships among Goods. Supplied original, 14 pages.

<https://www.iiia.csic.es/~jar/papers/2005/NAEC2005.camera.pdf>

[T03] ITU-R M.2160-0 (November 2023). Framework and overall objectives of the future development of IMT for 2030 and beyond.

<https://www.itu.int/rec/R-REC-M.2160-0-202311-I>

[T04] ITU-R P.525-5 (November 2024). Calculation of free-space attenuation.

<https://www.itu.int/rec/R-REC-P.525-5-202411-I/en>

[T05] ITU-R P.618-14 (August 2023). Propagation data and prediction methods required for the design of Earth-space telecommunication systems.

<https://www.itu.int/rec/R-REC-P.618/en>

[T06] 3GPP. Non-Terrestrial Networks (NTN): official technology overview and standardization context.

<https://www.3gpp.org/technologies/ntn-overview>

[T07] Shannon, C. E. (1948). A Mathematical Theory of Communication. Bell System Technical Journal 27, 379–423 and 623–656. Original paper, university-hosted copy.

<https://people.math.harvard.edu/~ctm/home/text/others/shannon/entropy/entropy.pdf>

[T08] ITU-T G.8271/Y.1366 (March 2020), Amendment 1 (August 2024). Time and phase synchronization aspects of telecommunication networks.

<https://www.itu.int/rec/t-rec-g.8271/en>

References · 2/2

Canonical source titles are retained in their original language. Online sources checked on 6 October 2026; versions used are identified below.

[T09] NOAA / NWS Space Weather Prediction Center. Space Weather Impacts.

<https://www.spaceweather.gov/impacts>

[T10] Ross, R. et al. (2021). Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. NIST SP 800-160, Volume 2, Revision 1.

<https://csrc.nist.gov/pubs/sp/800/160/v2/r1/final>

[T11] Rose, S. et al. (2020). Zero Trust Architecture. NIST SP 800-207.

<https://csrc.nist.gov/pubs/sp/800/207/final>

[T12] Burleigh, S.; Fall, K.; Birrane, E. (2022). Bundle Protocol Version 7. IETF RFC 9171.

<https://www.rfc-editor.org/rfc/rfc9171>

[T13] NASA. Small Spacecraft Technology State of the Art: 7.0 Thermal Control.

<https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/>

[T14] US Department of Energy (2022). Grid Energy Storage Technology Cost and Performance Assessment.

<https://www.energy.gov/cmei/2022-grid-energy-storage-technology-cost-and-performance-assessment>

[T15] Carilli, G. (2026). Marco técnico para la evolución resiliente del Sistema Global de Telecomunicaciones. Supplied technical report, 6 October 2026, 25 pages. Source for this elaboration.