

Scaling Bidirectional Telemetry for Large Aircraft Challenges and Solutions

Francisco M. Fernández Casuso¹
Airbus DS SAU, francisco.m.fernandez@airbus.com

Abstract:

Bidirectional telemetry is essential for the modern monitoring and control of aerospace platforms, enabling real-time data exchange and adaptive command transmission. While previous stages of this research successfully validated the system using laboratory environments and small-scale technology demonstrators, such as motor gliders, transitioning to large-scale aircraft introduces significant technical complexities. These challenges include managing increased cable lengths, mitigating signal shadowing caused by complex fuselage structures, and maintaining high bandwidth within demanding electromagnetic environments.

This paper proposes an integration approach designed to overcome these obstacles by leveraging advanced technologies, specifically multiple spatial streams and adaptive beamforming. We outline a comprehensive experimental test campaign aimed at assessing system performance on a large aircraft platform. The investigation focuses on coverage management, signal degradation due to platform topology, and the robustness of the bidirectional link. System effectiveness will be evaluated through key metrics, including signal strength, throughput, latency, and error rates. The results of this study aim to validate the scalability of the proposed system and provide critical insights into its behaviour under real-world operational conditions, marking a significant step toward more reliable and efficient aerospace telemetry solutions.

Key words: Bidirectional telemetry, aeronautical telemetry, C-band, MIMO, beamforming, IP-over-RF, air-to-ground link, link budget, flight test instrumentation,

1. Introduction

In recent years, the demand for bidirectional telemetry systems in aerospace platforms has grown significantly, driven by the need for real-time monitoring, video transmission, and remote command capabilities. Unlike traditional unidirectional telemetry, bidirectional systems enable direct interaction with onboard embedded systems, opening the door to adaptive flight test campaigns and closed-loop control architectures.

This work is part of a broader research effort that began with laboratory validation [1] and small-scale flight demonstrators, including motor glider platforms. That earlier phase concluded with the system reaching Technology Readiness Level 5, confirming its viability in a relevant flight environment [4]. The system deployed at that stage was intentionally reduced in scope — two radio nodes without amplification, transmitting at 1 W — yet it allowed full validation of the link budget, which proved conservative, providing more margin than initially predicted. Key findings directly inform the present work: spatial diversity — with both MIMO streams active — was identified as a critical factor for link stability, particularly during take-off and antenna

transition events; and the incompatibility of certain transport protocols with the intermittent nature of an air-to-ground link led to a fundamental revision of the communication stack, with direct implications for the architecture described in this paper. Building on these results, this paper addresses the next step in the scalability roadmap: the integration of the same architecture onto a large aircraft platform, where physical scale and operational constraints introduce a new tier of technical challenges.

The transition to large aircraft implies cable lengths with non-negligible RF losses, antenna placement dictated by structural constraints rather than RF optimization, and electromagnetic environments significantly more complex than those found in small-platform settings. Furthermore, integration must be carried out on an actively operated aircraft, requiring the reuse of pre-existing telemetry infrastructure wherever feasible.

This paper presents the design rationale, integration approach, and planned validation campaign for this phase. Final performance figures will be the subject of future publications once the flight test campaign is completed.

2. System Overview and Operational Objectives

The system establishes a bidirectional communication link between a large aircraft test platform and a ground station. Unlike conventional telemetry architectures, which are strictly unidirectional, the proposed design supports simultaneous transmission in both directions: downlink telemetry, video and sensor data from the aircraft, and uplink command and data reception from the ground station. This enables real-time interaction with onboard systems during flight test operations, without the constraints of post-flight data retrieval.

All services share a common IP transport layer. Telemetry, video and command traffic are encapsulated in standard IP packets and carried over a single RF link, which simplifies integration and allows the reuse of COTS networking equipment across different platform configurations. Connecting two fully operational IP networks — one airborne, one on the ground — does, however, introduce a non-trivial network management challenge: without adequate traffic isolation, broadcast and unintended traffic from either domain can propagate across the link, consuming bandwidth and degrading overall performance. Ensuring correct network boundary definition is therefore a fundamental requirement of the system design.

The system targets an end-to-end latency below 250 ms, consistent with the demands of real-time monitoring and command services. Transport is based on UDP, prioritising timely delivery over guaranteed reliability: for latency-sensitive applications, a late packet is operationally equivalent to a lost one, and retransmission mechanisms are deliberately excluded from the design. To meet this objective, strict Quality of Service policies have been applied at the network level, defining a traffic priority hierarchy across all services. Under this scheme, congestion is resolved through selective packet discarding rather than queuing: lower-priority flows — such as video — are dropped first, preserving bandwidth for critical telemetry streams, including positional data, whose integrity is essential to flight test operations.

3. Selected Architecture and Communication Technologies

The system is built around IP-native RF communication nodes operating in C-band, a frequency range that offers a suitable balance between propagation characteristics and available bandwidth for aeronautical telemetry applications [5]. Each node operates in a 2×2 MIMO configuration, providing two independent

spatial streams that improve both throughput and link robustness. Mesh networking protocols are incorporated at the network level, enhancing resilience under varying link conditions and laying the groundwork for future multi-node ground configurations.

Bidirectional amplifiers (BDAs) are deployed at both ends of the link — airborne and ground — to meet the coverage requirements imposed by the operational profile of a large aircraft. At the distances involved in this phase, the output power of the radio nodes alone is insufficient to maintain an adequate link margin, making external amplification a necessary element of the RF architecture. This represents a fundamental difference with respect to the previous project phase: on the motor glider platform, the expected operational range was considerably shorter, the airframe imposed stricter power constraints, and the proximity of the pilot to the antenna introduced additional safety considerations that ruled out the use of amplification. [4] On a large aircraft, none of these limitations apply, and the BDAs become a key enabler of the system. Their performance under operational conditions will therefore be one of the primary aspects evaluated during the test campaign. The architecture also supports an upgrade path to 4×4 MIMO should the results of the current phase indicate that additional spatial streams are warranted.

At the network level, services run predominantly over UDP, with TCP used only for short-lived transactions where reliability is required and latency tolerance permits, consistent with the design philosophy described in Section 2. Additionally, the impact of AES-256 encryption on system performance will be evaluated during the test campaign, with the aim of determining whether end-to-end encryption can be incorporated within the latency budget without compromising the integrity of critical data flows.

The radio nodes also support adaptive beamforming, which is of particular relevance on the airborne side. By combining the signals from the two antenna positions — tail and nose — with dynamically adjusted phase and amplitude weights, the system can concentrate radiated and received energy towards the ground station, partially compensating for the coverage limitations introduced by the non-optimal antenna placement. This capability operates continuously and transparently, adapting to the evolving geometry of the air-to-ground link throughout the flight. On the ground side, the directional antenna already provides a fixed beam oriented towards the aircraft, making adaptive beamforming redundant at that end. The interaction between beamforming and the

spatial diversity provided by the 2x2 MIMO configuration will be one of the aspects monitored during the test campaign, with the aim of quantifying its contribution to link margin under adverse geometries.

The ground station antenna completes the system architecture on the terrestrial side. A parabolic antenna is employed, providing the gain necessary to close the link budget at the operational distances involved. Unlike the airborne antennas, whose positions are constrained by the existing aircraft infrastructure, the ground antenna placement can be optimised for each test sortie, allowing the ground segment to partially compensate for the coverage limitations imposed by the airborne installation.

An additional aspect of particular relevance to this integration is the monitoring of the adaptive modulation scheme (MCS). The system dynamically adjusts its modulation and coding rate in response to link quality conditions, trading throughput for robustness as the link margin decreases. This behaviour makes the active modulation scheme a direct and continuous indicator of link health: degradation in signal quality manifests immediately as a downward shift in modulation order, providing real-time visibility into link capacity without requiring dedicated measurement infrastructure. In the context of a non-optimal integration — where cable losses, structural shadowing, and suboptimal antenna placement are expected to create variable link conditions throughout the flight envelope — tracking the modulation scheme over time will be one of the primary tools for characterising system behaviour and identifying the geometries and conditions under which performance is most affected.

4. Aircraft Integration Approach

The integration of the system into the aircraft has been one of the most demanding aspects of this project phase. Operating on an actively used flight test platform imposes significant constraints on the degree of modification that can be realistically achieved, and the approach has therefore been guided by a pragmatic philosophy: minimise structural intervention and maximise the reuse of existing infrastructure wherever possible.

In practice, this has meant working within the constraints of a pre-existing telemetry installation — both in terms of certified antenna positions and a significant portion of the RF cabling already routed through the aircraft. While this decision has reduced integration time and certification complexity, it introduces limitations that directly affect system performance and must be understood as a defining characteristic of this

phase. It should also be noted that prior to installation, the technology underwent a dedicated validation process involving the electrical and EMI/EMC departments, whose sign-off was required before any equipment could be fitted to the aircraft. This step, inherent to any integration on an operational platform, adds a layer of rigor to the process but also reflects the certification constraints within which the project operates.

Two airborne antennas are installed at positions inherited from the previous installation. The first is located on the lateral surface of the tail section, a placement that results in highly asymmetric coverage: only the illuminated side of the aircraft benefits from an adequate radiation pattern, while the opposite side falls entirely within the structural shadow of the fuselage. This is a known and accepted limitation of the current integration, and its effect on link availability across different flight geometries will be one of the primary subjects of the test campaign. The second antenna, positioned on the underside of the nose, is expected to provide more uniform ground coverage across a wider range of attitudes and headings, and will serve as the primary reference for link quality assessment.




Figure 1 – expected radiation diagram

A further constraint arises from the physical architecture of the amplification stage. A single BDA serves both antennas, and its location has been determined by the need to maintain a geometrically balanced cable distribution between the two antenna positions. This compromise results in cable runs exceeding 40 metres, introducing attenuation losses in the order of 8 dB — a significant but unavoidable penalty given the constraints of the existing installation. These losses are incorporated into the link budget [2][3] and represent one of the

key variables to be validated against real-world measurements during the test campaign.




Figure 2 – link budget summary

Taken together — suboptimal antenna placement, one-sided tail coverage, and elevated cable losses — these constraints define an integration scenario that is deliberately non-optimal. This is not considered a shortcoming of the phase, but rather an opportunity: operating under adverse, realistic conditions provides empirical data that idealised installations cannot. The results will form the basis for a more refined deployment in subsequent phases, and will allow the development of link models grounded in the practical realities of large aircraft integration.

5. Planned Validation and Test Campaign

The test campaign has been designed to validate the viability of the system under realistic operational conditions and to characterise its behaviour across the range of scenarios that a large aircraft integration entails. Given the constraints described in Section 4, the campaign serves a dual purpose: on one hand, to assess whether the system meets its functional requirements despite the non-optimal integration; on the other, to gather the empirical data needed to refine the link model and inform future phases of the project.

A further practical constraint shapes the structure of the campaign. Unlike dedicated flight test programmes where the test article is available on demand, the system must be evaluated within the existing flight schedule of an actively operated platform. This means adapting to available flight opportunities rather than defining an independent test matrix, which introduces variability in test conditions but also reflects the real-world constraints under which the technology will ultimately be deployed.

5.1 Objectives of the Test Campaign

The primary objective is to evaluate system performance in terms of latency, link robustness, coverage, and bidirectional throughput, while characterising the impact of the real aircraft integration on each of these parameters. A secondary but equally important objective is to identify the boundaries of acceptable performance — that is, to determine not only where the system works, but where and why it degrades.

The design of the test campaign has been directly shaped by the lessons drawn from the previous project phase, where three findings proved particularly influential.

The first concerns transport protocol behaviour. Long-lived TCP connections were found to be incompatible with the intermittent nature of an air-to-ground link, as connection interruptions left applications in a hung state from which recovery was slow and unreliable. UDP and short-lived TCP have therefore been adopted as the baseline transport configuration, and specific tests are included to verify that this behaviour scales correctly to the larger and more complex platform addressed here.

The second concerns spatial diversity. Link stability improved substantially when both MIMO streams were simultaneously active, an effect observed most clearly during take-off and during passes over the ground antenna. The test campaign therefore includes a dedicated evaluation of stream availability across different flight geometries, with particular attention to the conditions — shadow zones, antenna transitions, low-elevation passes — where single-stream operation is most likely.

The third concerns channel bandwidth. The 20 MHz configuration consistently outperformed narrower alternatives during the motor glider campaign and has been adopted as the baseline setting for this phase. Comparative measurements are included to quantify its advantage under the new integration conditions.

5.2 Test Categories

The campaign is structured into progressive test phases, each building on the results of the previous one.

Laboratory Tests establish the baseline before any aircraft involvement. Connectivity, latency, throughput, and system stability are verified in a controlled environment, both with and without AES-256 encryption. This phase ensures that any anomalies encountered during subsequent phases can be attributed to integration factors rather than to the equipment itself.

Ground Integration Tests mark the first time the airborne and ground infrastructures operate together as a single bidirectional system. The focus at this stage is not RF performance — link margin and antenna coverage are not the priority here — but rather the correct end-to-end integration of both domains. Tests verify that the desired traffic flows reach their intended destinations in both directions, and that no spurious or unintended traffic crosses the link boundary. Network isolation is validated under realistic conditions, confirming that the boundary management described in Section 2 behaves as designed when both IP networks are simultaneously active. Beyond connectivity, this phase covers the first functional validation of the use cases planned for the system: remote monitoring of onboard equipment from the ground station, bidirectional data exchange between airborne and ground nodes, and initial latency measurements under static conditions. These tests serve as the primary go/no-go gate before any airborne operation is attempted — if a use case does not work on the ground, it will not be cleared for flight.

Air Tests validate basic system operation in flight: continuous telemetry transmission, real-time video, uplink command reception, and bidirectional link stability across representative flight profiles. Within this phase, four categories of tests are distinguished: Performance Tests, Coverage and Geometry Tests, Bidirectional Use Case Tests, and Multi-node and Encryption Tests, each described in the following subsections.

Throughout all test phases, a consistent set of metrics will be recorded: end-to-end latency, jitter, packet loss rate, throughput, link availability, RSSI, and SNR. All measurements will be time-stamped and correlated with aircraft position and attitude data to enable spatial analysis of link performance.

5.3 RF Channel Characterisation Tests

A dedicated set of tests is planned to characterise the behaviour of the RF channel throughout the flight envelope. Unlike the ground integration phase, where the focus is on system interoperability, these tests are concerned with the physical layer of the link and its evolution under real airborne conditions.

The primary parameter under continuous monitoring will be the received signal strength (RSSI), recorded at both ends of the link as a function of aircraft position, heading, and altitude. Alongside RSSI, the signal-to-noise ratio (SNR) will be tracked to provide a more complete picture of link quality, particularly in scenarios where interference or multipath may

degrade performance independently of signal level.

Of particular interest is the evolution of the adaptive modulation scheme throughout the flight. Monitoring the active modulation and coding rate in real time provides a direct and continuous indicator of link health, and its variation across different flight geometries will be one of the primary datasets extracted from the campaign.

Time	SR(Rx/Tx)	F(MHz)	P(dB)	MCS	AirTime(s)	TotalRate(Mbit/s)	Gain	Drop
11:07:01:494	35/20	5140	30	13/14	9.837/07.06	0.120(85.00)	0	0.000
11:07:02:703	35/20	5140	30	13/14	7.830/07.06	0.074(85.00)	0	0.000
11:07:03:716	34/19	5140	29	13/14	10.501/07.06	0.112(85.00)	0	0.000
11:07:04:720	35/19	5140	30	13/14	9.229/07.06	0.080(85.00)	0	0.000
11:07:05:733	32/19	5140	29	13/14	9.764/07.06	0.110(85.00)	0	0.000
11:07:06:738	32/19	5140	30	13/14	9.766/07.06	0.110(85.00)	0	0.000
11:07:07:737	32/19	5140	29	13/14	8.156/07.06	0.074(85.00)	1	0.000
11:07:08:787	30/19	5140	30	12/14	9.399/07.06	0.100(85.00)	0	0.000
11:07:09:749	29/19	5140	30	12/14	9.755/07.06	0.097(85.00)	1	0.000
11:07:10:758	29/19	5140	29	12/13	12.289/08.00	0.129(75.00)	0	0.000
11:07:11:750	29/19	5140	31	12/13	9.056/08.00	0.094(75.00)	0	0.000
11:07:12:773	29/19	5140	29	12/13	8.885/08.00	0.103(75.00)	0	0.000
11:07:13:782	33/19	5140	29	12/13	9.574/08.00	0.090(75.00)	1	0.000
11:07:14:787	30/19	5140	30	5/07	16.060/12.64	0.106(47.45)	0	0.000

Figure 3 – Data being logged

Spatial diversity will also be systematically evaluated. The number of active MIMO spatial streams will be recorded throughout each flight, with particular attention to the conditions under which the system falls back to single-stream operation — take-off, shadow zone transitions, and low-elevation passes over the ground station. As demonstrated in the previous project phase, link stability improves substantially when both streams are simultaneously active, and replicating this analysis on a larger platform will determine whether the same effect holds under more demanding integration conditions.

Finally, the two antenna positions — tail and nose — will be compared across equivalent flight geometries to quantify the coverage difference introduced by their respective placements and to identify the scenarios in which each antenna provides the most reliable link.

5.4 Bidirectional Use Case Tests

Beyond RF characterisation, a second category of airborne tests addresses the functional validation of bidirectional data exchange between the ground station and the aircraft. These tests are organised around two complementary directions of operation, each with its own set of use cases.

Ground-to-aircraft tests focus on the ability to interact with onboard systems from the ground during flight. Planned activities include remote configuration changes to Flight Test Instrumentation equipment, remote access to onboard devices, and the upload of files of varying sizes to the aircraft. These tests are designed to validate not only connectivity but also the practical usability of the uplink under real flight conditions — latency, transfer stability,

and the behaviour of the system during link quality variations are all of interest.

Aircraft-to-ground tests exploit the downlink capacity beyond conventional telemetry streaming. In addition to standard parameter transmission, tests will include the download of files of varying sizes from the aircraft to the ground station, evaluating throughput and transfer reliability as a function of link quality. The ability to retrieve data from the aircraft in flight — rather than waiting for landing — represents one of the most operationally relevant capabilities of the system, with direct implications for flight test efficiency.

It should be noted that not all of these use cases reflect immediate operational requirements. Several are included as technology demonstrators, designed to explore the boundaries of what the system can support and to establish the empirical foundation for future use case development. The results will inform the definition of a validated use case catalogue for subsequent phases of the project.

5.5 Validation Hypotheses

The campaign is structured around five core hypotheses:

1. The proposed architecture is viable on a large aircraft, even under non-optimal integration conditions.
2. Reusing existing infrastructure allows functional operation within the defined performance requirements.
3. UDP transport is adequate for latency-sensitive services in this operational context.
4. A distributed antenna architecture improves link robustness compared to a single-antenna configuration.
5. The performance impact of AES-256 encryption remains within acceptable bounds relative to the system latency budget.

6. Expected Outcomes and Future Work

The primary expected outcome of this phase is the empirical validation — or informed refutation — of the five hypotheses defined in Section 5. In either case, the results will constitute a meaningful contribution: a system that performs as specified under the described integration constraints would demonstrate the viability of IP-based bidirectional telemetry on large aircraft platforms, while identified shortcomings would provide a well-characterised basis for targeted improvements in subsequent phases.

In terms of link performance, the system is expected to meet the 250 ms latency requirement — and low jitter — under nominal conditions, with degradation anticipated in geometries where the tail antenna shadow is dominant and the nose antenna cannot fully compensate. The 8 dB cable loss absorbed by the link budget represents a significant but calculated margin reduction, and one of the key outcomes of the campaign will be to determine how much operational margin remains under real flight conditions.

The coverage asymmetry introduced by the lateral tail antenna is expected to be one of the most visible findings of the campaign. Quantifying its extent — in terms of affected headings, distances, and link quality degradation — will directly inform the antenna placement strategy for the next integration phase, where a purpose-designed installation will no longer be constrained by pre-existing infrastructure.

Regarding the 4×4 MIMO upgrade path mentioned in Section 3, the results of this phase will provide the first empirical basis for assessing whether the additional spatial streams are necessary. If the 2×2 configuration proves sufficient within the operational envelope of the platform, the upgrade may be deferred; if link margin proves insufficient in a significant portion of the flight envelope, the 4×4 configuration will be prioritised for the next phase.

Finally, the evaluation of AES-256 encryption impact will determine whether end-to-end security can be incorporated into the system without exceeding the latency budget — a prerequisite for any future operational deployment beyond the flight test context.

From a technology maturity perspective, the previous phase demonstrated that the system had reached TRL 5 — technology validated in a relevant environment. The present phase is designed to advance this further: successful completion of the planned flight test campaign on a large aircraft platform, under the real operational constraints described throughout this paper, would constitute a meaningful step towards TRL 6.

References

- [1] Francisco Fernández Casuso: "Research into equivalent Bidirectional Telemetry", ETTC2024, Nuremberg 2024.
- [2] Francisco Fernández Casuso, Francisco Jiménez Martín: "Real time C Band Link Budget Model Calculation", ETTC2014 Toulouse, 2014.

- [3] Jose María Hernando Rábanos: " Transmisión por radio", chapter 5.6 basic parameters, 1993
- [4] Francisco Fernández Casuso: " Challenges and Solutions in Bidirectional Telemetry Integration", ETTC2025, Toulouse 2024.
- [5] WRC-23