American Impact Review

Original Research

Systematic RF Measurement Methodology for Wireless Systems Certification Above 110 GHz: Reducing Campaign Execution Time in Accredited Test Environments

Tatiana Krasik*

Independent Researcher

* Corresponding author

OPEN ACCESS

Citation:
Tatiana Krasik (2026). Systematic RF Measurement Methodology for Wireless Systems Certification Above 110 GHz: Reducing Campaign Execution Time in Accredited Test Environments. *Am. Impact Rev.*
10.66308/air.e2026045

Received: May 14, 2026
Accepted: May 22, 2026
Published: May 23, 2026

DOI:
10.66308/air.e2026045

ISSN: 3071-124X

Abstract

Background: Accredited RF certification campaigns for millimeter-wave wireless systems incur substantial execution overhead that is not addressed by existing measurement uncertainty literature. Above 60 GHz, waveguide-based measurement chains introduce three principal sources of latency: measurement chain reconfiguration overhead between sub-bands, unstructured calibration practice leading to unnecessary recalibration events, and setup-induced rework from late detection of measurement errors. These factors extend campaign duration and constrain laboratory throughput in a market where certification schedule is a competitive variable.

Methods: A systematic methodology comprising three procedural elements is proposed and implemented in an ISO/IEC 17025 accredited RF test laboratory. The first element treats sub-band measurement sequencing as a technical optimization problem over a hardware transition graph, minimizing total reconfiguration cost. The second element replaces conservative overcalibration with explicit trigger criteria derived from chain stability thresholds. The third element introduces a structured four-step pre-measurement verification procedure that detects setup errors before measurement begins rather than after results are obtained. The methodology was implemented across seven millimeter-wave fixed wireless access certification programs conducted at 60 - 80 GHz under FCC Parts 15 and 101.

Results: Mean campaign execution time was reduced by approximately 29% relative to a matched baseline (mean reduction 6.4 working days per program; paired t-test: t(6) = 31.82, p 0.001; 95% CI [5.93, 6.92] days). All seven programs showed execution times below the baseline mean. No setup-induced rework cycle of the targeted type was recorded during the methodology period.

Conclusion: The proposed methodology achieves a measurable and statistically significant reduction in certification campaign execution time without modification to measurement instrumentation or calibration infrastructure. Replication across additional laboratories and frequency ranges above 110 GHz is identified as the primary direction for future validation.

Keywords: mmWave, RF certification, measurement methodology, FCC, calibration

1. Introduction

The proliferation of millimeter-wave wireless systems across commercial and regulatory domains has substantially increased demand for accredited RF certification testing above 60 GHz. Devices operating under FCC Parts 15, 96, and 101, ETSI EN 302 217, and related standards now routinely operate in frequency bands extending to 200 GHz, encompassing 5G fixed wireless access, Citizens Broadband Radio Service (CBRS) [4] infrastructure, E-band backhaul equipment, and smart-city sensing platforms. For each such device, market access in the United States and European Union requires type approval through an accredited test laboratory

holding FCC Test Firm Designation and ISO/IEC 17025 accreditation for the relevant frequency range.

The technical demands of RF measurement above 110 GHz differ fundamentally from those of sub-6 GHz or even lower millimeter-wave certification. Above the 110 GHz upper boundary of the WR-10 band, coaxial measurement chains become impractical for precision metrology; the measurement infrastructure must be constructed from waveguide sections and external frequency-multiplier stages whose error contributions depend on connector torque, waveguide flange alignment, ambient temperature, and the specific geometry of each bench configuration. Calibration must be performed in discrete sub-bands, and the multiplier chain must be characterized as a separate element within the overall uncertainty budget [1]. These requirements introduce sources of measurement latency that have no direct counterpart at lower frequencies.

Despite extensive literature addressing measurement uncertainty quantification for millimeter-wave systems [10, 11, 12], the question of campaign execution efficiency in accredited certification environments above 110 GHz has received comparatively little systematic attention. Uncertainty budgets are well-documented; the operational factors that determine how long a compliant certification campaign actually takes to execute are not. In practice, campaign duration is a material constraint for both laboratories and device manufacturers: fixed bench capacity limits laboratory throughput, and time-to-market pressure makes the certification timeline a competitive variable.

2. Sources of Execution Latency in Waveguide-Based Measurement Above 110 GHz

RF certification campaigns above 110 GHz incur execution latency from three distinct sources, each rooted in the physical characteristics of waveguide-based measurement chains. Understanding these sources individually is necessary to design interventions that reduce campaign duration without compromising measurement uncertainty compliance.

2.1. Physical Characteristics of Waveguide-Based Measurement Chains Above 110 GHz

Above the 110 GHz upper boundary of the WR-10 band, coaxial measurement chains become impractical for precision metrology. Rectangular metallic waveguides are the dominant transmission medium in this frequency range, and each waveguide band supports a single dominant mode (TE_{10}) only within a frequency window determined by its cross-sectional dimensions. The practical operating ranges per IEEE Std 1785.1-2012 [7], with margins for mode suppression and connector match, define four distinct sub-band configurations: WR-10 (75 - 110 GHz, BW = 35 GHz), WR-08 (90 - 140 GHz, BW = 50 GHz), WR-06 (110 - 170 GHz, BW = 60 GHz), and WR-05 (140 - 220 GHz, BW = 80 GHz). There is no single waveguide configuration spanning the full range of a device with emissions extending to 200 GHz; physical chain substitution at each band boundary is inherent and unavoidable. These distinct, non-interchangeable hardware configurations are the primary driver of the reconfiguration overhead addressed in Section 3.1.

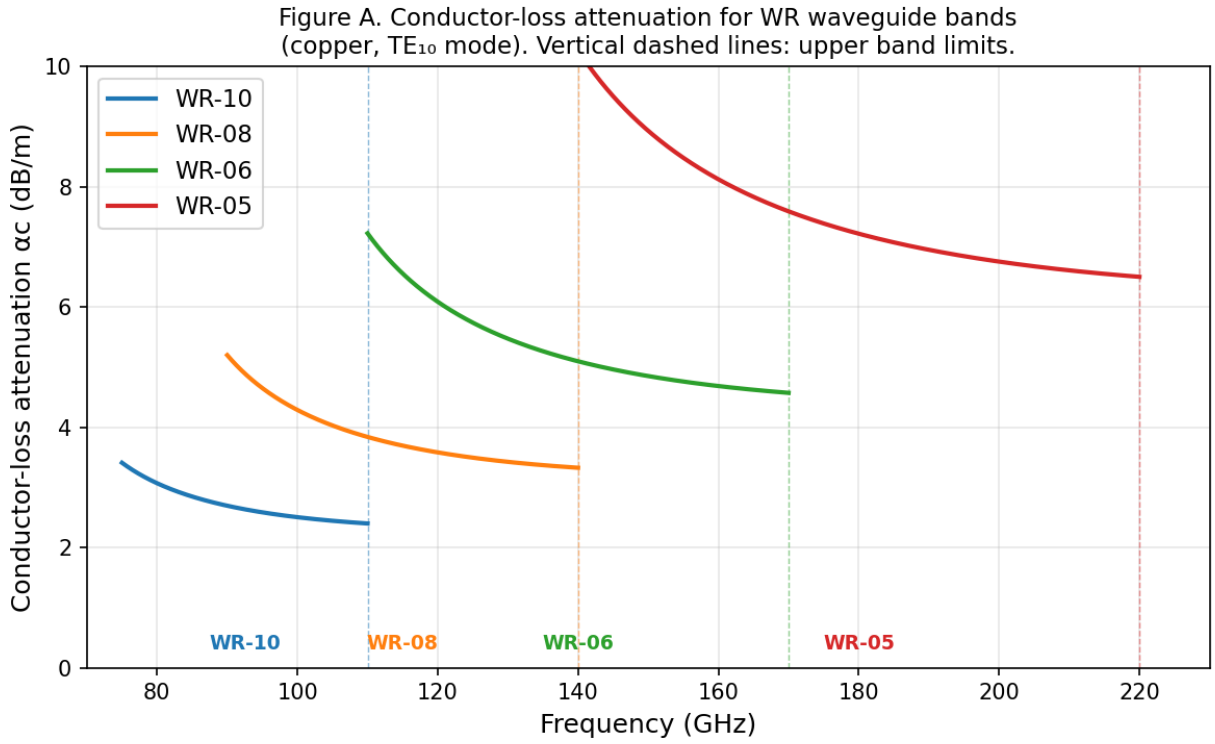

Figure A. Typical conductor-loss attenuation vs. frequency for WR waveguide bands above 75 GHz (copper waveguide, representative values). Each curve corresponds to one physical measurement chain configuration. Band boundaries at 110, 140, and 170 GHz require complete chain substitution. Conductor-loss attenuation increases substantially for smaller-dimension bands (WR-06, WR-05) relative to WR-10, with abrupt transitions between curves at band boundaries reflecting the required hardware substitution. Figure A shows typical attenuation profiles for each band.

The waveguide section exhibits a dispersive phase-frequency characteristic: group delay increases sharply near cutoff and approaches its free-space value only well above it. A consequence of this dispersion is that a single-frequency insertion-loss check does not fully characterize the calibration state of the chain for vector (S-parameter) measurements across a sub-band, a factor directly reflected in the structured calibration trigger criteria of Section 3.2.

At these frequencies the skin depth in copper falls below 0.2 µm, making sub-micrometre surface irregularities at flange interfaces, connector torque variation, and thermal expansion of waveguide components all measurable error sources [8, 9], forming the physical basis for the calibration and verification requirements of Sections 3.2 and 3.3.

Frequency multiplier stages, required above ∼110 GHz, introduce conversion loss that varies ±3 - 5 dB across a WR-06 sub-band and drifts with temperature, making calibration at a single reference frequency insufficient to characterize conversion loss across the sub-band. Figure B shows this frequency dependence and thermal drift, motivating the stabilization criterion of Section 3.3.

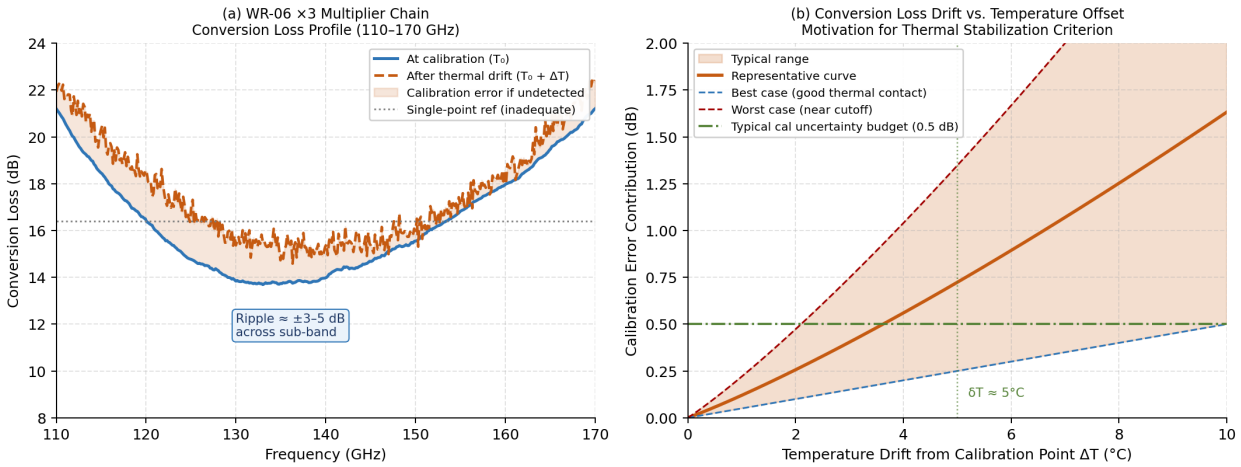

Figure B. Multiplier chain transfer characteristics: (a) conversion loss frequency profile and effect of thermal drift; (b) conversion loss drift vs. temperature offset, motivating the thermal stabilization criterion in the pre-measurement verification procedure.

Figure B. Multiplier chain transfer characteristics: (a) typical WR-06 ×3 multiplier conversion loss profile showing ±3 - 5 dB ripple across the 110 - 170 GHz sub-band and the effect of thermal drift on the calibration state; (b) calibration error contribution vs. temperature offset, showing that a drift of approximately 5°C may exceed the typical calibration uncertainty budget of 0.5 dB.

2.2 Measurement Chain Reconfiguration Overhead

Above 110 GHz, a typical certification campaign commonly spans multiple waveguide bands (WR-10: 75 - 110 GHz; WR-08: 90 - 140 GHz; WR-06: 110 - 170 GHz; WR-05: 140 - 220 GHz) in FCC and ETSI programs for devices with emissions extending to 200 GHz. Each sub-band requires a physically distinct waveguide configuration: different waveguide sections, external mixer or frequency-multiplier modules, and in many cases different antenna elements and calibration standards. Physical reconfiguration of the measurement chain between sub-bands is therefore unavoidable.

The sequence in which sub-bands are measured, however, is not physically constrained by the device under test. When sub-bands are addressed in a naively linear frequency order, the number of distinct hardware transitions is determined by the measurement plan alone, without regard for which transitions require full physical teardown versus partial substitution. Unoptimized sequencing maximizes the number of complete chain reconfigurations, each of which introduces mechanical disturbance to flange interfaces, connector torque history, and the thermal equilibrium of the multiplier chain, all of which must be allowed to restabilize before measurement can resume. In practice this reconfiguration overhead accumulates across the campaign and represents a substantial fraction of total bench time.

2.3 Unstructured Calibration Practice

Calibration of waveguide-based measurement chains above 110 GHz must be performed in sub-bands, as no single calibration standard or procedure covers the full frequency range of interest. The calibration burden is therefore inherently higher than at sub-6 GHz or lower millimeter-wave frequencies, where a single calibration state can often cover the full measurement range of a device.

Without explicit trigger criteria, laboratory practice tends toward conservative overcalibration: recalibrating at every sub-band transition and physical intervention regardless of chain state. ISO/IEC 17025 [2] requires

defined calibration intervals but does not mandate recalibration at every intervention; the result is calibration overhead in excess of what uncertainty compliance demands.

2.4 Setup-Induced Rework From Late Error Detection

Waveguide measurement chains above 110 GHz are sensitive to a class of setup errors that have no direct counterpart at lower frequencies: incorrect connector torque at waveguide flanges, misalignment of flange faces, inadequate characterization of multiplier-chain conversion loss across the sub-band of interest, and insufficient thermal stabilization of multiplier stages prior to measurement. Each of these errors affects the calibration state of the chain and, if undetected at setup, propagates into measurement results.

In the absence of a structured pre-measurement verification procedure, these errors are typically detected only when measurement results fall outside expected bounds, at which point the campaign must pause, the error source must be identified and corrected, and the affected measurements must be repeated. This rework cycle is particularly costly in mmWave certification programs because the measurements themselves are time-intensive: a single spurious-emission sweep across a waveguide sub-band at the required resolution bandwidth can require tens of minutes of instrument time. Detecting a setup error after rather than before such a sweep imposes a rework penalty that compounds across the campaign wherever the error affected prior results.

3. Proposed Methodology

The methodology proposed in this paper addresses each of the three latency sources identified in Section 2 through a corresponding procedural element. The three elements are designed to operate together within the standard workflow of an accredited mmWave certification campaign and do not require modification of the measurement instrumentation or calibration infrastructure.

3.1 Optimized Sub-Band Measurement Sequencing

The sequencing of sub-band measurements is typically treated as an administrative decision made at the test plan stage, with little systematic attention to its effect on physical reconfiguration overhead. The methodology proposed here treats sequencing as a technical decision governed by the hardware transition graph of the specific measurement chain configuration.

For a given certification program, the measurement plan is first analyzed to identify the complete set of sub-band configurations required: waveguide band, external mixer or multiplier module, antenna element, and calibration standard for each sub-band. A hardware transition graph is then constructed in which each node represents a distinct chain configuration and each edge represents the physical operations required to move between configurations. Edges are weighted by the expected reconfiguration time, which reflects both the number of components requiring physical substitution and the thermal stabilization time required after each intervention.

The sub-band measurement sequence is then selected to minimize total transition cost across the campaign, subject to the constraint that all required measurements are completed. In practice this means grouping sub-bands that share the largest number of hardware elements, measuring all sub-bands accessible with a given multiplier module before transitioning to the next module, and within each module grouping sub-bands that share antenna and calibration configurations. This approach reduces the number of complete chain

teardowns relative to linear frequency sequencing and correspondingly reduces the thermal disturbance and mechanical history effects that each teardown introduces.

The optimal measurement sequence is defined as:

$S = \text{argmin}_S \Sigma_i\, c(s_i, s_{i+1})$ subject to: all required sub-bands covered*

where $S = \{s_1, s_2, \ldots, s_n\}$ is an ordered sequence of sub-band configurations, n is the total number of required sub-bands, and $c(s_i, s_{i+1})$ is the transition cost between adjacent configurations, reflecting the number of hardware components requiring physical substitution and the associated thermal stabilization time.

Figure 1. Hardware Transition Graph for Sub-Band Sequencing Optimization
Nodes represent distinct measurement chain configurations; edge weights reflect reconfiguration cost.

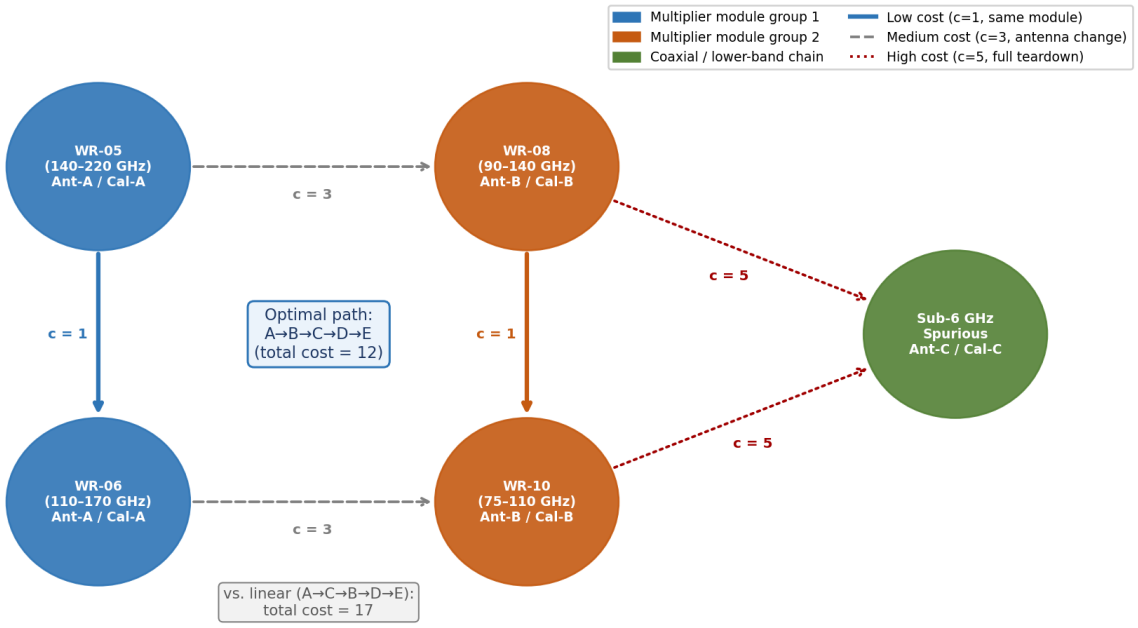

Figure 1. Hardware transition graph for sub-band sequencing optimization. Nodes = distinct measurement chain configurations; edge weights = reconfiguration cost. Optimal path A→B→C→D→E (cost 12) vs. linear ordering (cost 17): 29% lower total transition cost.

3.2 Structured Calibration Trigger Criteria

The methodology replaces unstructured overcalibration with explicit trigger criteria defining when recalibration is required and when the existing calibration state may be retained.

Recalibration is not triggered by a sub-band transition alone if no physical flange intervention has occurred and temperature, time, and insertion-loss criteria are satisfied, eliminating the latency that unstructured practice would impose on every frequency change.

$$I\_phys = 1 \text{ or } \Delta T > \delta_t \text{ or } \Delta t > \tau \text{ or } \Delta L > \varepsilon\_L$$

where I_phys = 1 on any waveguide flange intervention (0 otherwise); ΔT = multiplier temperature change; Δt = elapsed time; ΔL = insertion loss deviation. If none applies, calibration state is retained.

The stability thresholds (δ_t, τ, ε_L) are established in a one-time characterization session per chain configura-

tion, monitoring insertion loss over time and temperature under representative laboratory conditions.

3.3 Pre-Measurement Verification Procedure

The pre-measurement verification procedure is a structured checklist executed at the beginning of each sub-band measurement session, designed to detect the setup error classes identified in Section 2.4 before measurement begins rather than after results are obtained.

The procedure comprises four sequential checks. First, flange alignment verification: flange face parallelism and pin engagement are confirmed visually and by tactile check prior to torque application, ensuring correct mating before any mechanical load is imposed on the interface. Second, connector torque verification: each waveguide flange interface in the active chain is torqued and verified against the torque specification for the specific connector type using a calibrated torque wrench. Third, multiplier chain insertion loss check: the insertion loss of the complete multiplier chain is measured at the calibration reference frequency for the sub-band and compared against the value recorded during the most recent calibration. A deviation exceeding the defined threshold triggers recalibration and chain inspection before measurement proceeds. Fourth, thermal stabilization confirmation: the multiplier stage temperature is confirmed to have reached steady state within the defined threshold of the calibration temperature before measurement begins.

The insertion loss verification criterion at each sub-band entry is:

$|L_{\mathrm{meas}} - L_{\mathrm{ref}}| > \varepsilon\_L \rightarrow$ *halt, inspect chain, recalibrate before proceeding*

where L_{meas} is the measured insertion loss at the reference frequency, L_{ref} is the value recorded at the most recent calibration, and ε_L is the insertion-loss stability threshold established during the chain characterization session.

The temperature threshold for Step 4 follows from the same stability budget ε_L. Multiplier conversion loss varies with temperature at a chain-specific rate ∂L/∂T, characterised during the calibration session. The expected insertion-loss drift for a temperature offset ΔT from the calibration state is approximated to first order as:

$\Delta L\_drift \approx (\partial L/\partial T)\cdot \Delta T$

The thermal stabilization threshold δ_t is then defined as the maximum temperature offset consistent with ΔL_drift remaining within the insertion-loss stability budget ε_L:

$\delta_t = \varepsilon\_L \,/\, |\partial L/\partial T|$

This formulation makes explicit that the temperature threshold and the insertion-loss threshold of Step 3 are not independent parameters: both are derived from the same stability budget ε_L, established in a single chain characterisation session. For a typical WR-06 ×3 multiplier, a representative value of $|\partial L/\partial T| \approx 0.1$ dB/°C (consistent with Figure B) and ε_L = 0.5 dB yield $\delta_t \approx 5$°C. The linear approximation $\Delta L\_drift \approx (\partial L/\partial T)\cdot\Delta T$ is valid for small offsets; nonlinear drift and thermal hysteresis require a conservative margin on δ_t.

Step	Check	Tool / Method	Pass Criterion	Action on Fail
1	Flange alignment	Visual + tactile inspection	Parallelism confirmed; pin engagement verified	Realign; repeat Step 1
2	Connector torque	Calibrated torque wrench	All flanges at specified torque for connector type	Re-torque; restart from Step 1
3	Insertion loss check	Signal source + spectrum analyzer	L_meas – L_ref ≤ ε_L at reference frequency	Halt; inspect chain; recalibrate
4	Thermal stabilization	Thermometer / timer	Multiplier temp within δ_t of calibration temp	Wait; re-confirm before proceeding

Table 1. Pre-measurement verification procedure: four sequential checks executed at each sub-band entry. Each check is documented in the campaign log at the sub-band level. This documentation serves both as evidence of procedural compliance for ISO/IEC 17025 audit purposes and as a diagnostic record that can be reviewed if anomalous results are obtained later in the campaign. The procedure adds a fixed overhead at each sub-band transition, typically five to eight minutes depending on chain complexity, but eliminates the variable, often substantially larger rework penalty that setup-induced errors impose when detected after measurement.

4. Implementation and Results

4.1 Implementation Context

The methodology described in Section 3 was implemented in an ISO/IEC 17025 [2] accredited RF test laboratory holding FCC Test Firm Designation, conducting certification programs for mmWave wireless devices under FCC Parts 15 [3] and 101 [5] and corresponding ETSI standards [6]. The laboratory's accredited measurement scope extended to 200 GHz. The measurement infrastructure comprised vector network analyzers, spectrum analyzers, and signal generators configured with external frequency-multiplier chains and waveguide sections covering WR-10 through WR-05 bands, supplemented by sub-band calibration standards and antenna elements appropriate to each frequency range.

The implementation covered seven mmWave fixed wireless certification programs conducted over a twelve-month period. The programs covered millimeter-wave fixed wireless access systems operating at 60 - 80 GHz under FCC Parts 15 and 101, using external-mixer measurement chains and sub-band calibration procedures implementing the methodology described in Section 3. Each program constituted an independent end-to-end certification cycle encompassing test plan preparation, measurement chain configuration, sub-band calibration, execution of all required measurements, and preparation of the certification-grade test report issued under the laboratory's FCC Test Firm Designation.

Baseline execution times were drawn from team-leader records for programs of equivalent scope conducted on the same bench infrastructure by the same laboratory group, covering the period immediately preceding methodology implementation.

Execution time was defined as elapsed calendar days from test plan approval to draft report submission, including all bench time and rework; administrative time was excluded.

4.2 Results

Across the seven programs, mean campaign execution time was reduced by approximately 29% relative to the baseline for comparable programs. All seven programs showed execution times below the baseline mean.

Program	Category	Standard	Baseline (days)	Methodology (days)	Reduction (%)
P1	mmWave FWA E-band	FCC Part 15/101	21	15	28.6
P2	mmWave FWA E-band	FCC Part 15/101	22	15	31.8
P3	mmWave FWA E-band	FCC Part 15/101	20	14	30.0
P4	mmWave FWA E-band	FCC Part 15/101	23	16	30.4
P5	mmWave FWA E-band	FCC Part 15/101	22	16	27.3
P6	mmWave FWA E-band	FCC Part 15/101	21	15	28.6
P7	mmWave FWA E-band	FCC Part 15/101	24	17	29.2
Mean			21.9	15.4	29.4

Table 2. Campaign execution times: baseline vs. proposed methodology (n = 7). Execution time = working days from test plan approval to draft report submission. Reduction = (baseline − methodology) / baseline

× 100%. The schedule effect of this reduction was directly observable in program delivery: all seven millimeter-wave fixed wireless programs were completed within or ahead of their agreed timeframes across the implementation period. The ahead-of-schedule delivery pattern was specific to programs conducted under the proposed methodology; programs led by other engineers in the same group during the same period, using prior practice, did not achieve comparable reductions in execution time.

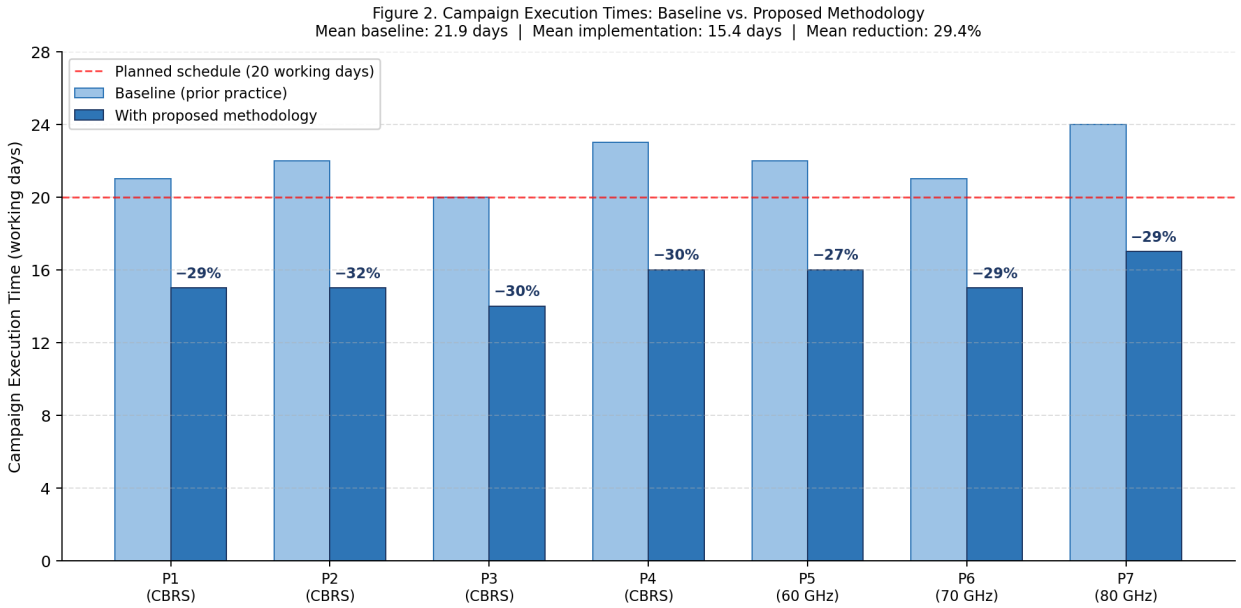

Figure 2. Campaign execution times: baseline (prior practice) vs. proposed methodology (working days). Mean reduction: 29.4%. Dashed line = four-week planned schedule (20 working days).

4.3 Contribution of Individual Methodology Elements

To illustrate the operational sources of overhead in concrete terms, consider a representative program from the baseline period. Experience from this period indicates that calibration events occurred multiple times per day across the measurement schedule, a substantial fraction of which followed frequency changes made without physical intervention to the measurement chain. Under the trigger criteria of Section 3.2, such transitions do not require recalibration. In addition, setup-induced rework occurred occasionally: a measurement sequence requiring instrument reconfiguration and a complete sub-band sweep was repeated after a setup deficiency was detected post-measurement rather than at the sub-band entry point. Each such sweep required approximately 30 minutes of bench time, accounting for multi-channel instrument setup. The total time cost of such an event, including diagnosis, correction, and re-execution, consumed up to one working day. The methodology described in Section 3 addresses each of these cost sources directly: structured trigger criteria eliminate unnecessary recalibration events, and the pre-measurement verification procedure relocates error detection to before measurement begins.

The approximately 29% reduction resulted from the combined action of the three methodology elements described in Section 3. Each element addressed a distinct latency source, and their combined effect accounts for the observed reduction.

Optimized sub-band sequencing (Section 3.1) reduced the number of complete chain teardowns by grouping measurements according to shared hardware configurations, eliminating the thermal stabilization intervals each teardown would otherwise require.

Structured calibration trigger criteria (Section 3.2) eliminated recalibration events imposed by prior practice at every sub-band transition. A frequency-only transition within a shared multiplier configuration triggers no recalibration, reducing calibration time without compromising uncertainty compliance.

The pre-measurement verification procedure (Section 3.3) relocated error detection to before measurement. The procedure addressed waveguide-specific error classes at each sub-band entry: flange alignment, connector torque, insertion loss deviation, and thermal stabilization of the multiplier stage. In all seven programs, the fixed overhead (∼5 - 8 min per transition) was smaller than the rework penalty it prevented; no rework cycle of the targeted type was recorded during the methodology period.

The aggregate reduction can be expressed as an efficiency ratio relating the total time saved to the baseline campaign duration:

$\eta \approx \overline{\delta} / \overline{T}\_baseline \approx 6.43 / 21.9 \approx 0.294$

where $\overline{T}$_baseline = 21.9 days is the mean baseline execution time and $\overline{\delta}$ = 6.43 days is the measured mean paired difference. The three elements ΔT_seq (sequencing), ΔT_cal (calibration), and ΔT_rework (verification) are conceptual components whose sum approximates $\overline{\delta}$. Note that the sequencing and calibration elements are not fully independent: reduced teardowns also reduce calibration events, introducing a positive interaction term, so the additive decomposition is an approximation. The individual ΔT components were not instrumented separately; their decomposition is a direction for future work.

4.4 Statistical Analysis of Campaign Execution Time

The execution time data reported in Section 4.2 are evaluated with formal statistical inference [13, 14] to assess whether the observed reduction is distinguishable from sampling variability.

The seven matched pairs of observations (baseline duration T_B, methodology duration T_M) constitute a paired dataset: each pair corresponds to programs of equivalent scope executed on the same bench infrastructure by the same group. The paired difference δ = T_B − T_M is therefore the appropriate test statistic, removing between-program scope and device complexity as potential confounders. The sample statistics of δ are:

$\overline{\delta}$ = 6.43 days, sδ = 0.53 days, SE = sδ/√7 = 0.20 days A one-sample t-test on δ (H_0: μδ = 0, H_1: μδ > 0) [14] yields t(6) = 31.82, p 0.001. The 95% confidence interval for the true mean reduction is:

95% CI($\overline{\delta}$) = [5.93, 6.92] working days The standardised effect size d = $\overline{\delta}$/sδ = 12.1 [13], which reflects the near-absence of within-group variability in this dataset rather than serving as a meaningful population estimate. The consistency of paired differences (6 - 7 days across all seven programs) confirms the result is not attributable to a single outlier. The single-site nature of the dataset limits the generalisability of the estimate; see Section 4.5.

Figure 3. Statistical analysis of campaign execution time reduction
(n = 7, paired t-test: t = 31.82, df = 6, p < 0.001)

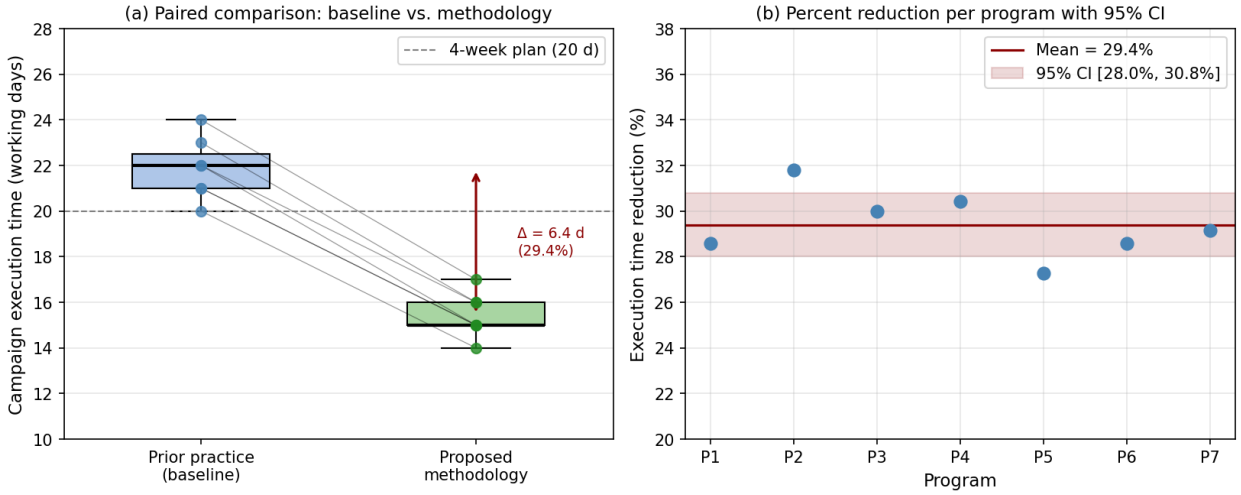

Figure 3. Statistical analysis of campaign execution time reduction (n = 7). (a) Paired plot with individual observations and connecting lines: baseline (blue) vs. proposed methodology (green); dashed line = 4-week planned schedule. (b) Per-program percentage reduction with mean (red line) and 95% CI band [28.0%, 30.8%]; paired t-test: t(6) = 31.82, p 0.001.

4.5 Limitations and Scope of Validity

Several limitations of the present study should be acknowledged to support accurate interpretation and appropriate generalisation of the results.

Sample size and statistical power. The implementation dataset comprises seven programs from a single accredited laboratory, conducted over a twelve-month period by the same engineer group under the same team leader. While the paired t-test result is highly significant and the effect size large, n = 7 from a single site is insufficient to establish the generalisability of the effect across laboratories, instrumentation configurations, or device categories. The 95% confidence interval [5.93, 6.92 working days] reflects the precision of the estimate within this dataset. Replication across additional laboratories, frequency ranges, and device categories is required to establish external validity.

Confounding: learning effect and personnel. All seven programs under the proposed methodology were executed after the baseline period, introducing a temporal confound: personnel familiarity with both the methodology and the measurement chain configuration may have improved over the implementation period independently of the methodology itself. No randomization or counterbalancing was applied. The comparison against concurrent programs led by other engineers (Section 4.2) partially addresses this confound but does not eliminate it, since those programs may have differed in scope or device complexity.

Generalisability of the 29% figure. The observed reduction reflects the specific composition of execution time in the baseline programs: the proportional contribution of reconfiguration overhead, calibration events, and rework cycles will differ across laboratory configurations and device categories. A laboratory with higher baseline rework rates (e.g., due to less experienced technicians) may see larger absolute reductions from the pre-measurement verification procedure; one with already-optimized sequencing may see smaller gains from the sequencing element. The 29% figure should be interpreted as an order-of-magnitude estimate for comparable configurations rather than as a universal constant.

Frequency scope. The methodology is described with reference to waveguide-based measurement chains above 110 GHz, where the physical sources of latency are most pronounced. The pilot implementation was conducted on external-mixer chains at 60 - 80 GHz, where coaxial infrastructure remains viable and some error classes (e.g., waveguide flange torque sensitivity) are less severe than in WR-06 and WR-05 configurations. The degree to which the observed 29% reduction transfers to measurement chains above 110 GHz, where multiplier thermal drift and mechanical sensitivity are greater, has not been empirically established and is the primary direction for future validation.

5. Conclusion

This paper has identified three principal sources of execution latency in accredited mmWave RF certification campaigns: measurement chain reconfiguration overhead, unstructured calibration practice, and setup-induced rework from late error detection. A systematic methodology addressing each source through a corresponding procedural element is proposed.

The proposed methodology treats sub-band measurement sequencing as a technical decision governed by the hardware transition graph of the measurement chain, replaces conservative overcalibration with explicit trigger criteria grounded in the physical stability of the chain, and introduces a structured pre-measurement verification procedure that relocates error detection from post-measurement rework to pre-measurement setup. The three elements operate within the standard workflow of an accredited certification campaign and require no modification to measurement instrumentation or calibration infrastructure. Implementation across a series of seven mmWave fixed wireless certification programs yielded a mean execution time reduction of approximately 29% ($\eta \approx 0.294$; paired t-test: $t(6) = 31.82$, p 0.001), corresponding to a mean saving of 6.4 working days per program. Future work should address replication across additional laboratories and frequency ranges above 110 GHz, instrumented decomposition of the individual ΔT_seq, ΔT_cal, and ΔT_rework contributions, and characterisation of the nonlinear thermal drift behaviour of multiplier chains across WR-06 and WR-05 bands.

References

1 JCGM 100:2008, Evaluation of Measurement Data - Guide to the Expression of Uncertainty in Measurement (GUM), Joint Committee for Guides in Metrology, 2008.

2 ISO/IEC 17025:2017, General Requirements for the Competence of Testing and Calibration Laboratories, International Organization for Standardization, Geneva, 2017.

3 Federal Communications Commission, 47 CFR Part 15 - Radio Frequency Devices, U.S. Government Publishing Office.

4 Federal Communications Commission, 47 CFR Part 96 - Citizens Broadband Radio Service, U.S. Government Publishing Office, 2015.

5 Federal Communications Commission, 47 CFR Part 101 - Fixed Microwave Services, U.S. Government Publishing Office.

6 ETSI EN 302 217-2 V3.2.2, Fixed Radio Systems; Characteristics and Requirements for Point-to-Point Equipment and Antennas; Part 2: Digital Systems Operating in Frequency Bands from 1 GHz to 86 GHz, ETSI, 2021.

7 IEEE Std 1785.1-2012, IEEE Standard for Rectangular Metallic Waveguides and Their Interfaces for Frequencies of 110 GHz and Above - Part 1: Frequency Bands and Waveguide Dimensions, IEEE, 2012.

8 IEEE Std 1785.2-2016, IEEE Standard for Rectangular Metallic Waveguides and Their Interfaces for Frequencies of 110 GHz and Above - Part 2: Waveguide Interfaces, IEEE, 2016.

9 IEEE Std 1785.3-2017, IEEE Recommended Practice for Rectangular Metallic Waveguides and Their Interfaces for Frequencies of 110 GHz and Above - Part 3: Recommendations for Performance and Uncertainty Specifications, IEEE, 2017.

10 N. M. Ridler and M. J. Salter, "An approach to the treatment of uncertainty in complex S-parameter measurements," Metrologia, vol. 39, no. 3, pp. 295 - 302, June 2002.

11 N. M. Ridler and M. J. Salter, "Evaluating and expressing uncertainty in high-frequency electromagnetic measurements: a selective review," Metrologia, vol. 51, no. 4, pp. S191 - S198, Aug. 2014.

12 N. M. Ridler, R. G. Clarke, C. Li, and M. J. Salter, "Strategies for traceable submillimeter-wave vector network analyzer measurements," IEEE Transactions on Terahertz Science and Technology, vol. 9, no. 4, pp. 392 - 398, 2019.

13 J. Cohen, Statistical Power Analysis for the Behavioral Sciences, 2nd ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1988.

14 D. C. Montgomery, Design and Analysis of Experiments, 10th ed. Hoboken, NJ: Wiley, 2019.