TECHNICAL DESIGN REPORT

Wideband S-Band Feed and Ecartometer Design Report

From first antenna baselines to the V5A integrated feedhead blueprints

Prepared for	Wideband S-band feedhead and ecartometer design review
Prepared on	14 June 2026
Simulation scope	2.200-2.400 GHz S-band feed, selectable CP center feed, four-sensor ecartometer ring
Final package referenced	V5A_integrated_wideband_feedhead
Primary source folders	Run1-Run12 outputs plus FabPackage/V5_fusion_clean_components

Reading order: The report starts with the original antenna question and follows the evidence forward. The selected V5A dimensions and ISO-style blueprints appear after the simulation story, so the reader sees why the design ended there instead of being handed the endpoint first.

Executive Summary

The project began with a basic S-band patch question: what antenna geometry can provide a usable payload feed, circular polarization behavior, and enough bandwidth around 2.3 GHz while still leaving room for tracking sensors and fabrication hardware. The first runs tested the obvious patch families and exposed the main tension: single-plane geometries could match well at one frequency, but the band-edge match, circular polarization margin, and sensor coupling constraints did not all close together.

The winning RF path was a lower driven circular patch plus an upper floating circular parasitic patch. The passive-stack tests proved that this family was worth pursuing, the center-feed consolidation selected the architecture source, and the final height-distribution sweep moved the build dimensions to a 3.5 mm ground-to-lower gap and a 5.0 mm lower-to-upper gap, with lower radius 35.75 mm and upper radius 31.50 mm.

The mechanical path followed the RF results instead of inventing a new antenna. V3 translated the Run7/Run8 architecture into modular tiles. V4 updated the package to the Run12 height and radius selection and added a windowed carrier/backplate concept. V5A cleaned the Fusion import by exporting separate components, explicit RF windows, keepouts, feed-driver pin references, drawings, and validation manifests.

Objective Roadmap

The simulation campaign was not a single monolithic optimization. It was a staged narrowing process: first identify viable electromagnetic families, then select the center feed, then combine it with the ecartometer ring, then test reflector integration, then lock down selectable polarization and manufacturable heights. Run IDs are kept as trace tags only; the design logic is easier to read by objective.

Objective gate	Trace source	Engineering question	Outcome
Patch feasibility	Run1-Run3	Can a simple patch, dual-feed patch, or one-plane cluster serve the payload and sensors?	Useful baselines, but band-edge match, CP margin, or cluster symmetry remained weak.
Matching and sensor ring	Run4	Can feed matching or symmetry refinement fix the center and sensors?	Matching network did not beat the baseline; the R130 mm ecartometer ring became the tracking reference.
Stacked center feed	Run5A-Run7	Does a passive upper patch improve the center feed without unbuildable complexity?	Yes. The passive stacked center became the selected center-feed family.
Architecture compile	Run8	What is the best combined architecture from available solved data?	Selected passive center plus selected ecartometer ring, with an explicit caveat that the exact hybrid was not fully solved.
Reflector integration	Run9-Run10	Where does the feed sit in the dish, and what gain/blockage/tolerance tradeoffs matter?	F/D 0.45 carried through; school dish works as baseline, 1.50 m reflector ranks higher if allowed.
CP selection	Run11-Run11B	How do RHCP/LHCP feed choices map into hardware?	Three-feed selectable option found: common -Y, switch +X/-X, amp 0.85, phase +90 deg.
Manufacturable center	Run12	Can height distribution improve the selected passive stack?	Yes. Final center dimensions moved to R35.75, h 3.50 mm, gap 5.00 mm.
Blueprints	V3-V5A	How do RF coordinates become CAD/KiCad/Fusion drawings?	V5A is the latest integrated package with separated components and validated windows/placement manifests.

How the Main Body Is Ordered

The main body does not send the reader to a graph dump. It starts with the first antenna baseline, then follows each rejected and accepted branch in order. Within each objective, the reference geometry appears before the plots, and every graph is introduced with the reason it matters to the next design decision.

Figure 1 gives the reader the full path before the detailed evidence starts: baseline patch work, CP and ecartometer separation, passive-stack selection, reflector checks, selectable CP, and then V5A packaging.

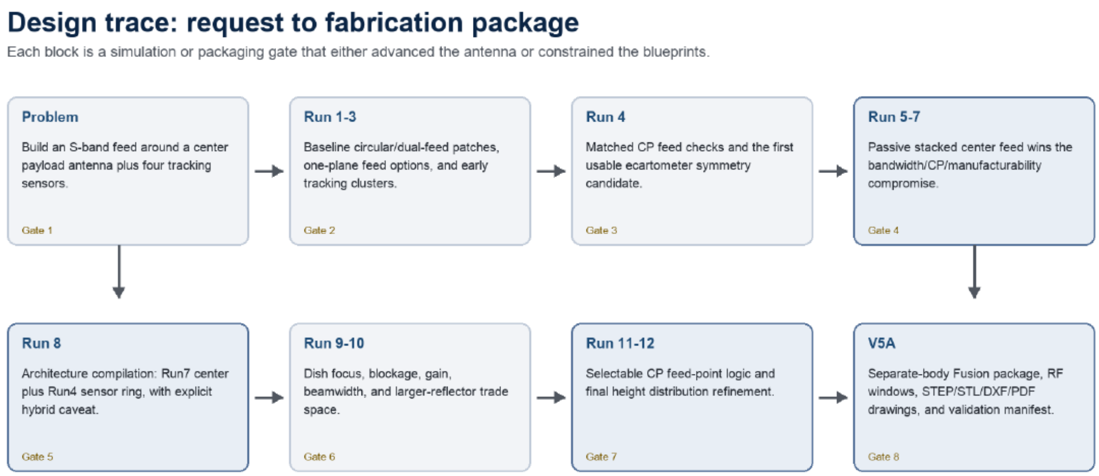

Figure 1. End-to-end trace map from the first antenna question to V5A fabrication assets.

Design Story and Evidence

The rest of the report follows the design in the same order the decisions were made. Each objective starts with the question we were trying to answer, shows the reference geometry first, then shows the plots that changed the next decision. The run number is only a trace source.

Figure 2 is the high-level center-feed thread: the driven-only reference gets replaced by the passive stack, and the final height-distribution sweep tightens the build dimensions before V5A.

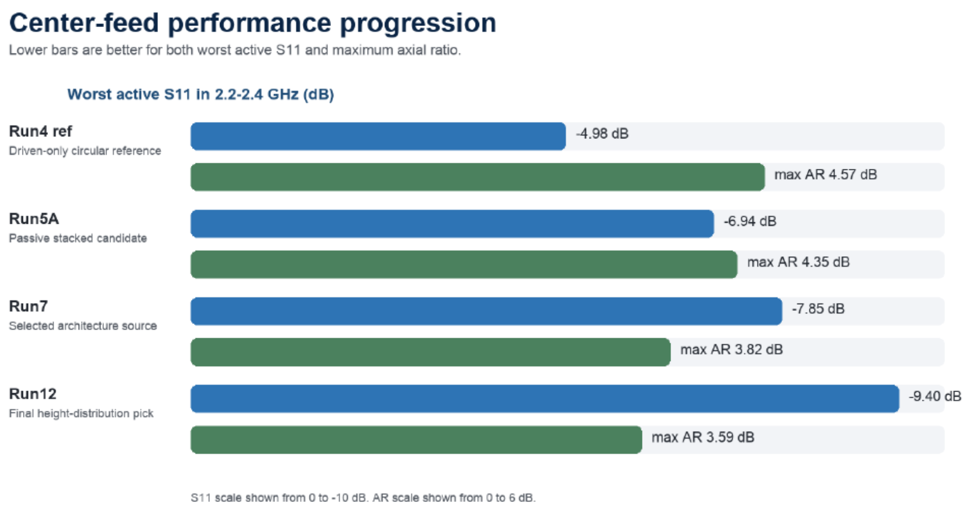

Figure 2. Story overview: center-feed progression from driven-only reference to the final height distribution.

1. Single-Patch Resonance Baseline

The design question was: Can the simplest circular patch hit the 2.3 GHz S-band center frequency?

We tried it because we needed a known-good electromagnetic baseline before spending effort on circular polarization, sensors, or mechanical packaging.

Trace source: Run 1 / circular_patch_s11_sweep

Scope

- Swept circular patch radius and feed offset over an air dielectric above a common ground plane.
- Judged both center-frequency match and worst-case match over 2.2-2.4 GHz.
- Generated heatmaps to show whether the answer was a broad design region or a narrow tuning point.

Evidence and Interpretation

In Figure 3, this is the first physical reference: a single circular patch over ground. It establishes the radius/feed-offset scale used by later CP and stack searches. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the simplest circular patch hit the 2.3 GHz S-band center frequency?

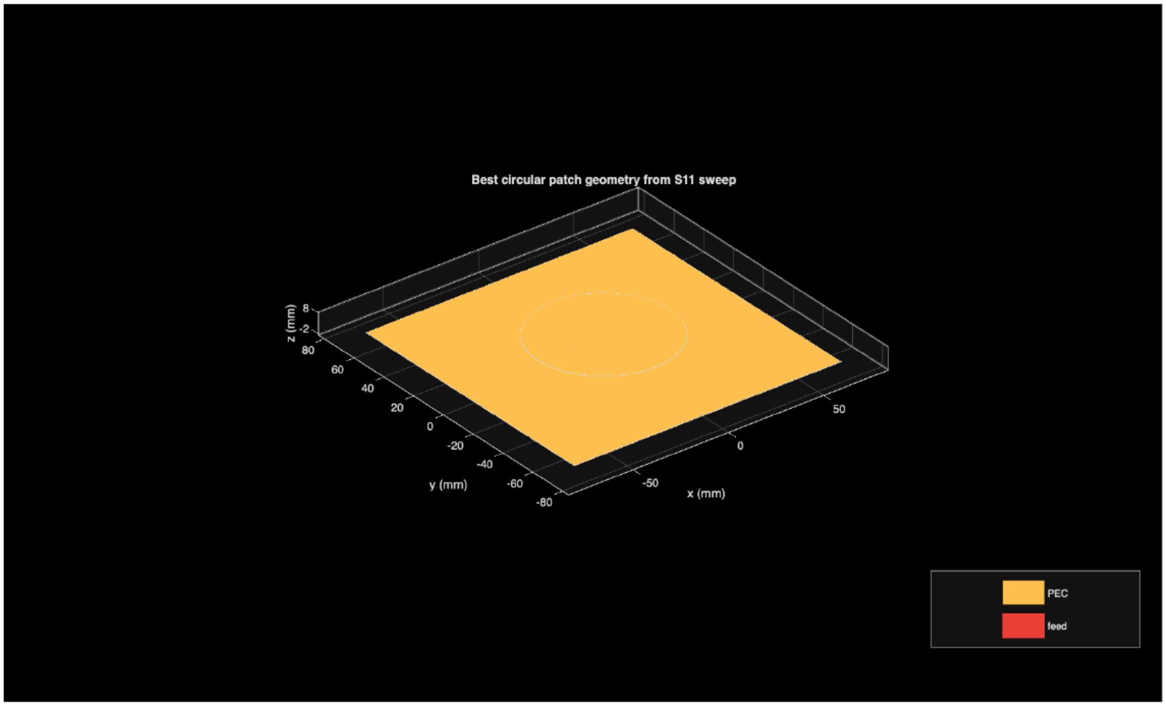

Figure 3. Single-Patch Resonance Baseline - Reference geometry: Best patch geometry.

In Figure 4, this heatmap shows that the 2.3 GHz match was tunable rather than accidental. The useful region is visible, but it only answers the center-frequency question. Read against the section objective, it supports the decision at the end of this block: The patch can match very deeply at 2.3 GHz, but band-edge S11 was only about -6.4 dB, so a single-feed patch was not enough.

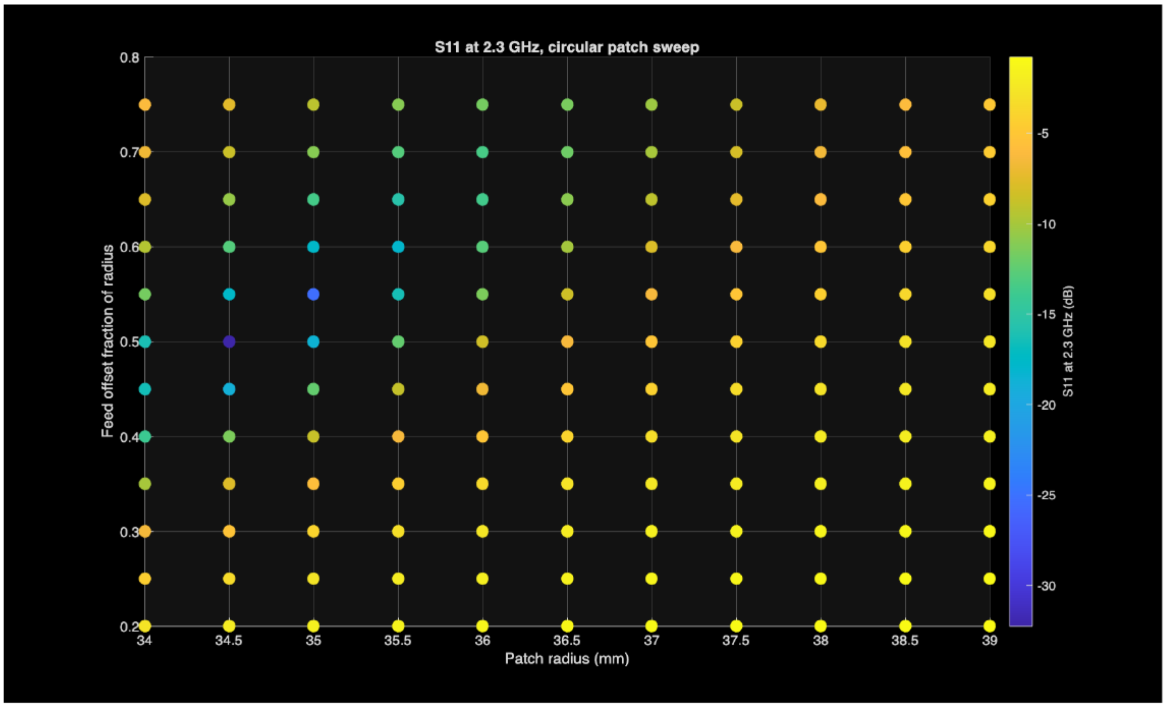

Figure 4. Single-Patch Resonance Baseline - Graph evidence: Heatmap s11 at 2.3GHz.

In Figure 5, this plot is why the simple patch could not be frozen: the center frequency can look good while the 2.2-2.4 GHz band edges remain weak. Read against the section objective, it supports the decision at the end of this block: The patch can match very deeply at 2.3 GHz, but band-edge S11 was only about -6.4 dB, so a single-feed patch was not enough.

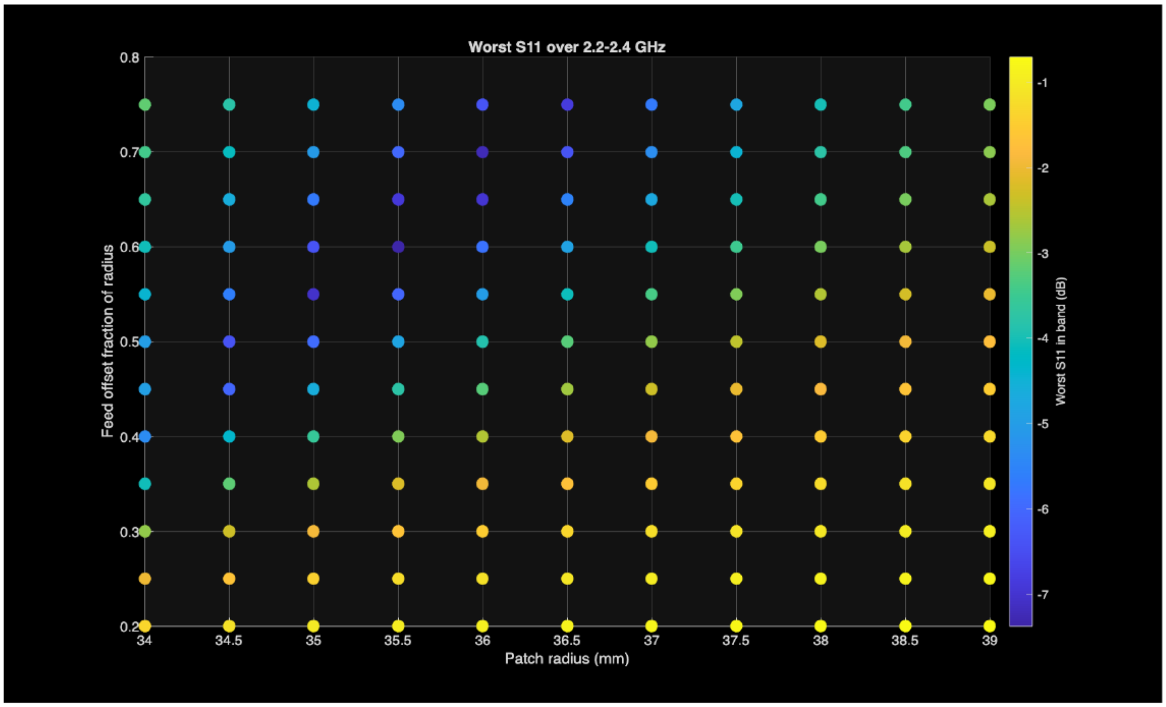

Figure 5. Single-Patch Resonance Baseline - Graph evidence: Heatmap worst s11 band.

Design decision: The patch can match very deeply at 2.3 GHz, but band-edge S11 was only about -6.4 dB, so a single-feed patch was not enough. Next step: Keep R34.5 mm as the first center-frequency anchor and test dual-feed CP excitation.

2. Dual-Feed Circular Polarization Probe

The design question was: Can two orthogonal feeds on the same circular patch create usable circular polarization?

We tried it because the payload feed needed CP behavior, so we tested whether a direct two-port circular patch could produce it before adding stacked layers.

Trace source: Run 1 / dualfeed_cp_patch

Scope

- Swept radius, feed radius fraction, and +/-90 deg phase excitation.
- Measured active S11 and axial ratio across the 2.2-2.4 GHz band.

- Kept the geometry intentionally simple so CP behavior could be understood before mechanical complexity entered.

Evidence and Interpretation

In Figure 6, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can two orthogonal feeds on the same circular patch create usable circular polarization?

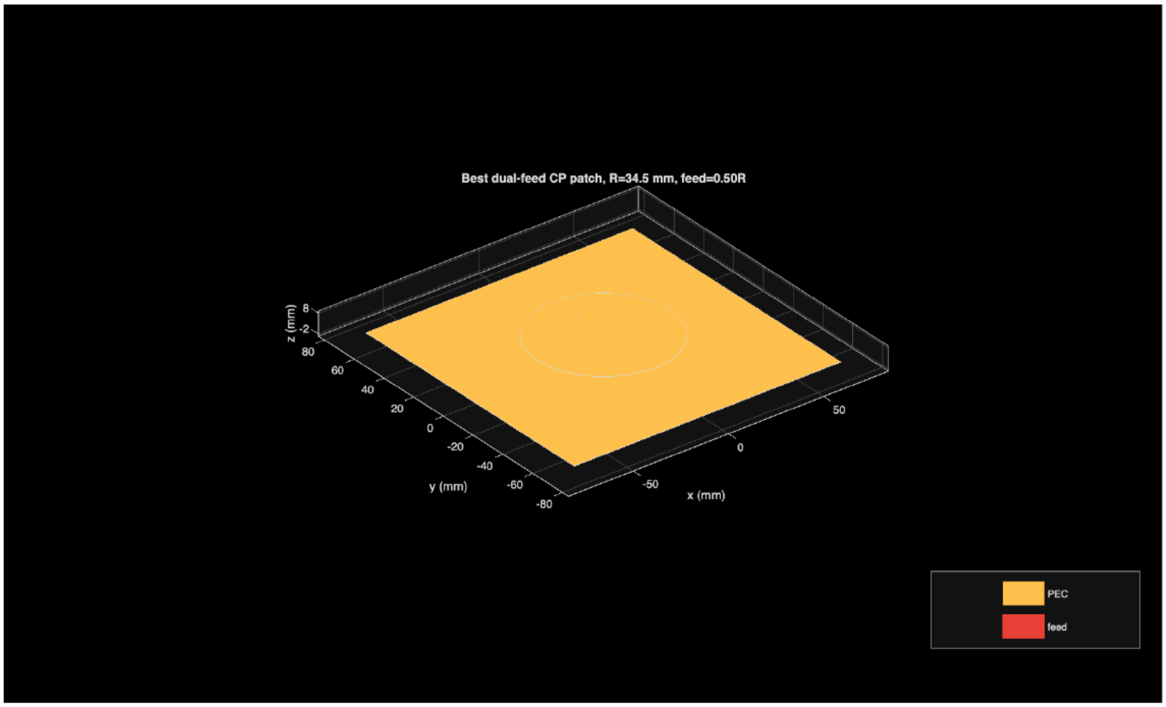

Figure 6. Dual-Feed Circular Polarization Probe - Reference geometry: Best geometry.

In Figure 7, the active S11 plot shows whether the chosen excitation is actually usable as an antenna input across the band, not just at one frequency. Read against the section objective, it supports the decision at the end of this block: CP was real, but worst active S11 and max axial ratio did not close the full band margin.

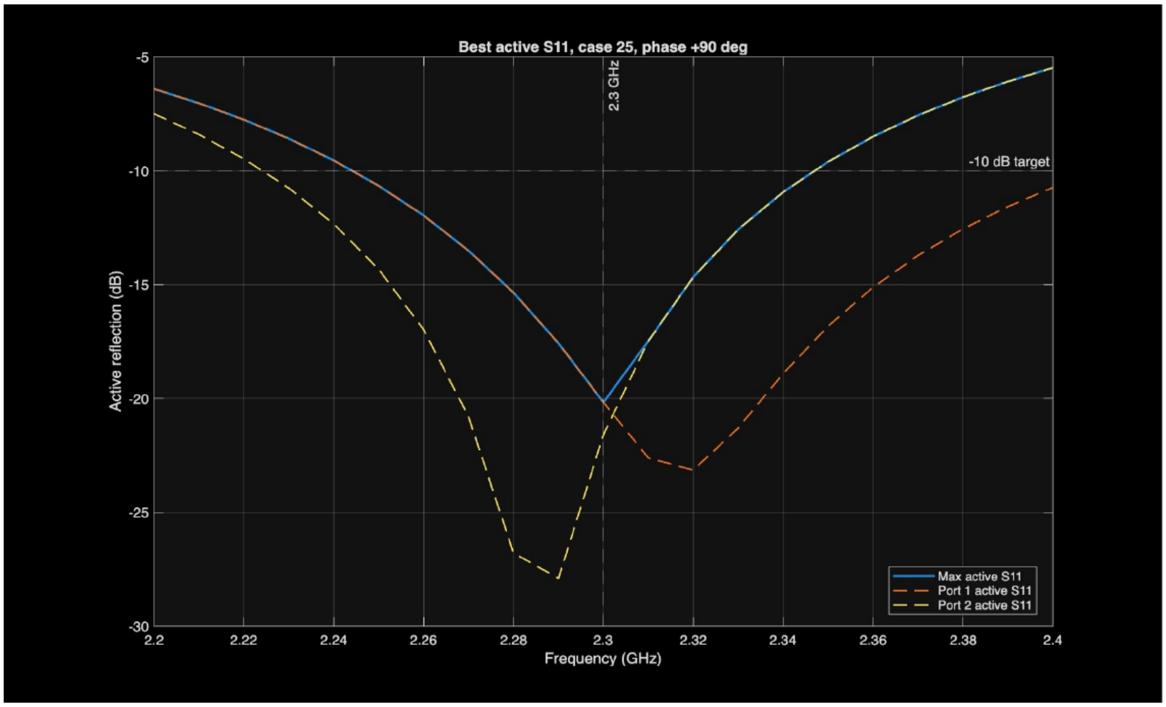

Figure 7. Dual-Feed Circular Polarization Probe - Graph evidence: Best active s11.

In Figure 8, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: CP was real, but worst active S11 and max axial ratio did not close the full band margin.

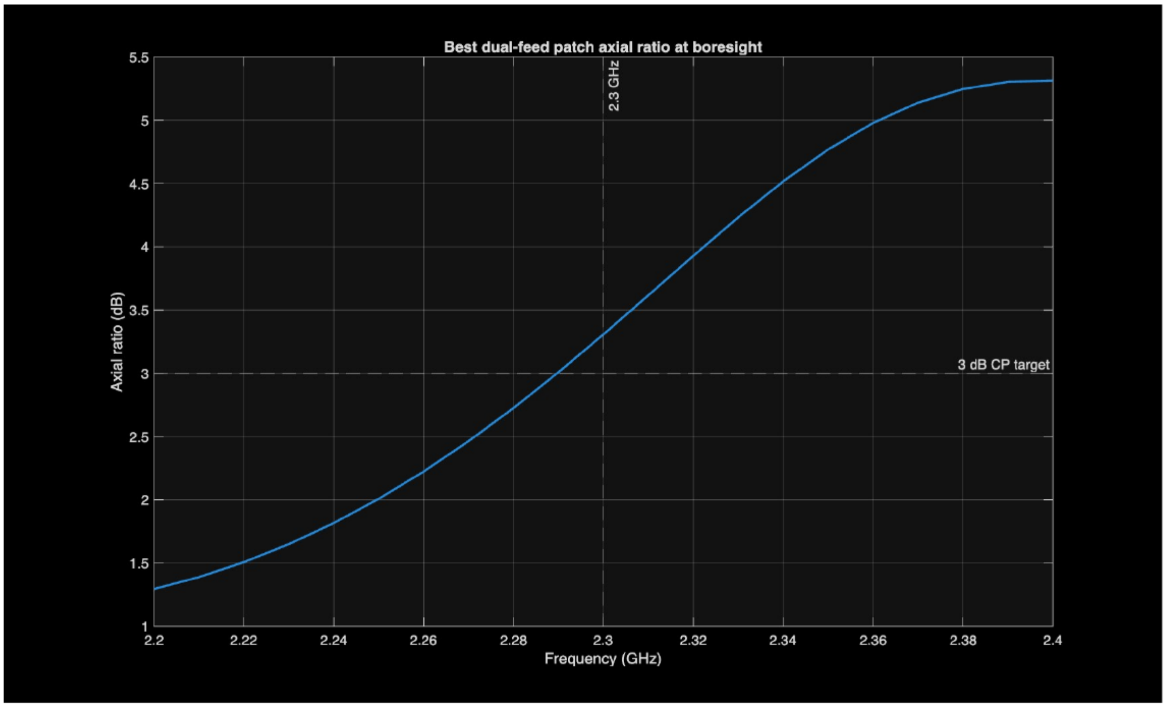

Figure 8. Dual-Feed Circular Polarization Probe - Graph evidence: Best axial ratio.

In Figure 9, the two-port response confirms the orthogonal-feed CP experiment was electrically plausible before we judged active match and axial ratio. Read against the section objective, it supports the decision at the end of this block: CP was real, but worst active S11 and max axial ratio did not close the full band margin.

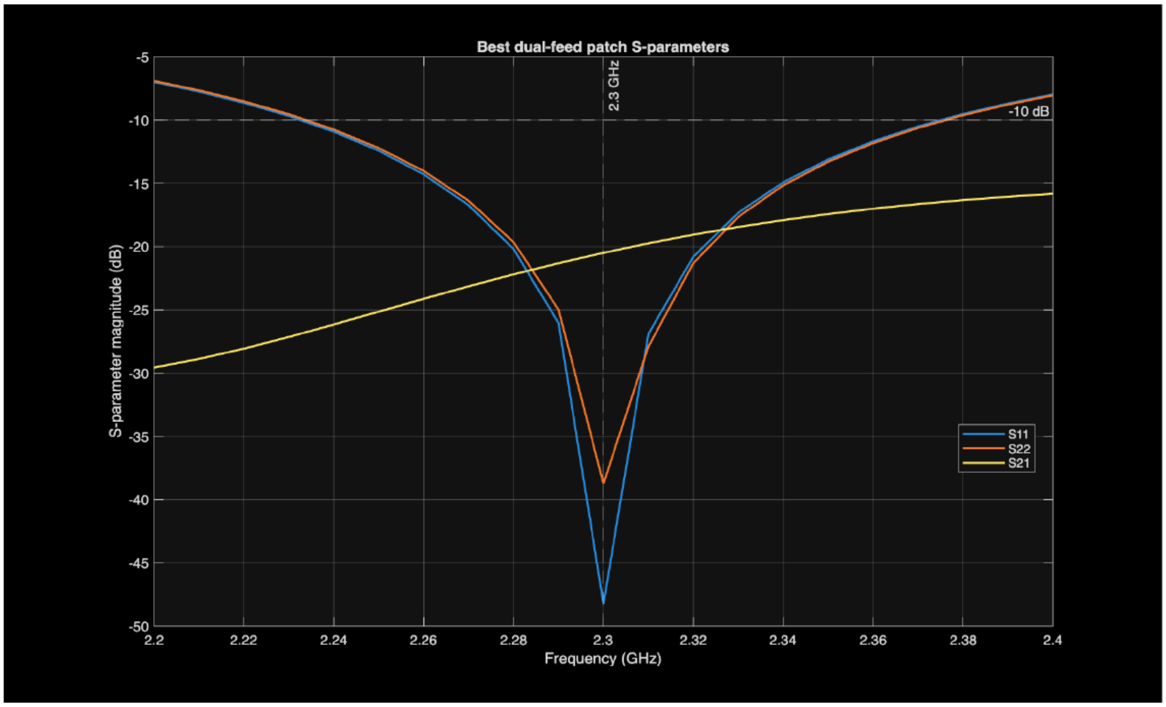

Figure 9. Dual-Feed Circular Polarization Probe - Graph evidence: Best sparams.

Design decision: CP was real, but worst active S11 and max axial ratio did not close the full band margin. Next step: Use the dual-feed result as the CP baseline and test whether a parasitic upper patch could broaden the useful behavior.

3. First Stacked/Parasitic Bandwidth Attempt

The design question was: Can a floating upper patch broaden the center feed without abandoning the circular-patch family?

We tried it because the single and dual-feed patches were too narrow at the band edges; a passive parasitic layer was the simplest bandwidth lever.

Trace source: Run 1 / stacked_patch_wideband

Scope

- Swept lower radius, upper radius, and inter-patch spacing over a 2.0-2.6 GHz wide frequency span.
- Compared the selected stacked case against the single-patch baseline.
- Tracked coverage below -10 dB as the practical measure of bandwidth improvement.

Evidence and Interpretation

In Figure 10, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can a floating upper patch broaden the center feed without abandoning the circular-patch family?

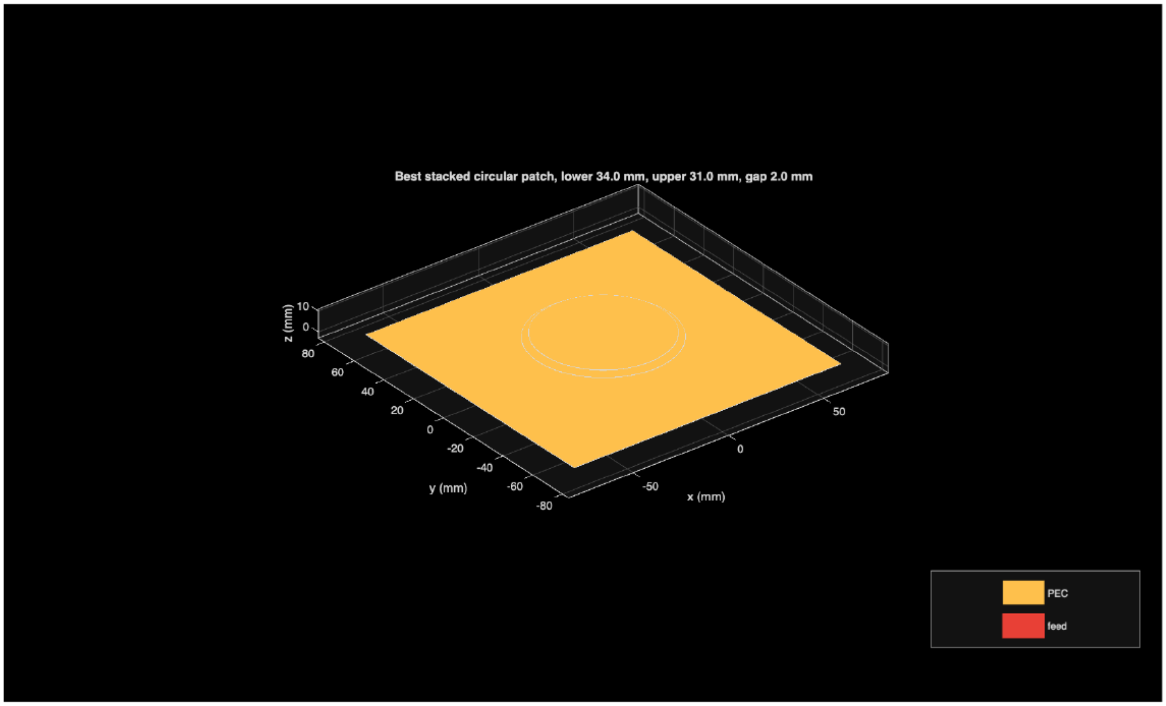

Figure 10. First Stacked/Parasitic Bandwidth Attempt - Reference geometry: Best geometry.

In Figure 11, the wideband sweep shows whether the stacked response is broad enough to be worth carrying forward instead of being a narrow tuned artifact. Read against the section objective, it supports the decision at the end of this block: The passive idea improved coverage but was not yet optimized for CP or mechanical stack height.

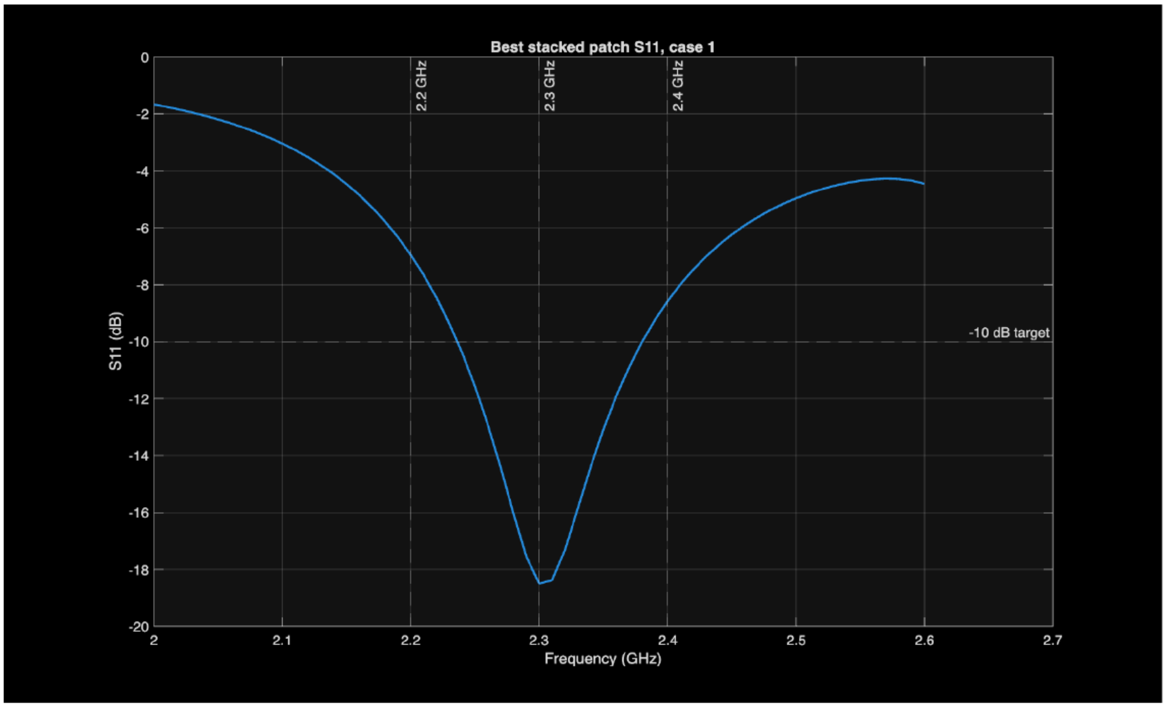

Figure 11. First Stacked/Parasitic Bandwidth Attempt - Graph evidence: Best s11 wide.

In Figure 12, this direct comparison is the first evidence that a passive upper patch can broaden the useful response relative to the single-patch baseline. It is included here because it affected the next step in the design path: Carry the passive-stack concept forward, but first check whether sensors can coexist with the feed.

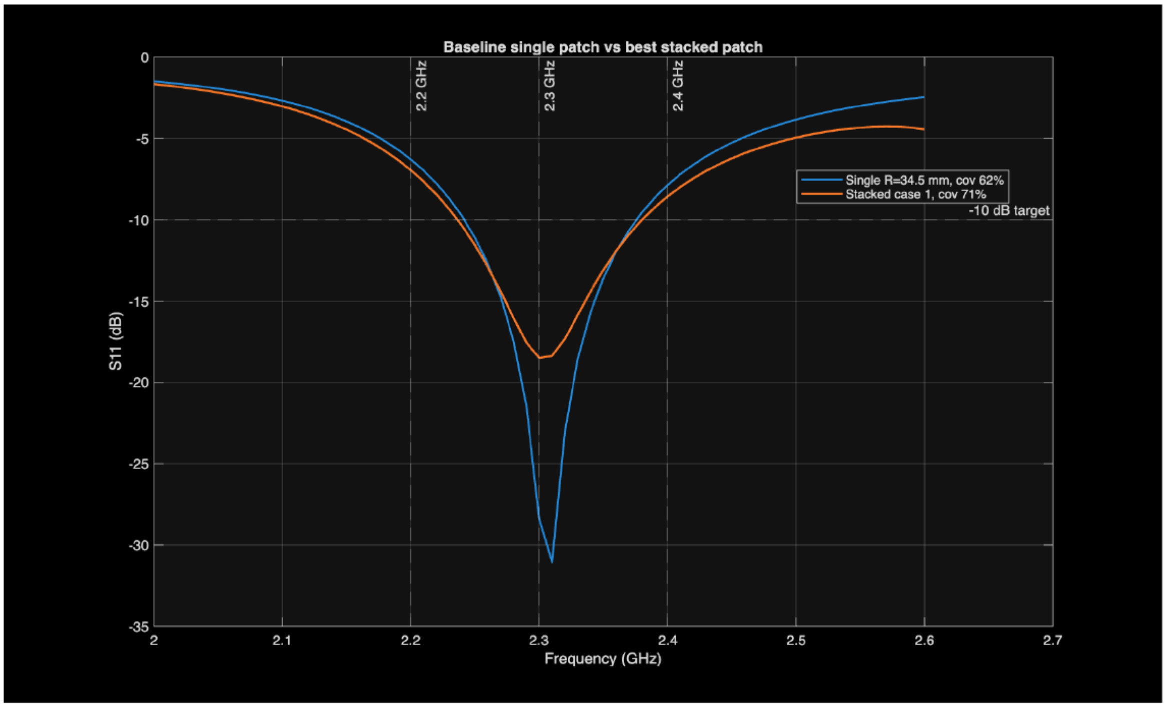

Figure 12. First Stacked/Parasitic Bandwidth Attempt - Generated evidence: Comparison single vs stacked.

Design decision: The passive idea improved coverage but was not yet optimized for CP or mechanical stack height. Next step: Carry the passive-stack concept forward, but first check whether sensors can coexist with the feed.

4. First Four-Sensor Tracking Cluster

The design question was: Can four nearby sensor patches surround the center feed without ruining match or coupling?

We tried it because the concept needed tracking/elevation error sensing, so we tested a compact four-patch ring around the payload feed early.

Trace source: Run 1 / 4patch_tracking_cluster

Scope

- Swept sensor patch radius and sensor-center radius around the R34.5 mm center patch.
- Measured center match degradation, center-to-sensor coupling, and copper gap.
- Saved coupling matrices to expose whether symmetry and isolation were plausible.

Evidence and Interpretation

In Figure 13, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can four nearby sensor patches surround the center feed without ruining match or coupling?

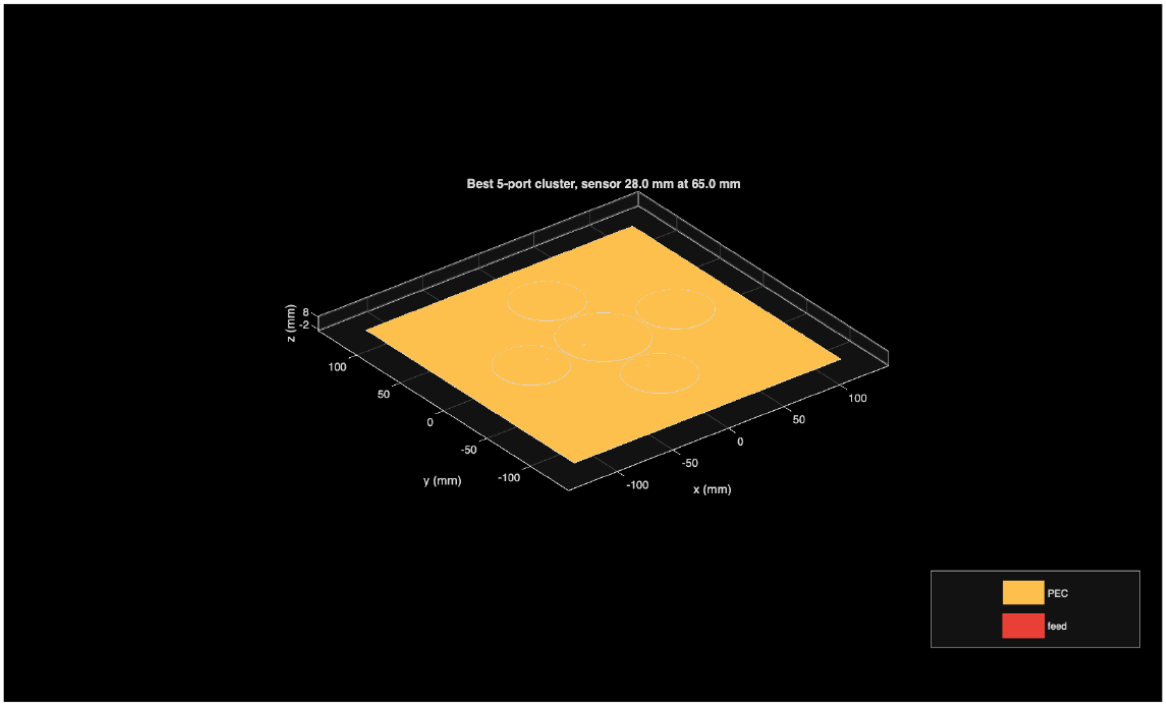

Figure 13. First Four-Sensor Tracking Cluster - Reference geometry: Best geometry.

In Figure 14, this comparison shows how much the surrounding sensor cluster damages the center feed. It is the main reason compact rings were not accepted. Read against the section objective, it supports the decision at the end of this block: The compact cluster was too coupled and too tight mechanically; max center-to-sensor coupling was only about -16.8 dB.

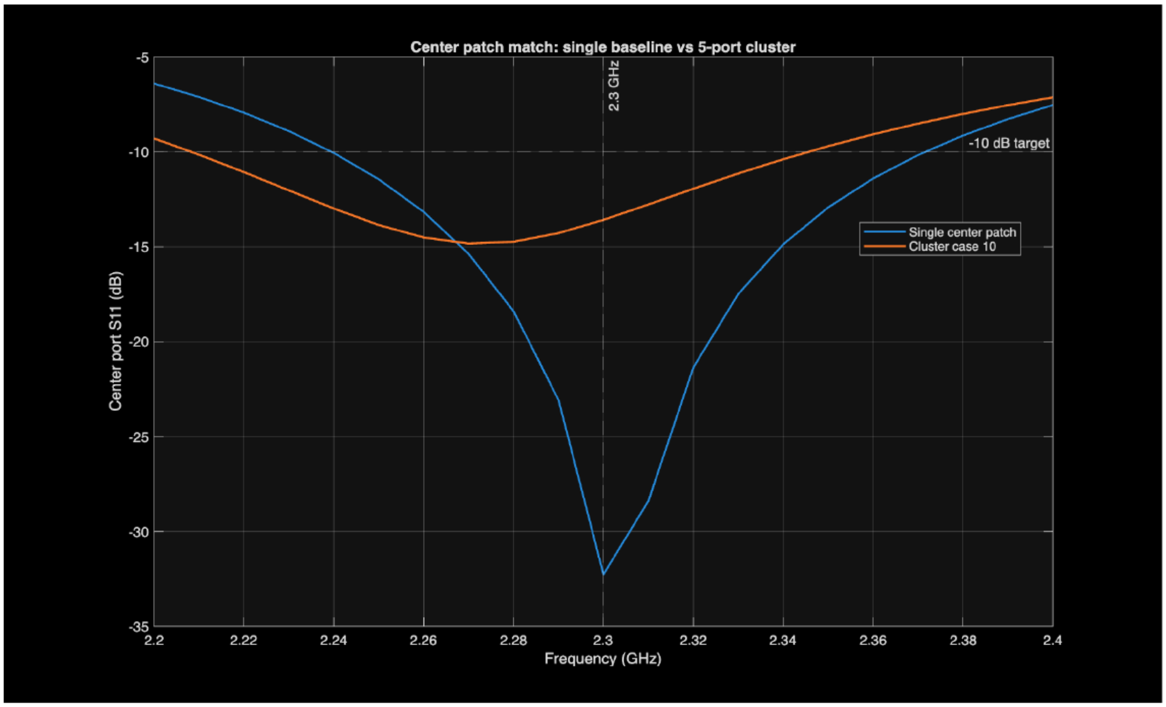

Figure 14. First Four-Sensor Tracking Cluster - Graph evidence: Center s11 comparison.

In Figure 15, the coupling matrix is the sensor-coexistence test: useful tracking sensors need isolation and symmetry, not just individual match. Read against the section objective, it supports the decision at the end of this block: The compact cluster was too coupled and too tight mechanically; max center-to-sensor coupling was only about -16.8 dB.

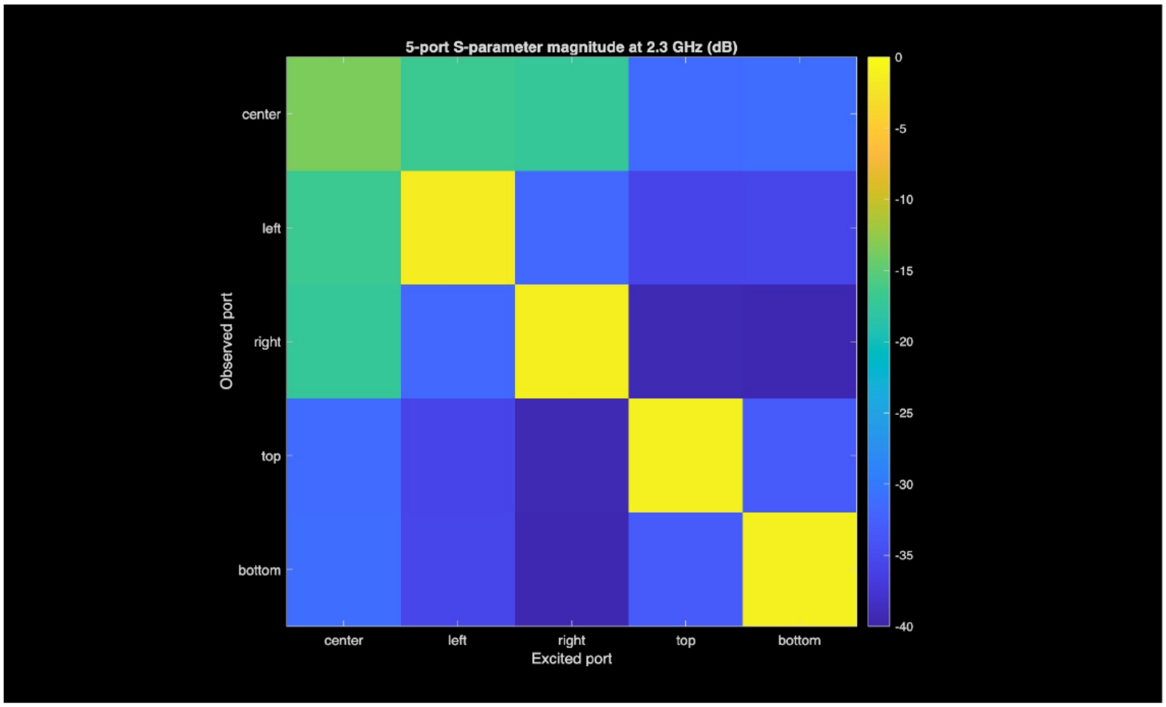

Figure 15. First Four-Sensor Tracking Cluster - Graph evidence: Coupling matrix 2.3GHz.

Design decision: The compact cluster was too coupled and too tight mechanically; max center-to-sensor coupling was only about -16.8 dB. Next step: Increase sensor size/spacing and treat the sensor ring as a separate design variable.

5. One-Plane Payload Feed

The design question was: Can the whole payload feed be made as one top copper layer over one ground plane?

We tried it because a one-plane PCB would be easier to manufacture, so we tested circular, elliptical, dual-circular, and dual-square variants before committing to a stack.

Trace source: Run 2 / oneplane_payload_feed

Scope

- Swept height, equivalent radius, feed fraction, phase, and amplitude ratio for multiple one-plane architectures.
- Used active S11 and axial ratio as hard ranking drivers.

- Kept feed-network routing separate from the radiator result.

Evidence and Interpretation

In Figure 16, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the whole payload feed be made as one top copper layer over one ground plane?

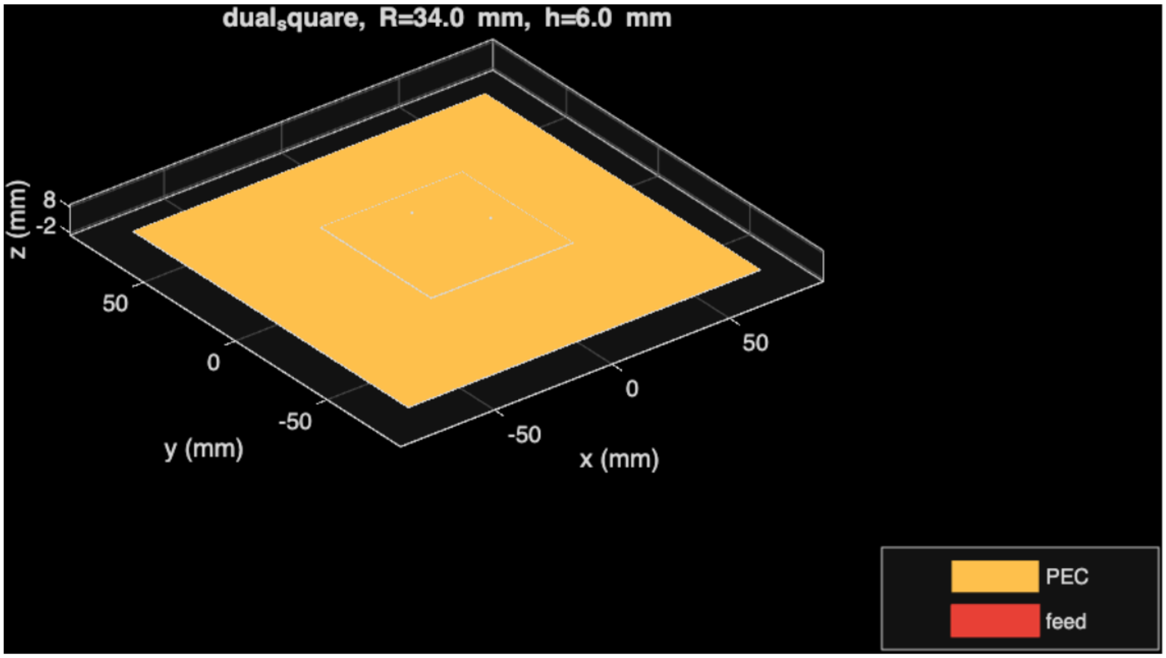

Figure 16. One-Plane Payload Feed - Reference geometry: Best geometry.

In Figure 17, the active S11 plot shows whether the chosen excitation is actually usable as an antenna input across the band, not just at one frequency. Read against the section objective, it supports the decision at the end of this block: The one-plane feed had useful 2.3 GHz behavior but did not satisfy the full-band CP/match target.

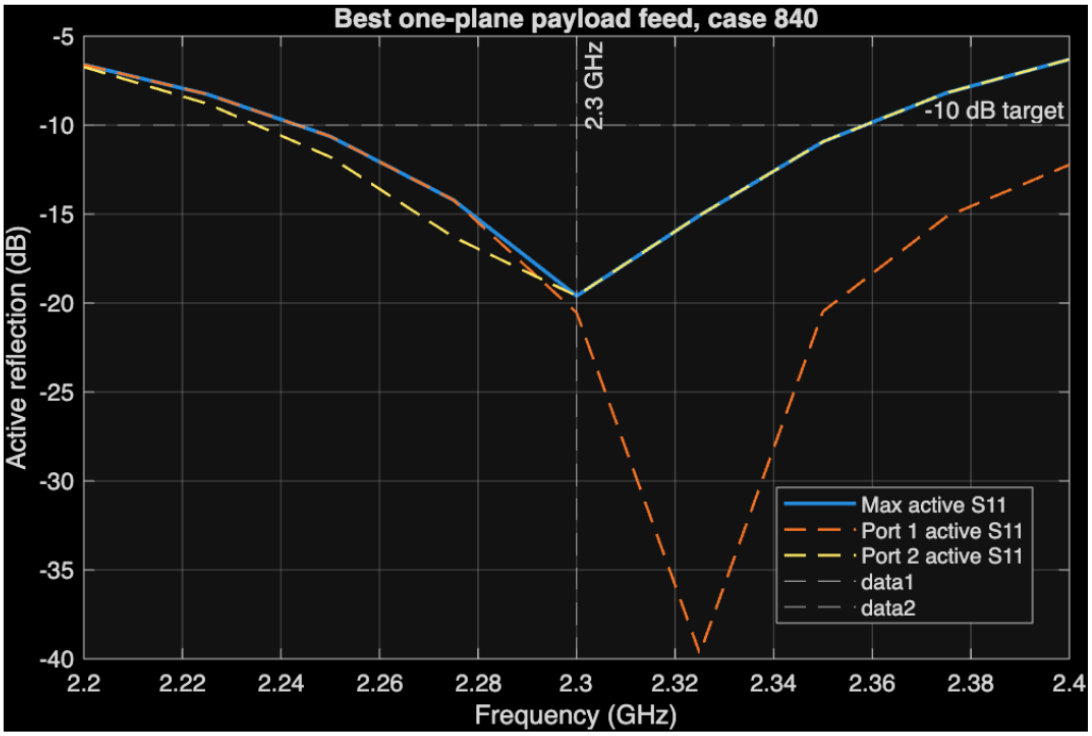

Figure 17. One-Plane Payload Feed - Graph evidence: Best active s11.

In Figure 18, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: The one-plane feed had useful 2.3 GHz behavior but did not satisfy the full-band CP/match target.

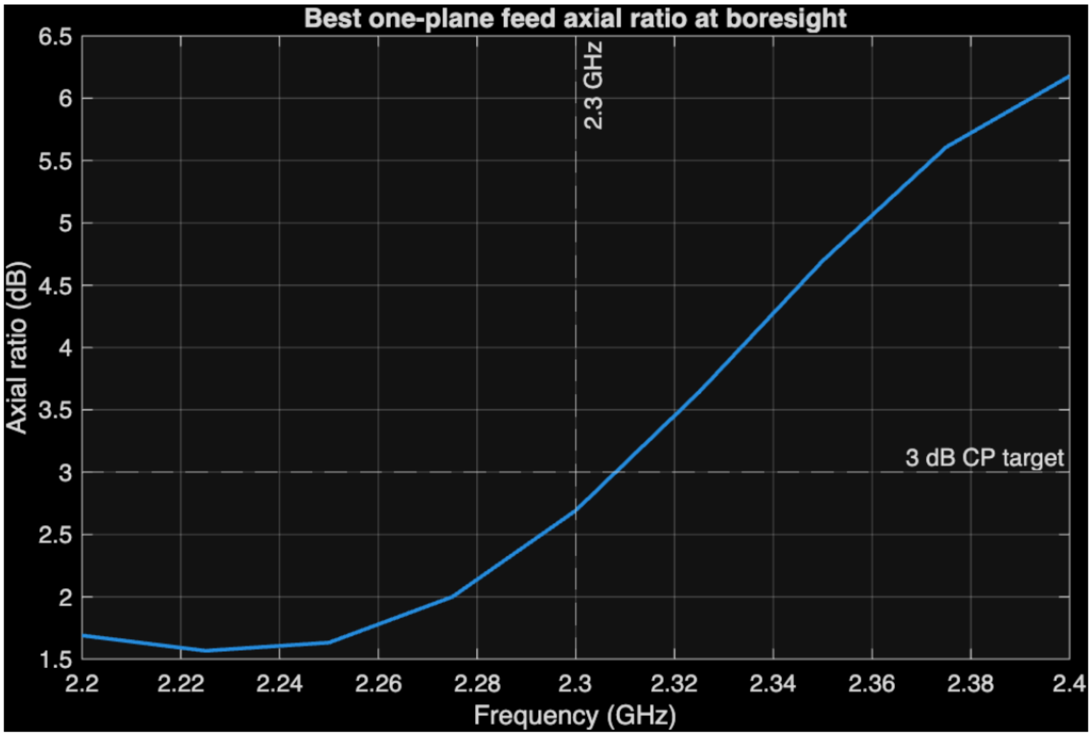

Figure 18. One-Plane Payload Feed - Graph evidence: Best axial ratio.

> **Design decision:** The one-plane feed had useful 2.3 GHz behavior but did not satisfy the full-band CP/match target. Next step: Do not freeze a one-plane payload feed; test whether one-plane sensors can still be useful.

6. One-Plane Feed Network Sizing

The design question was: What PCB feed-line sizes would a one-plane dual-feed concept require?

We tried it because even if the one-plane antenna was not final, we needed to know whether the feed network dimensions were physically reasonable.

Trace source: Run 2 / feed_network

Scope

- Computed closed-form microstrip dimensions because RF Toolbox was unavailable.
- Used the one-plane payload feed active impedances as the load reference.
- Exported 50-ohm, branch-line, and quarter-wave transformer dimensions.

Evidence and Interpretation

This supporting calculation produced sizing data rather than PNG plots; the conclusion is carried in the source traceability list.

> **Design decision:** The numbers were useful for feasibility but too substrate- and model-dependent for final hardware. Next step: Keep feed routing as a later packaging task; focus the RF selection on radiator geometry.

7. One-Plane Four-Sensor Ring

The design question was: Can the sensor ring work on one copper plane if spacing is increased?

We tried it because the first compact cluster was too close; this test asked whether a larger same-plane ring could recover isolation.

Trace source: Run 2 / oneplane_4sensor_cluster

Scope

- Tuned isolated sensor radius/feed first, then swept sensor-center distances from 85 to 130 mm.
- Scored center degradation, sensor match, coupling, symmetry, and copper clearance.
- Compared the cluster against the isolated center baseline.

Evidence and Interpretation

In Figure 19, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the sensor ring work on one copper plane if spacing is increased?

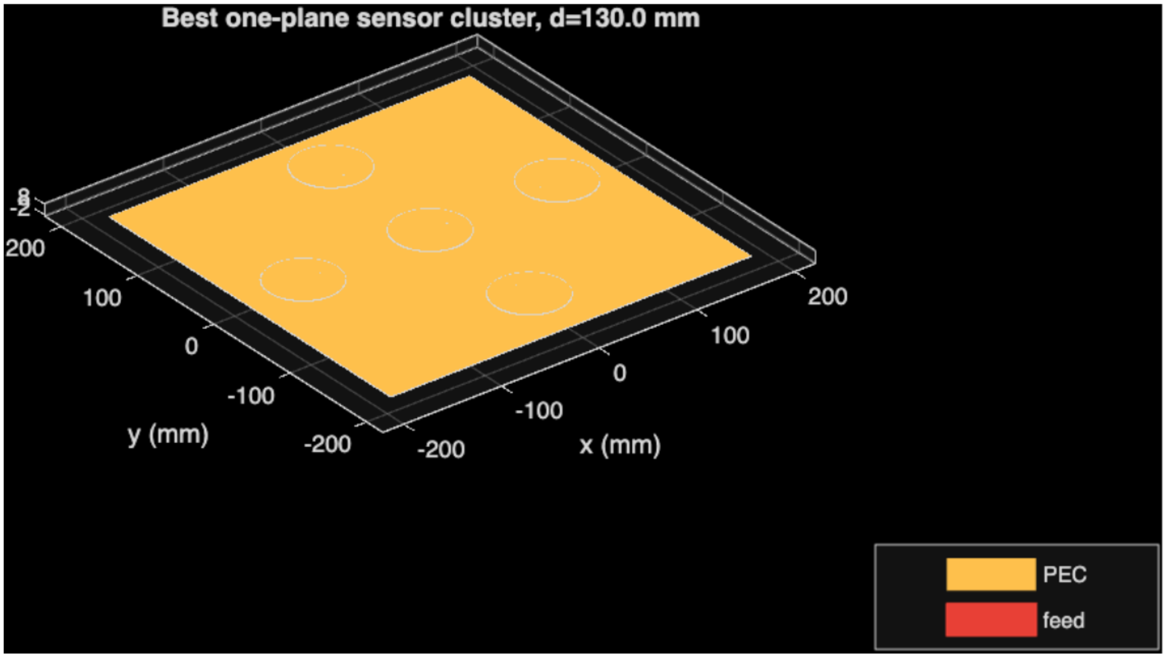

Figure 19. One-Plane Four-Sensor Ring - Reference geometry: Best geometry.

In Figure 20, this comparison shows how much the surrounding sensor cluster damages the center feed. It is the main reason compact rings were not accepted. Read against the section objective, it supports the decision at the end of this block: The larger ring improved isolation enough to become a serious ecartometer direction, even though symmetry still needed work.

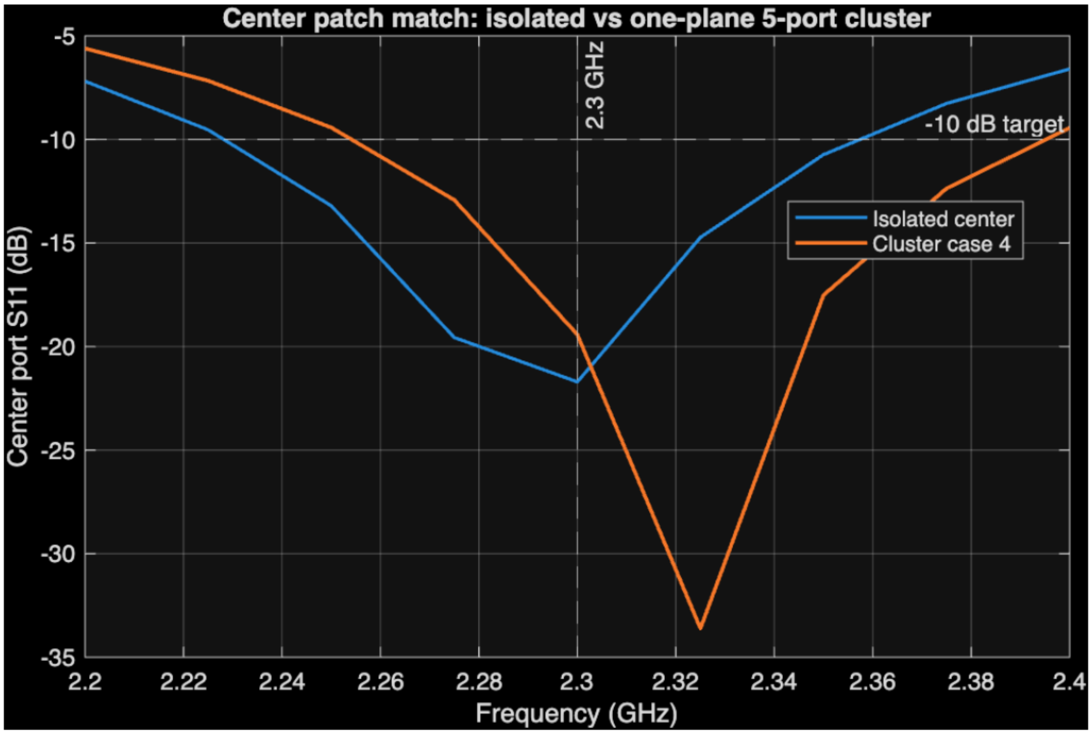

Figure 20. One-Plane Four-Sensor Ring - Graph evidence: Center s11 comparison.

In Figure 21, the coupling matrix is the sensor-coexistence test: useful tracking sensors need isolation and symmetry, not just individual match. Read against the section objective, it supports the decision at the end of this block: The larger ring improved isolation enough to become a serious ecartometer direction, even though symmetry still needed work.

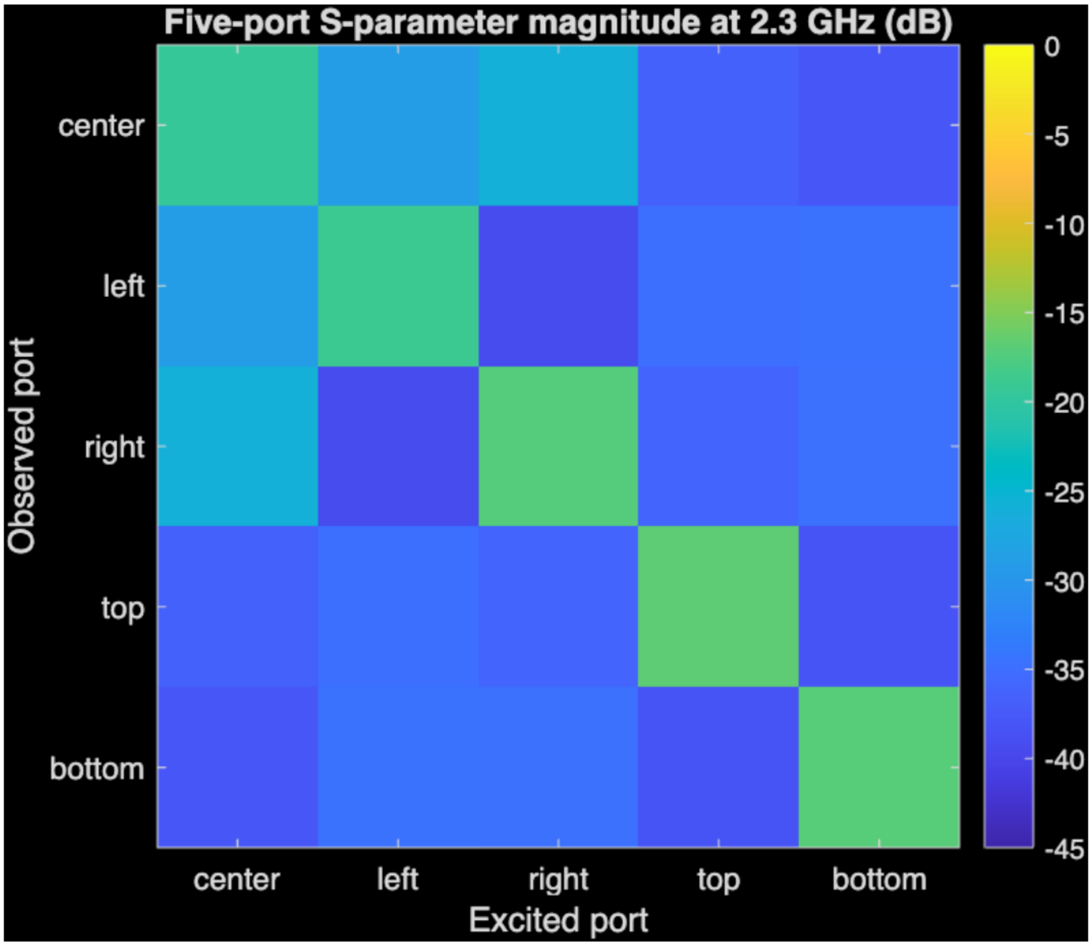

Figure 21. One-Plane Four-Sensor Ring - Graph evidence: Coupling matrix 2.3GHz.

Design decision: The larger ring improved isolation enough to become a serious ecartometer direction, even though symmetry still needed work. Next step: Carry the R130 mm sensor-ring idea into later ecartometer-specific refinement.

8. One-Plane Wideband Feedhead Integration

The design question was: Can one top copper layer solve the center feed and the tracking ring together?

We tried it because before splitting the system into layers/tiles, we tried the simplest integrated mechanical idea: everything on one plane.

Trace source: Run 3 / oneplane_wideband_feedhead

Scope

- Evaluated center-feed candidates, sensor candidates, and cluster distances under one-plane constraints.
- Included slot and annular geometries where the modeling API could represent them.
- Marked hard rejects for lost coverage, bad symmetry, and degraded sensor behavior.

Evidence and Interpretation

In Figure 22, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can one top copper layer solve the center feed and the tracking ring together?

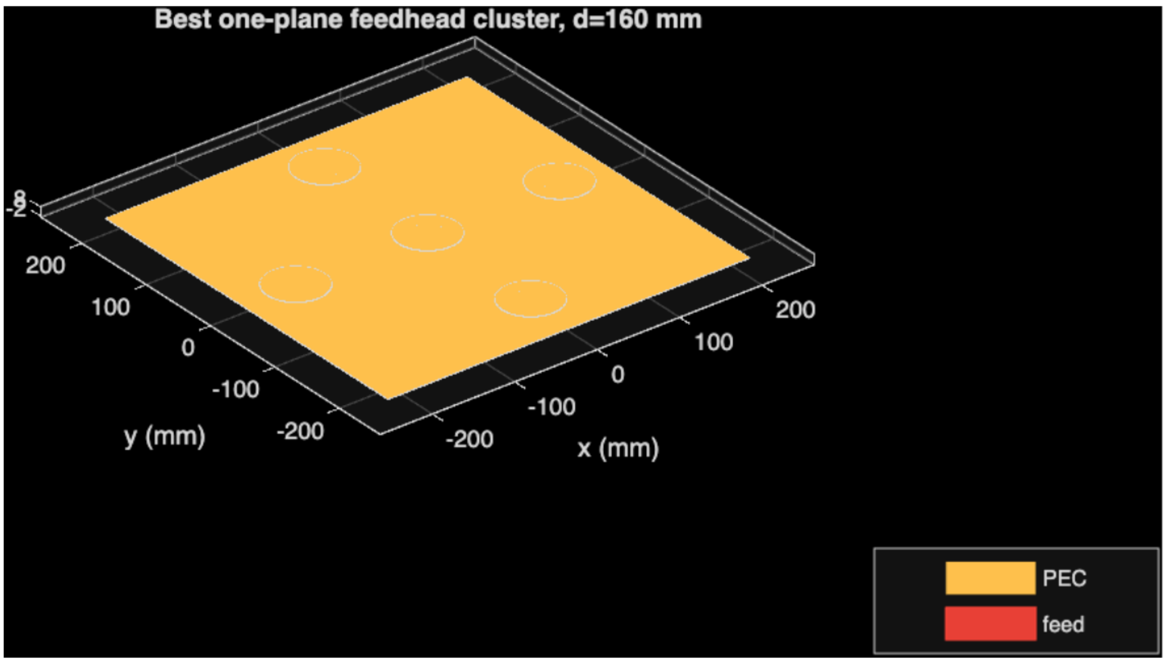

Figure 22. One-Plane Wideband Feedhead Integration - Reference geometry: Best geometry.

In Figure 23, this plot checks whether the center feed still holds its match once the one-plane integration constraints are applied. Read against the section objective, it supports the decision at the end of this block: The integrated one-plane cluster was rejected for degraded center coverage and severe left/right/top/bottom asymmetry.

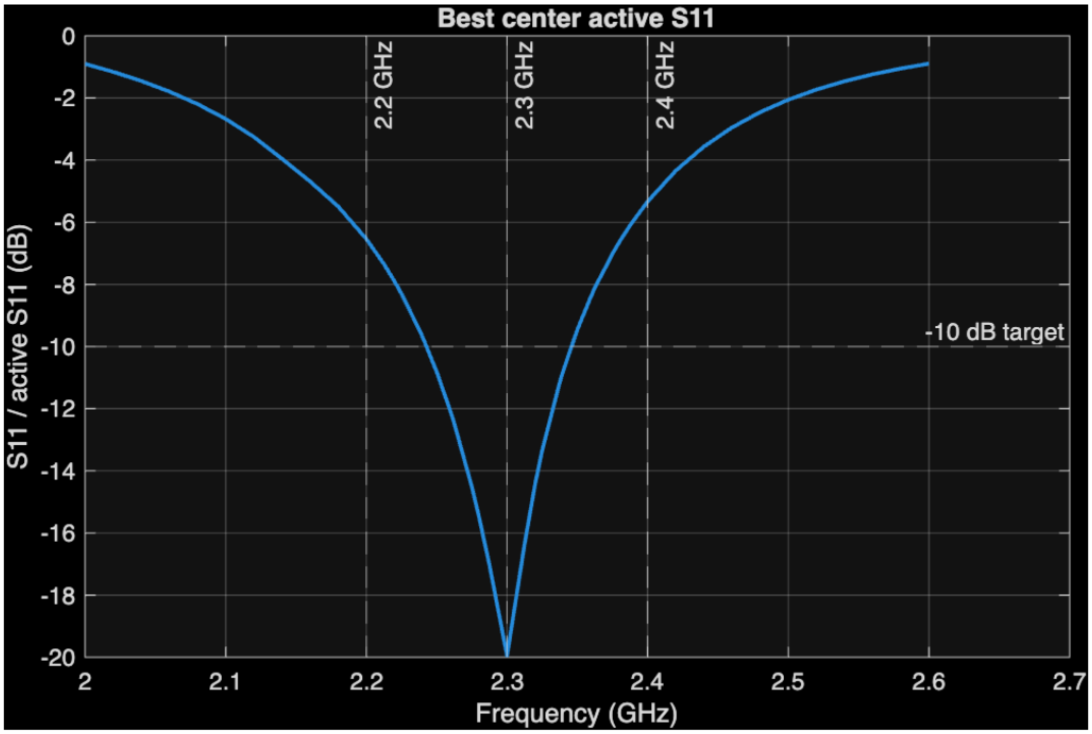

Figure 23. One-Plane Wideband Feedhead Integration - Graph evidence: Best center s11.

In Figure 24, this plot checks whether the sensor element remains tuned before judging it as part of a four-sensor ring. Read against the section objective, it supports the decision at the end of this block: The integrated one-plane cluster was rejected for degraded center coverage and severe left/right/top/bottom asymmetry.

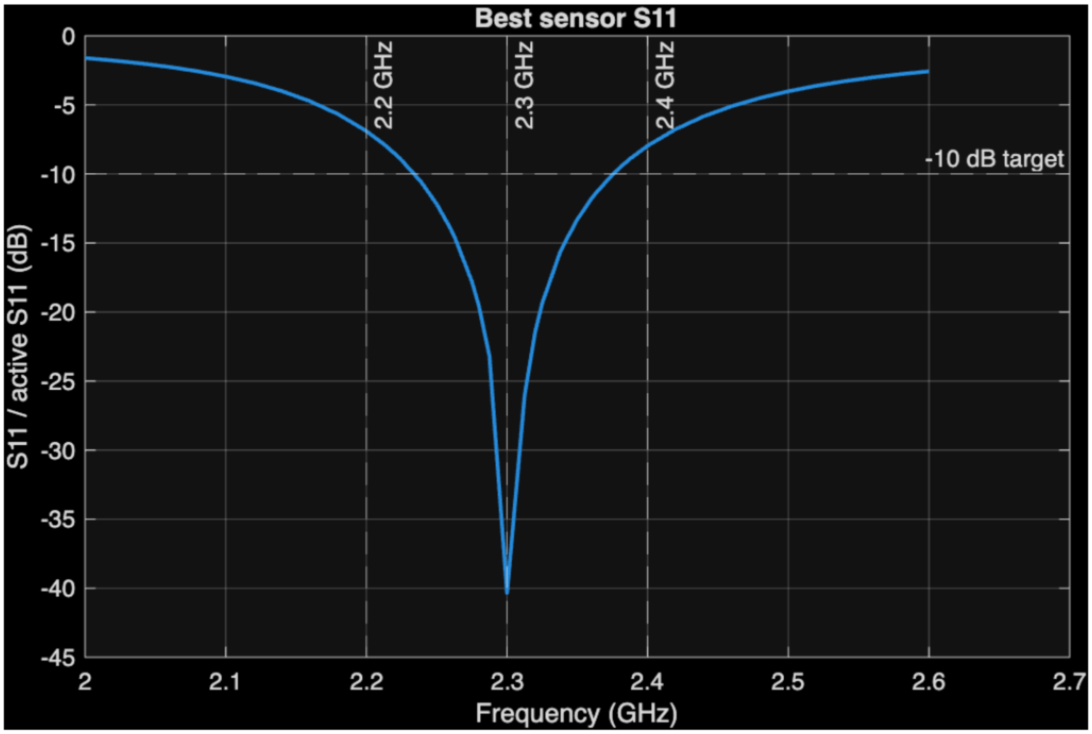

Figure 24. One-Plane Wideband Feedhead Integration - Graph evidence: Best sensor s11.

In Figure 25, this graph is one of the objective checks for one-plane wideband feedhead integration; it supports the conclusion that the integrated one-plane cluster was rejected for degraded center coverage and severe left/right/top/bottom asymmetry. Read against the section objective, it supports the decision at the end of this block: The integrated one-plane cluster was rejected for degraded center coverage and severe left/right/top/bottom asymmetry.

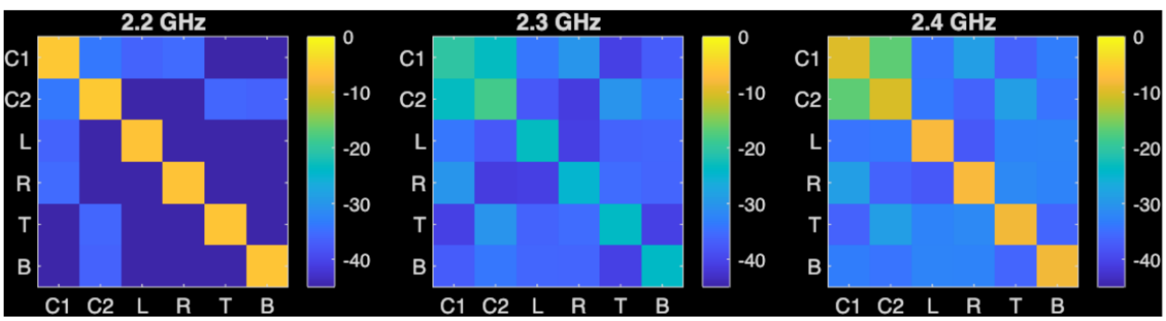

Figure 25. One-Plane Wideband Feedhead Integration - Graph evidence: Cluster coupling 2.2 2.3 2.4.

Design decision: The integrated one-plane cluster was rejected for degraded center coverage and severe left/right/top/bottom asymmetry. Next step: Separate center-feed optimization from ecartometer ring optimization.

9. Matching Network Check

The design question was: Can a practical matching network fix the driven center patch without changing the antenna family?

We tried it because before adding layers, we tested the cheaper electrical fix: matching the existing center feed.

Trace source: Run 4 / center_matching_network

Scope

- Computed port impedance over frequency and checked before/after active S11 and axial ratio.
- Ranked candidate matching dimensions against the active reflection target.
- Kept the same geometry so any improvement could be attributed to matching only.

Evidence and Interpretation

In Figure 26, this geometry is the reference shape for the matching network check objective; the following plots should be read against this physical layout. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can a practical matching network fix the driven center patch without changing the antenna family?

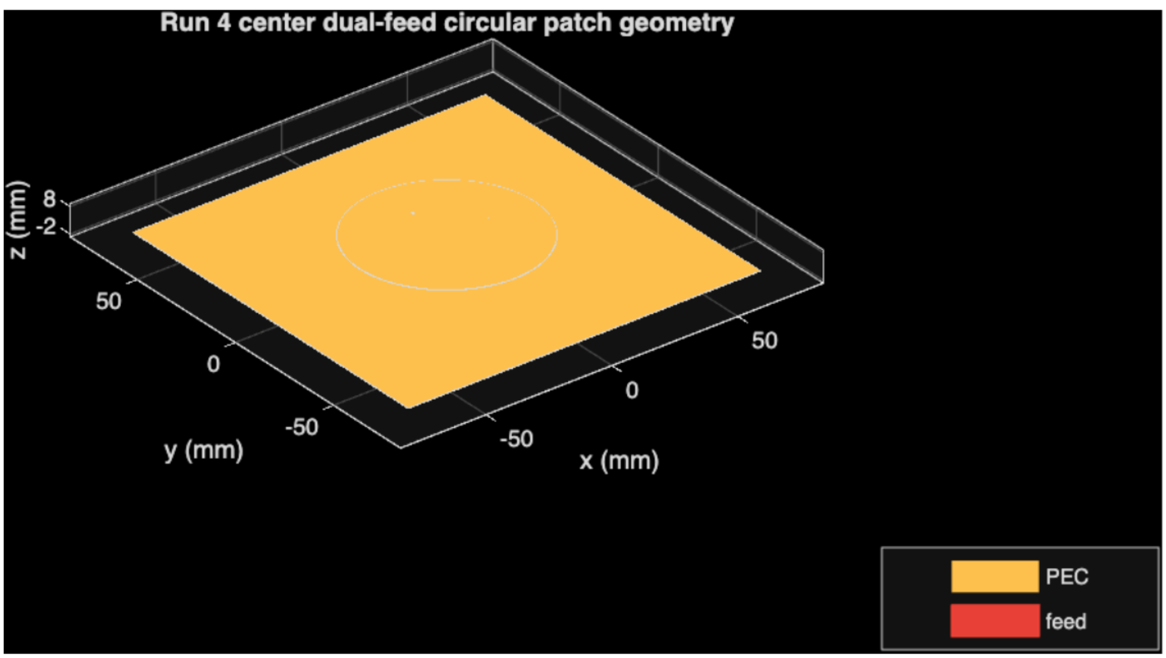

Figure 26. Matching Network Check - Reference geometry: Best center geometry.

In Figure 27, the active S11 plot shows whether the chosen excitation is actually usable as an antenna input across the band, not just at one frequency. Read against the section objective, it supports the decision at the end of this block: The best verified answer did not justify a matching-network freeze; geometry changes still mattered more.

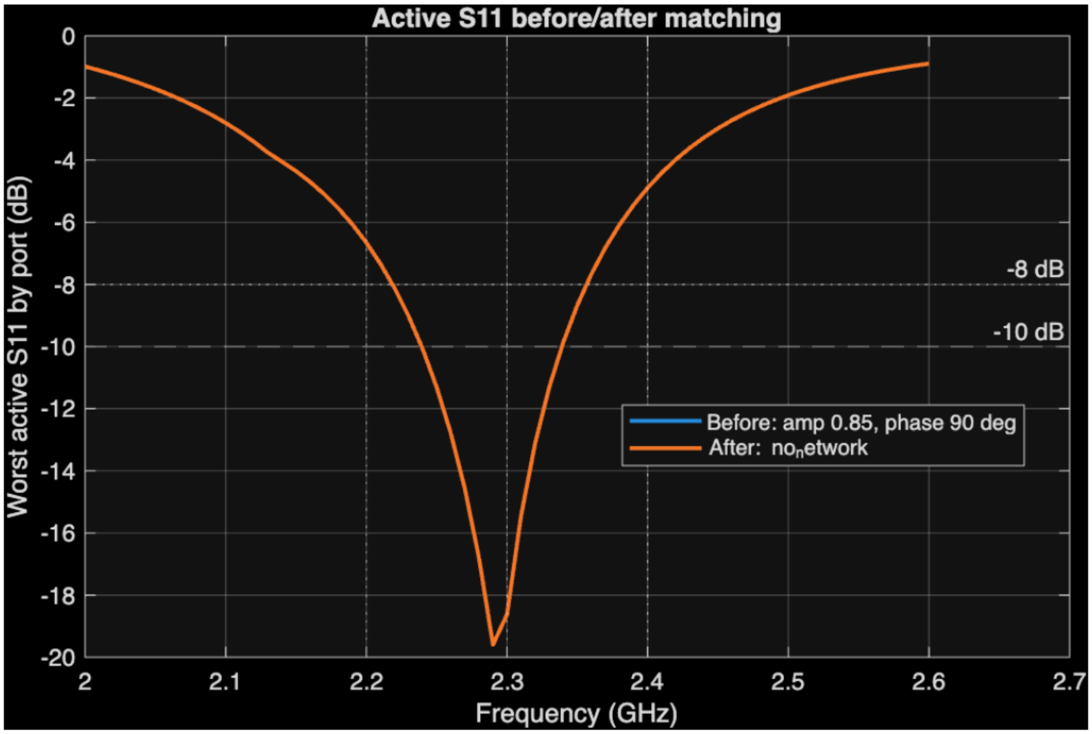

Figure 27. Matching Network Check - Graph evidence: Best active s11 before after.

In Figure 28, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: The best verified answer did not justify a matching-network freeze; geometry changes still mattered more.

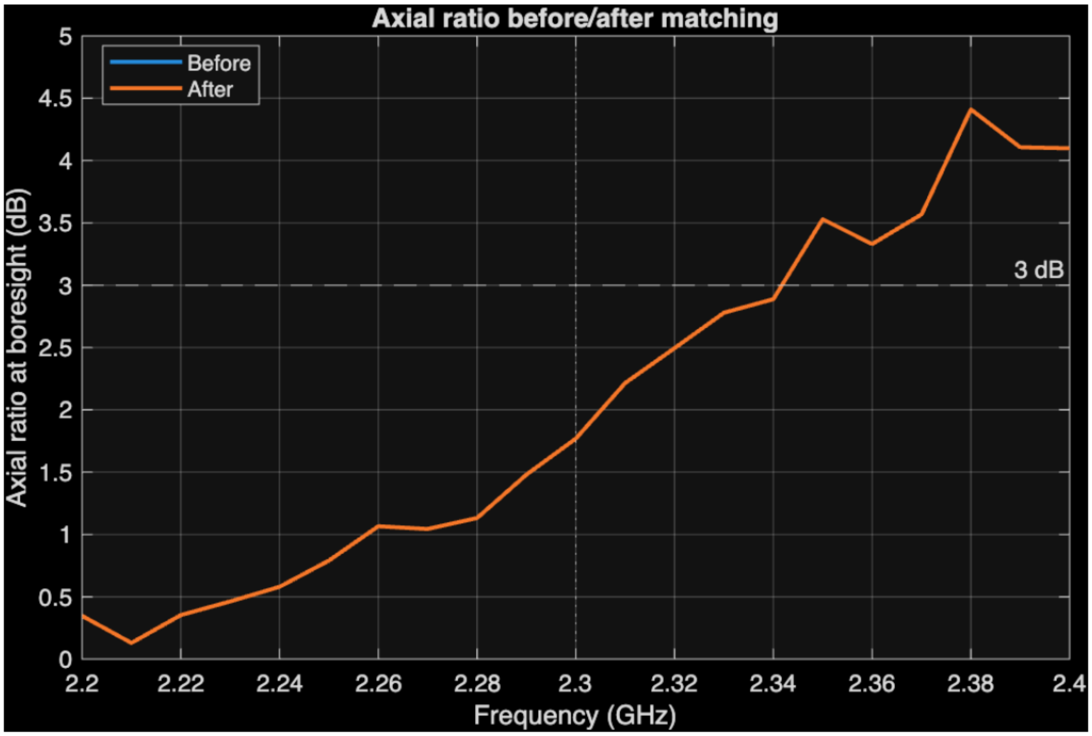

Figure 28. Matching Network Check - Graph evidence: Best axial ratio before after.

In Figure 29, the impedance trace explains why a matching network was considered. It shows the electrical target before asking whether the network is worth the added hardware. Read against the section objective, it supports the decision at the end of this block: The best verified answer did not justify a matching-network freeze; geometry changes still mattered more.

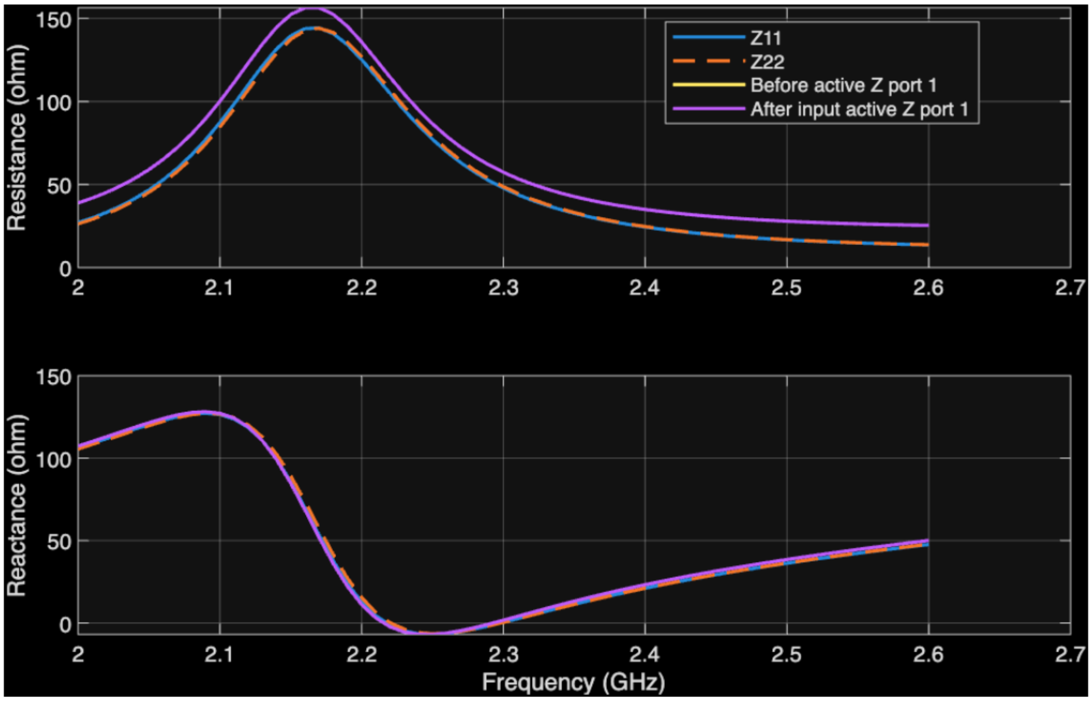

Figure 29. Matching Network Check - Graph evidence: Port impedance vs frequency.

Design decision: The best verified answer did not justify a matching-network freeze; geometry changes still mattered more. Next step: Stop trying to solve the center with matching alone and test geometry perturbations/stacks.

10. One-Plane CP Slot Perturbations

The design question was: Can slots or shape perturbations produce better CP on one layer?

We tried it because slots are a common single-layer CP tool, so this was the alternative to stacked patches.

Trace source: Run 4 / cp_slots_oneplane

Scope

- Generated slotted candidate geometries and compared best slot S11/axial-ratio against the baseline.
- Recorded mesh failures separately instead of hiding invalid candidates.
- Used hard rejects for candidates that did not produce stable or useful EM outputs.

Evidence and Interpretation

In Figure 30, this slot geometry records the single-layer CP alternative. It matters because the passive stack was only acceptable after credible one-layer alternatives were tested. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can slots or shape perturbations produce better CP on one layer?

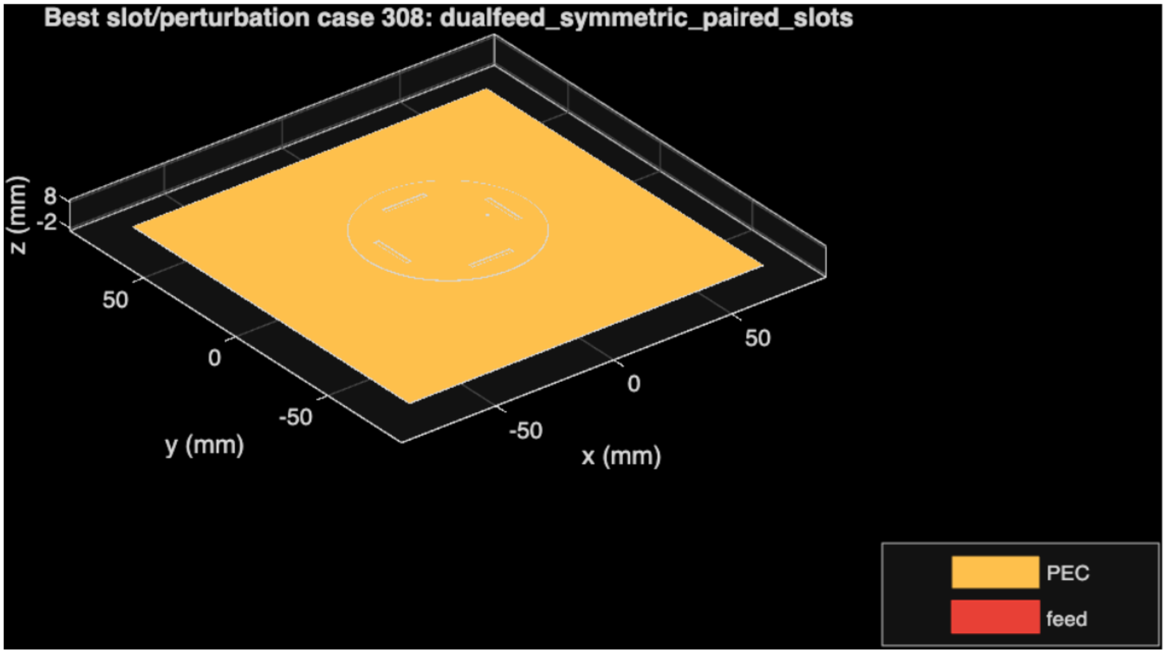

Figure 30. One-Plane CP Slot Perturbations - Reference geometry: Best slot geometry.

In Figure 31, this baseline comparison shows whether slots beat the reference antenna family. The result kept slots in the explored-but-not-selected column. It is included here because it affected the next step in the design path: Keep slots as explored/rejected and concentrate on passive stacked center-feed work.

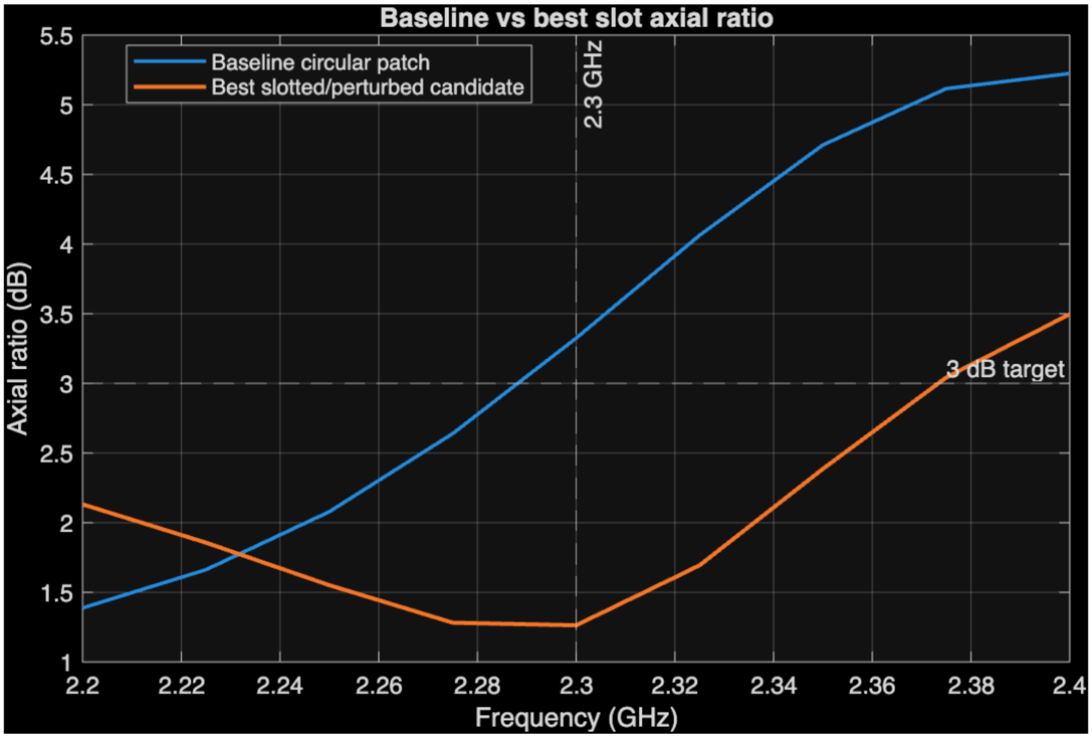

Figure 31. One-Plane CP Slot Perturbations - Generated evidence: Baseline vs best slot ar.

In Figure 32, this baseline comparison shows whether slots beat the reference antenna family. The result kept slots in the explored-but-not-selected column. Read against the section objective, it supports the decision at the end of this block: Slots did not produce a cleaner path than the passive stack for this feedhead.

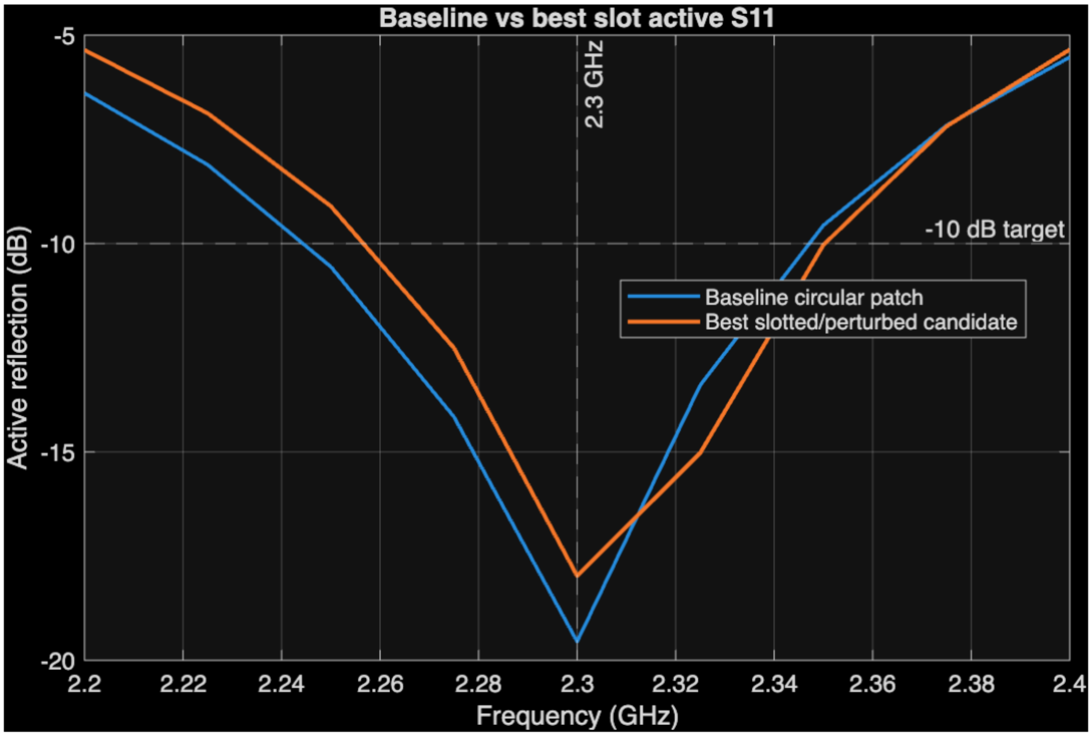

Figure 32. One-Plane CP Slot Perturbations - Graph evidence: Baseline vs best slot s11.

In Figure 33, this graph is one of the objective checks for one-plane cp slot perturbations; it supports the conclusion that slots did not produce a cleaner path than the passive stack for this feedhead. Read against the section objective, it supports the decision at the end of this block: Slots did not produce a cleaner path than the passive stack for this feedhead.

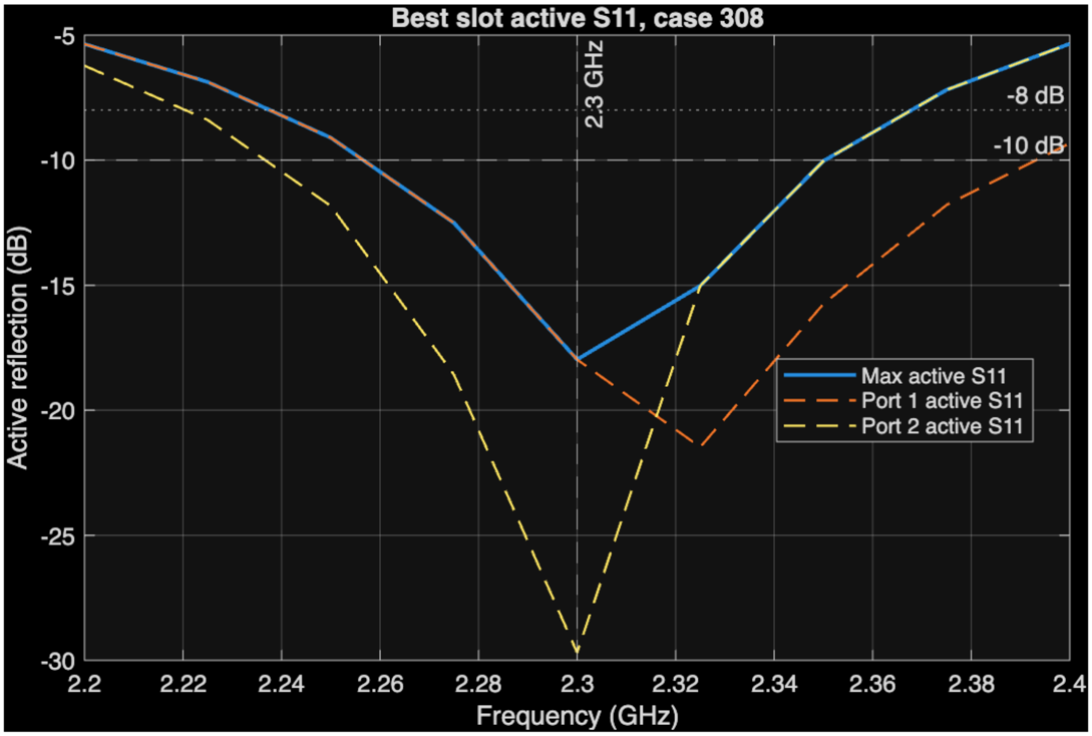

Figure 33. One-Plane CP Slot Perturbations - Graph evidence: Best slot active s11.

In Figure 34, this graph is one of the objective checks for one-plane cp slot perturbations; it supports the conclusion that slots did not produce a cleaner path than the passive stack for this feedhead. Read against the section objective, it supports the decision at the end of this block: Slots did not produce a cleaner path than the passive stack for this feedhead.

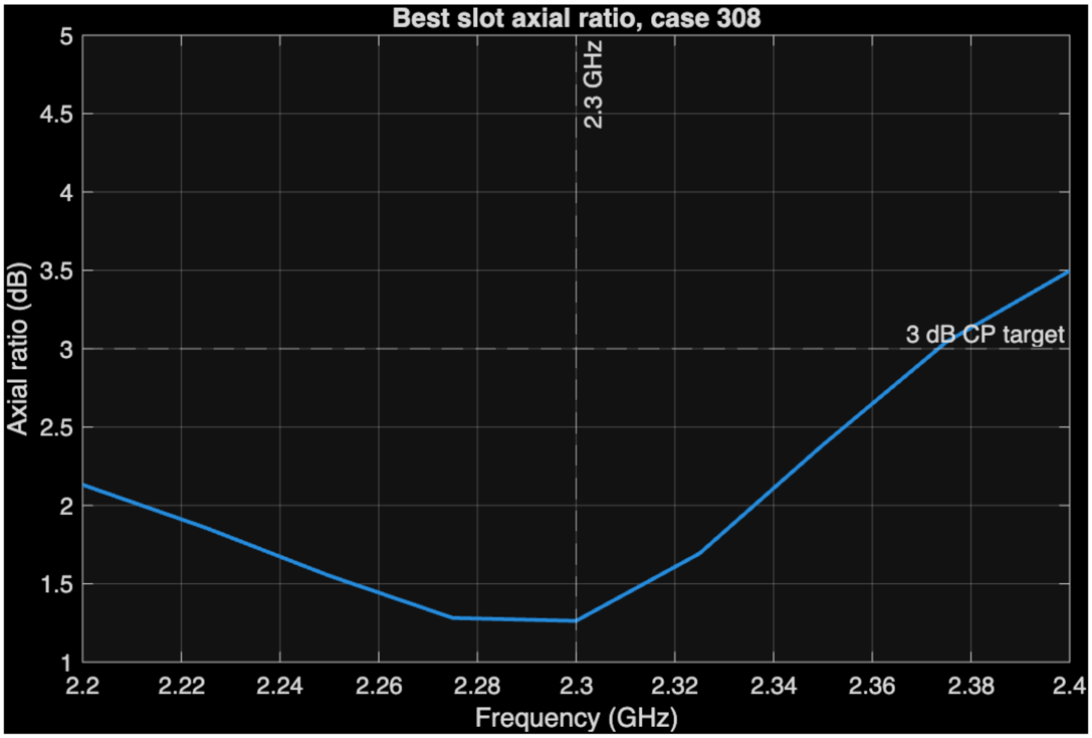

Figure 34. One-Plane CP Slot Perturbations - Graph evidence: Best slot axial ratio.

Design decision: Slots did not produce a cleaner path than the passive stack for this feedhead. Next step: Keep slots as explored/rejected and concentrate on passive stacked center-feed work.

11. Ecartometer Symmetry Refinement

The design question was: What sensor-ring geometry gives usable isolation and symmetry for tracking?

We tried it because the payload feed and tracking ring were competing for space; the sensor ring needed a stable reference independent of the center feed.

Trace source: Run 4 / ecartometer_symmetry

Scope

- Swept sensor ring geometry and compared isolated sensor behavior, center degradation, and coupling matrices.
- Tracked coupling at 2.2, 2.3, and 2.4 GHz.
- Selected the ring as a subsystem reference, not as a final combined hybrid solve.

Evidence and Interpretation

In Figure 35, this geometry reference shows the selected sensor-ring layout. The point of the surrounding plots is to test whether that ring is isolated and symmetric enough to keep. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What sensor-ring geometry gives usable isolation and symmetry for tracking?

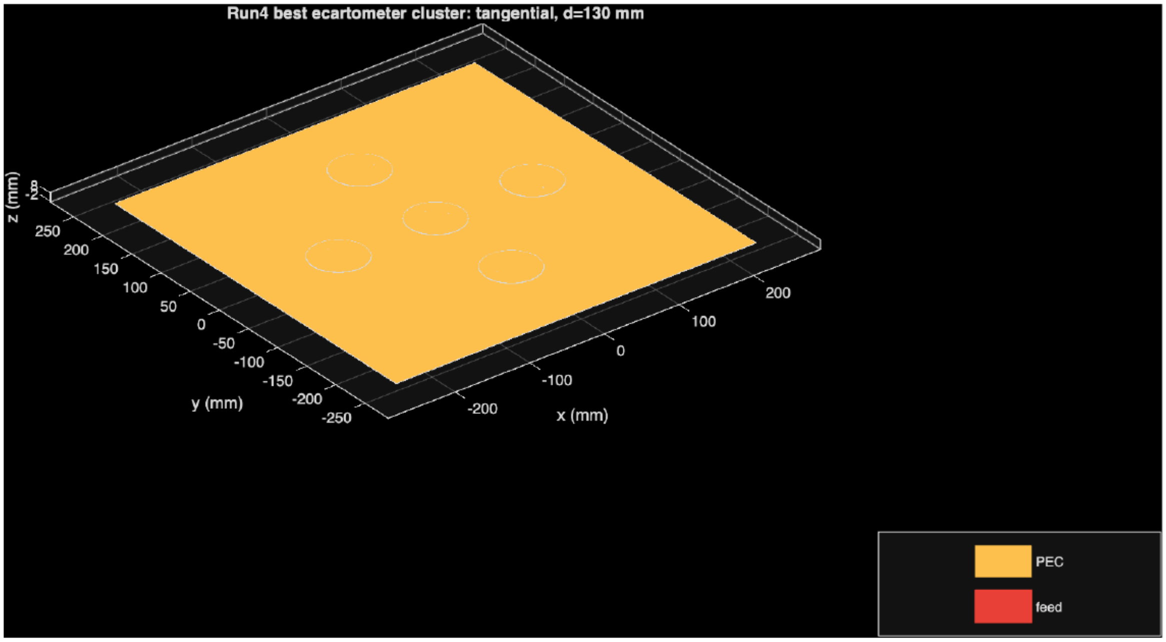

Figure 35. Ecartometer Symmetry Refinement - Reference geometry: Best cluster geometry.

In Figure 36, this plot measures the price paid by adding the tracking ring around the payload feed. Lower degradation means the ring can remain a separate subsystem. It is included here because it affected the next step in the design path: Combine the selected sensor ring later with the best center source, while preserving the caveat that they come from separate solves.

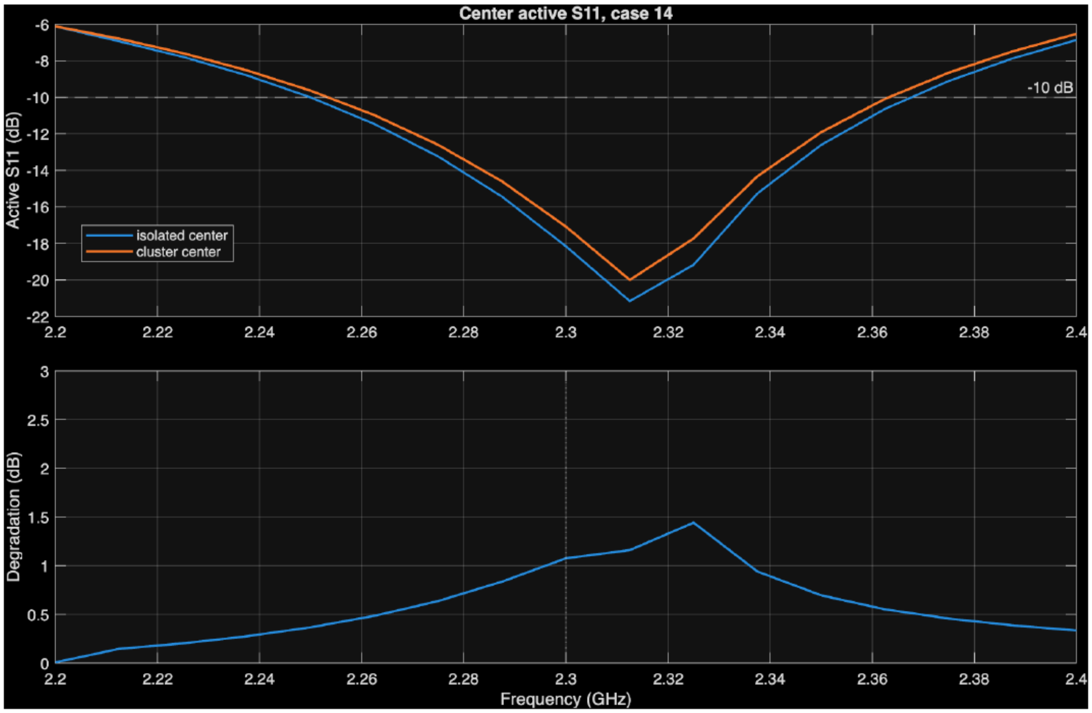

Figure 36. Ecartometer Symmetry Refinement - Generated evidence: Center degradation.

In Figure 37, the coupling matrix is the sensor-coexistence test: useful tracking sensors need isolation and symmetry, not just individual match. Read against the section objective, it supports the decision at the end of this block: Run4 case 14 became the sensor-ring source carried into Run8 and V5A.

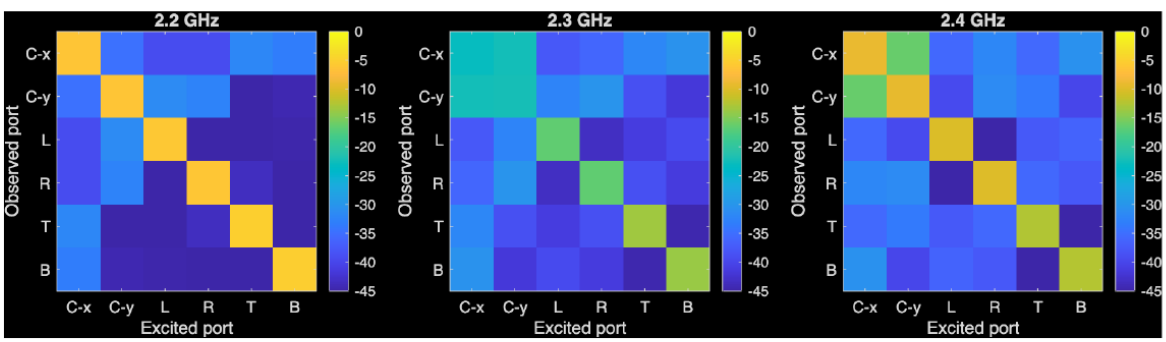

Figure 37. Ecartometer Symmetry Refinement - Graph evidence: Coupling matrix 2.2 2.3 2.4.

In Figure 38, this plot checks that the selected sensor ring did not solve coupling by detuning the sensors themselves. Read against the section objective, it supports the decision at the end of this block: Run4 case 14 became the sensor-ring source carried into Run8 and V5A.

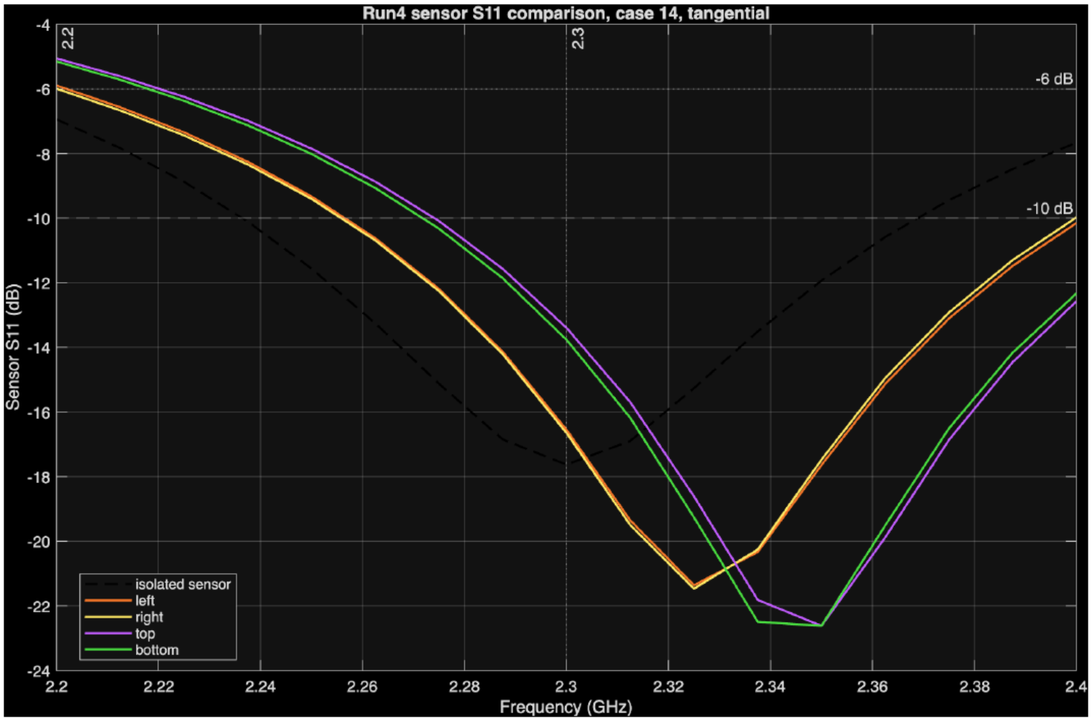

Figure 38. Ecartometer Symmetry Refinement - Graph evidence: Sensor s11 comparison.

Design decision: Run4 case 14 became the sensor-ring source carried into Run8 and V5A. Next step: Combine the selected sensor ring later with the best center source, while preserving the caveat that they come from separate solves.

12. Passive Stacked Center Feed

The design question was: Can a floating upper patch solve the center feed better than driven-only or slotted one-plane approaches?

We tried it because the first stacked test was promising; Run5A made it a focused center-feed optimization.

Trace source: Run 5A / passive_stacked_center_feed

Scope

- Swept lower radius, upper radius, ground-to-lower height, gap, and feed excitation.
- Compared the selected passive stack directly against the baseline active S11 and axial ratio.
- Checked manufacturability notes for whether the stack could be built cleanly.

Evidence and Interpretation

In Figure 39, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can a floating upper patch solve the center feed better than driven-only or slotted one-plane approaches?

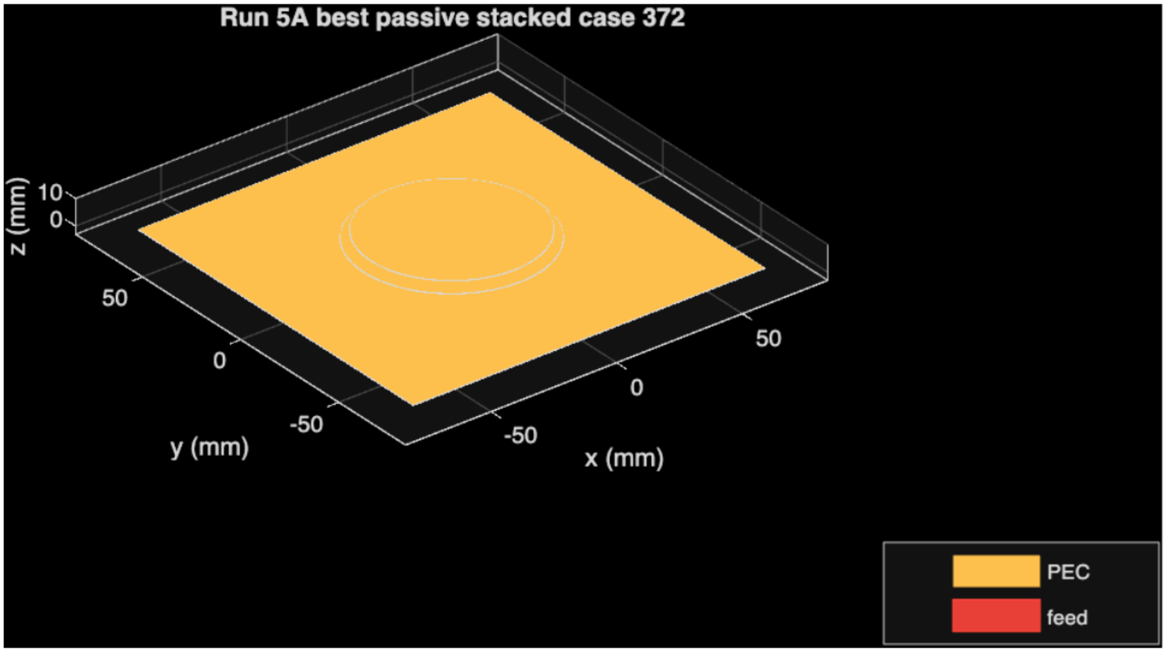

Figure 39. Passive Stacked Center Feed - Reference geometry: Best geometry.

In Figure 40, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: The passive stack improved the right metrics with less hardware risk than slots or dual-active stacking.

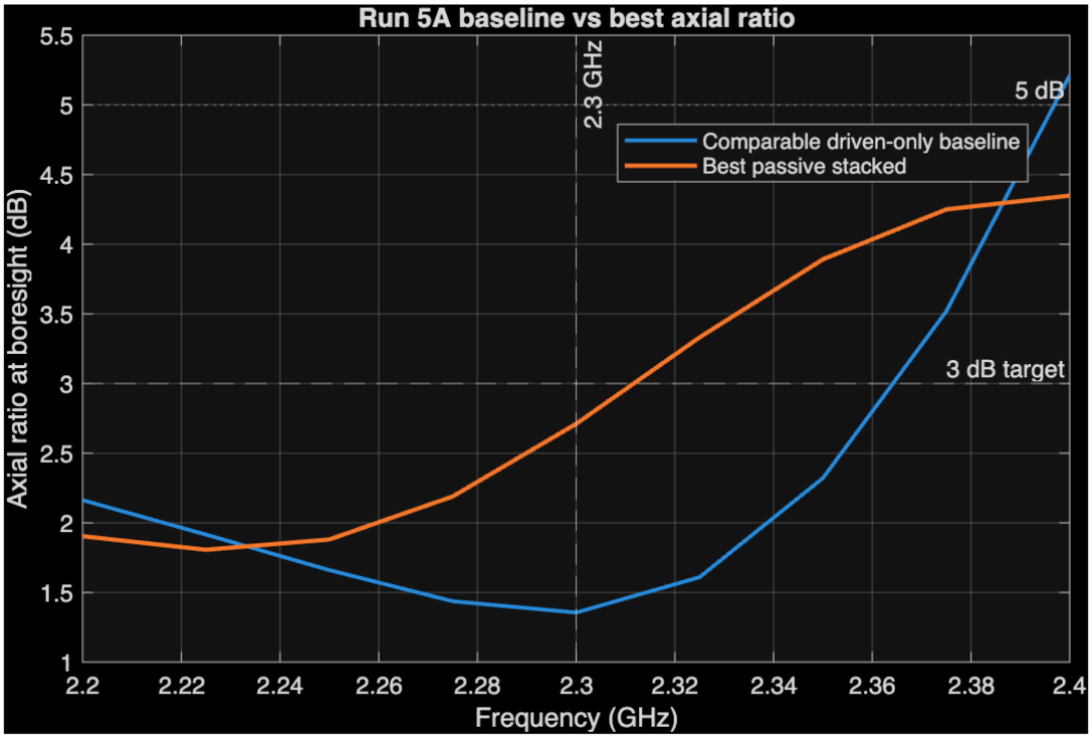

Figure 40. Passive Stacked Center Feed - Graph evidence: Baseline vs best axial ratio.

In Figure 41, this plot is a decision plot: it shows whether the new candidate improves active match over the stored baseline across the working band. Read against the section objective, it supports the decision at the end of this block: The passive stack improved the right metrics with less hardware risk than slots or dual-active stacking.

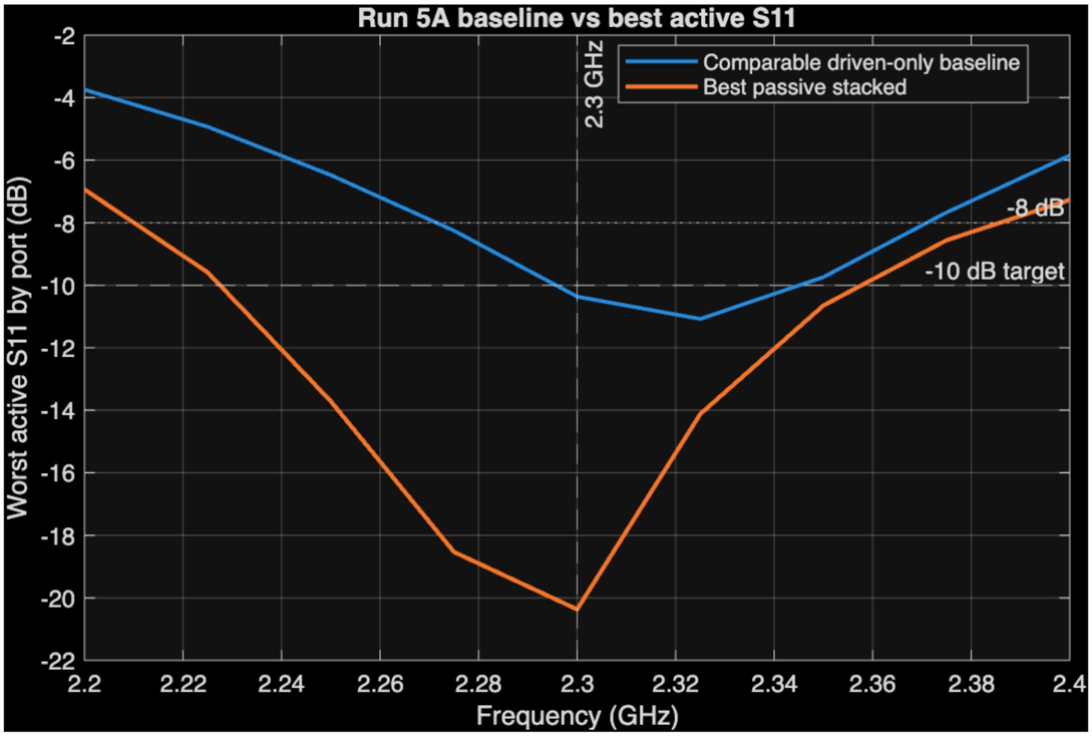

Figure 41. Passive Stacked Center Feed - Graph evidence: Baseline vs best s11.

Design decision: The passive stack improved the right metrics with less hardware risk than slots or dual-active stacking. Next step: Use passive stacking as the preferred center-feed family and test whether active upper feeding adds enough value.

13. Dual-Active Stacked Center Feed

The design question was: Does actively feeding both stack layers outperform a passive upper patch enough to justify complexity?

We tried it because a second driven layer might broaden or shape CP response, but it adds feed routing and assembly complexity.

Trace source: Run 5B / dual_active_stacked_center_feed

Scope

- Ranked dual-active stack candidates and tracked contribution by frequency.
- Compared combined S11 and axial-ratio behavior against the simpler passive stack.

- Evaluated whether added control offsets added build risk without enough performance gain.

Evidence and Interpretation

In Figure 42, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Does actively feeding both stack layers outperform a passive upper patch enough to justify complexity?

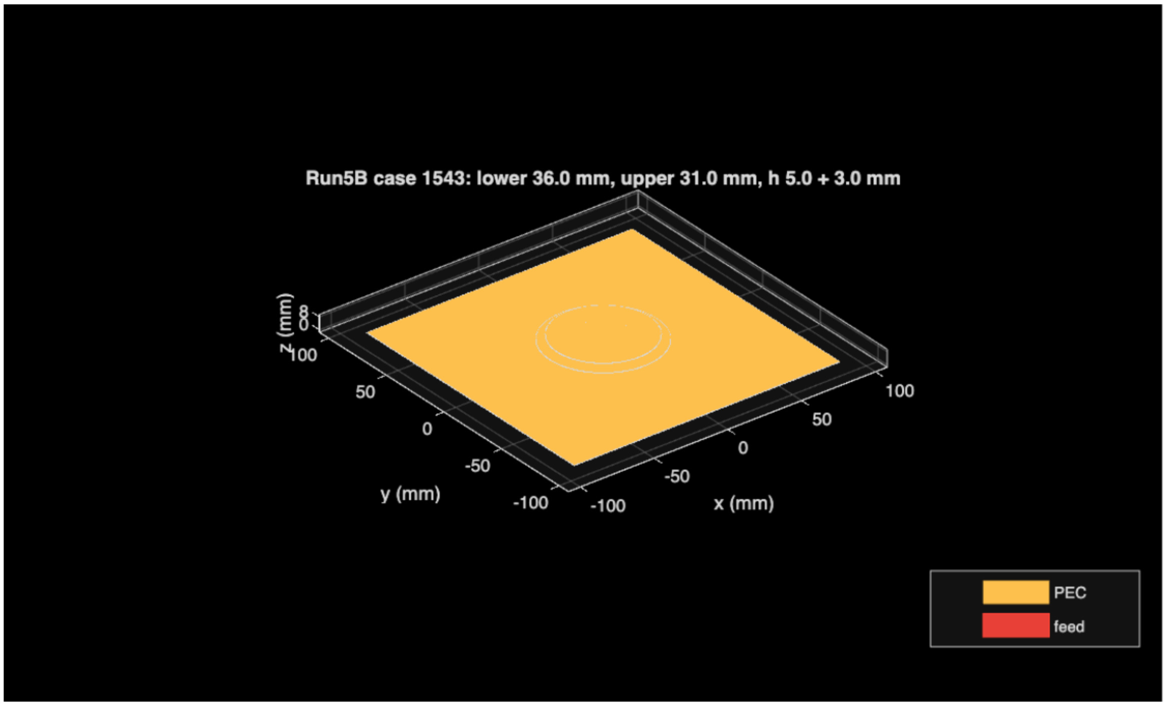

Figure 42. Dual-Active Stacked Center Feed - Reference geometry: Best geometry.

In Figure 43, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: Dual-active stacking was more complicated than the passive result and did not become the final path.

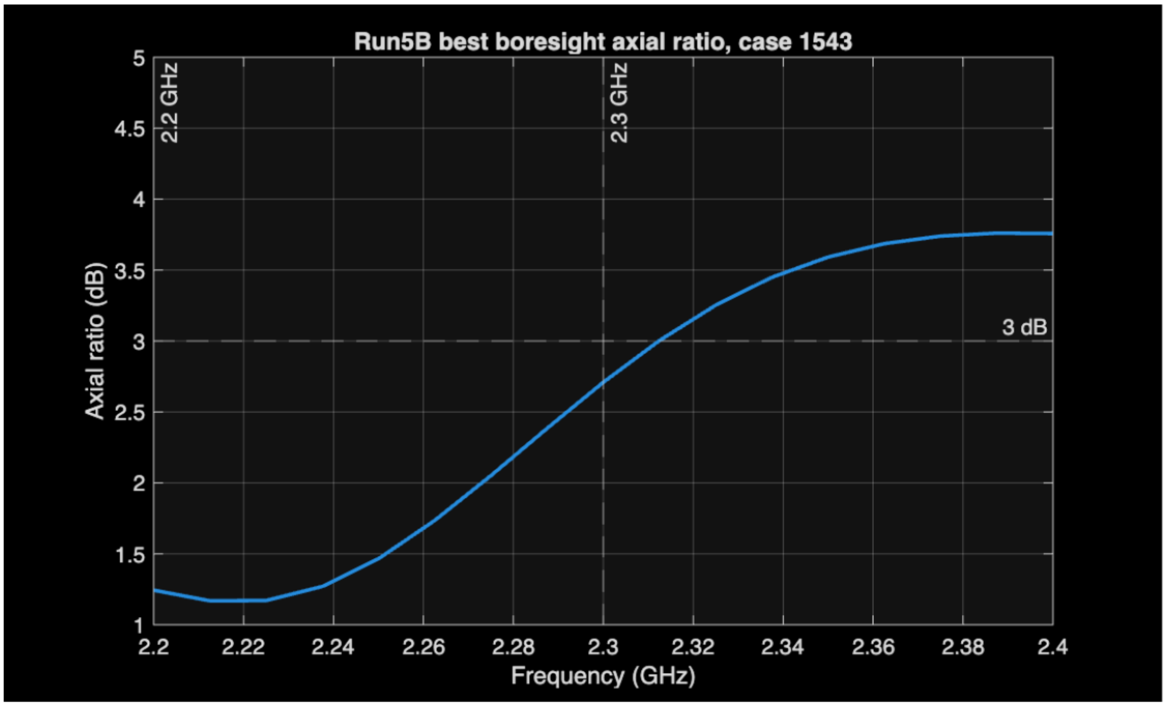

Figure 43. Dual-Active Stacked Center Feed - Graph evidence: Best axial ratio.

In Figure 44, for the dual-active stack, this plot shows whether feeding both layers gives enough active-match benefit to justify the extra feed hardware. Read against the section objective, it supports the decision at the end of this block: Dual-active stacking was more complicated than the passive result and did not become the final path.

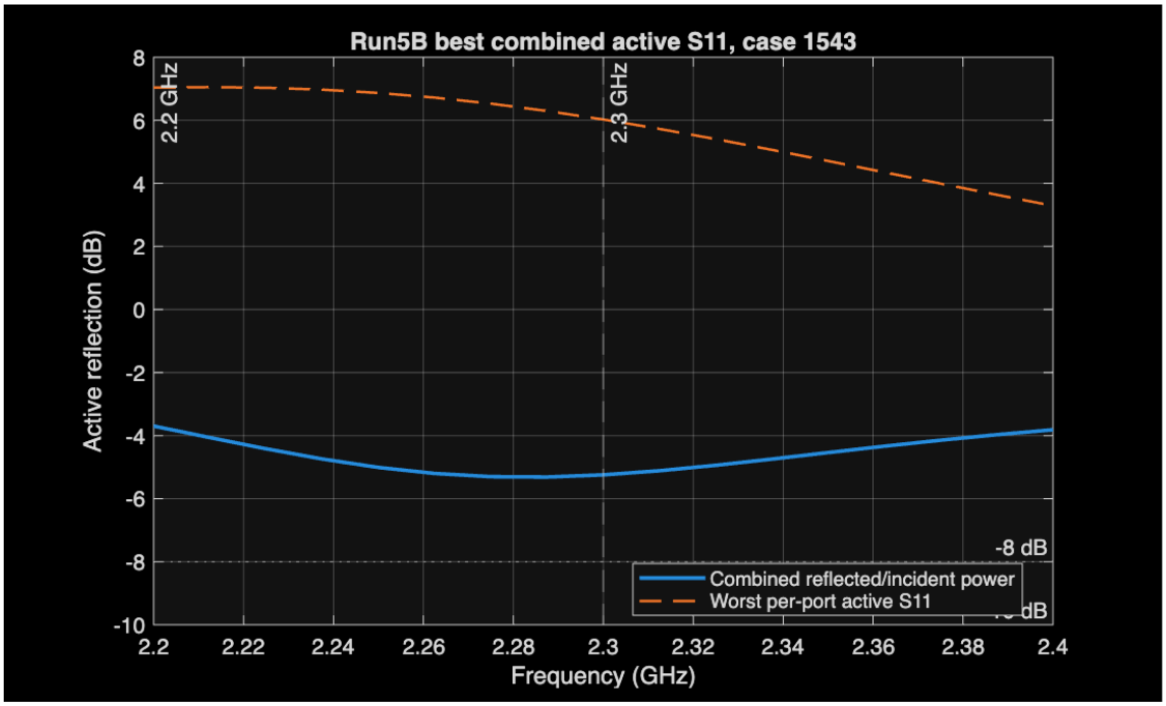

Figure 44. Dual-Active Stacked Center Feed - Graph evidence: Best combined s11.

In Figure 45, this contribution plot explains how the two active stack layers share the response. It is useful diagnostically, but it also exposes extra complexity. It is included here because it affected the next step in the design path: Keep the passive upper patch as the center-feed architecture.

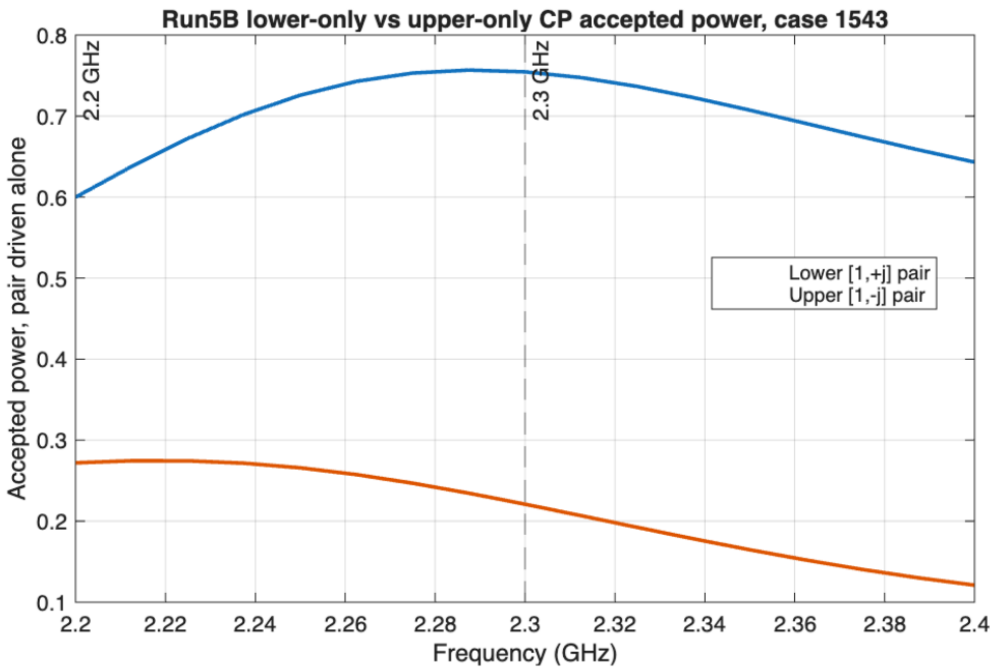

Figure 45. Dual-Active Stacked Center Feed - Generated evidence: Contribution vs frequency.

Design decision: Dual-active stacking was more complicated than the passive result and did not become the final path. Next step: Keep the passive upper patch as the center-feed architecture.

14. Slot-Based Center Feed Suite

The design question was: Can driven or parasitic slots beat the passive circular stack if searched more broadly?

We tried it because before freezing the passive stack, we made a broader slot suite so the rejection was based on evidence, not preference.

Trace source: Run 6 / center_slotted_patch_suite

Scope

- Ran staged searches across driven slots, parasitic slots, and refined candidates.
- Saved separate best driven-slot, parasitic-slot, and overall geometry references.
- Compared baseline versus best S11 and axial ratio.

Evidence and Interpretation

In Figure 46, this geometry isolates the driven-layer slot idea, which was the most direct competitor to the passive driven/lower plus floating/upper stack. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can driven or parasitic slots beat the passive circular stack if searched more broadly?

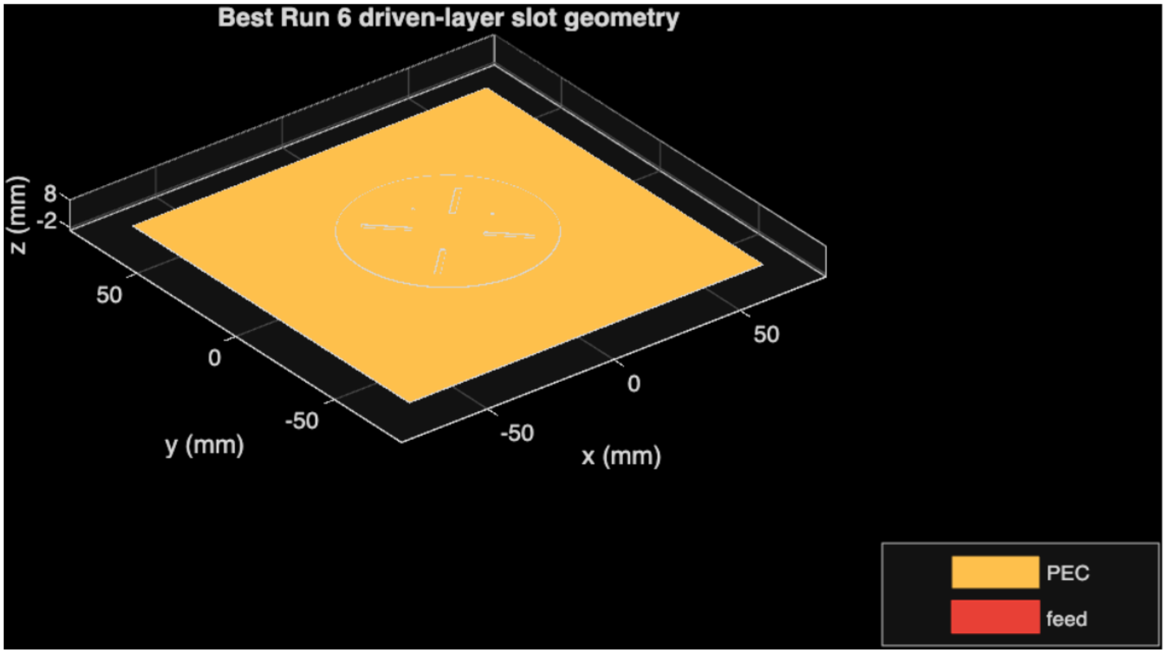

Figure 46. Slot-Based Center Feed Suite - Reference geometry: Best driven slot geometry.

In Figure 47, this is the best slot-suite geometry after screening. It documents that the slot family was judged from its strongest candidate, not from a weak example. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can driven or parasitic slots beat the passive circular stack if searched more broadly?

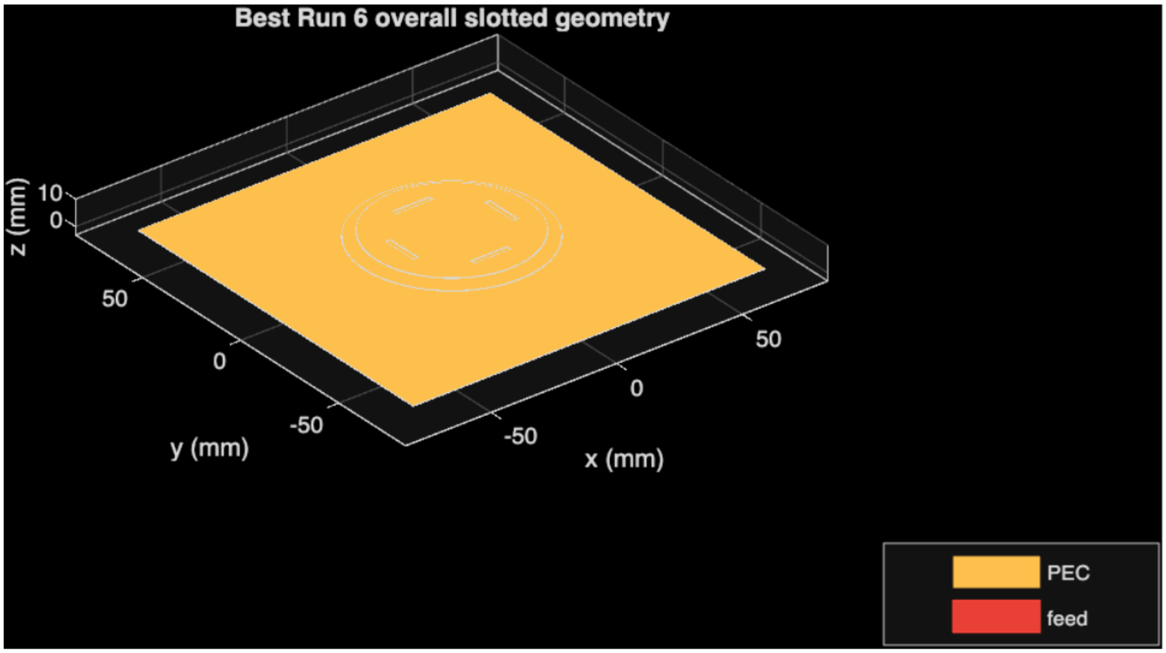

Figure 47. Slot-Based Center Feed Suite - Reference geometry: Best overall geometry.

In Figure 48, this geometry isolates the parasitic-slot idea: a nearby alternative to the passive upper patch that still did not become cleaner than the circular stack. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can driven or parasitic slots beat the passive circular stack if searched more broadly?

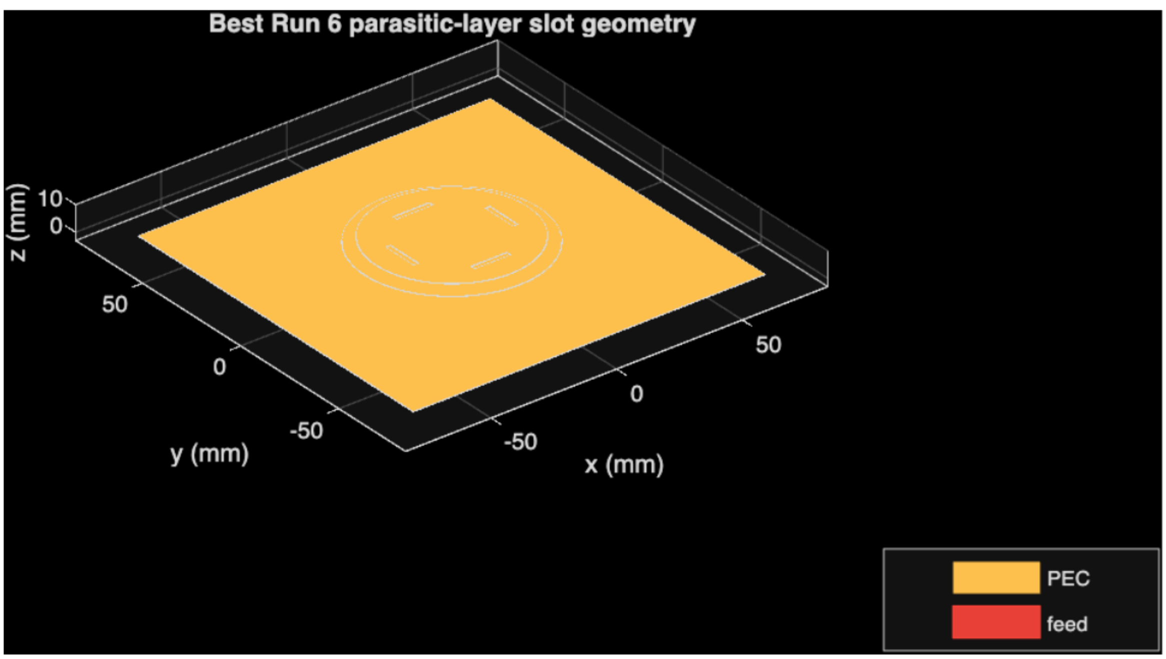

Figure 48. Slot-Based Center Feed Suite - Reference geometry: Best parasitic slot geometry.

In Figure 49, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: The slot families produced interesting cases but did not provide a cleaner validated path than passive stacking.

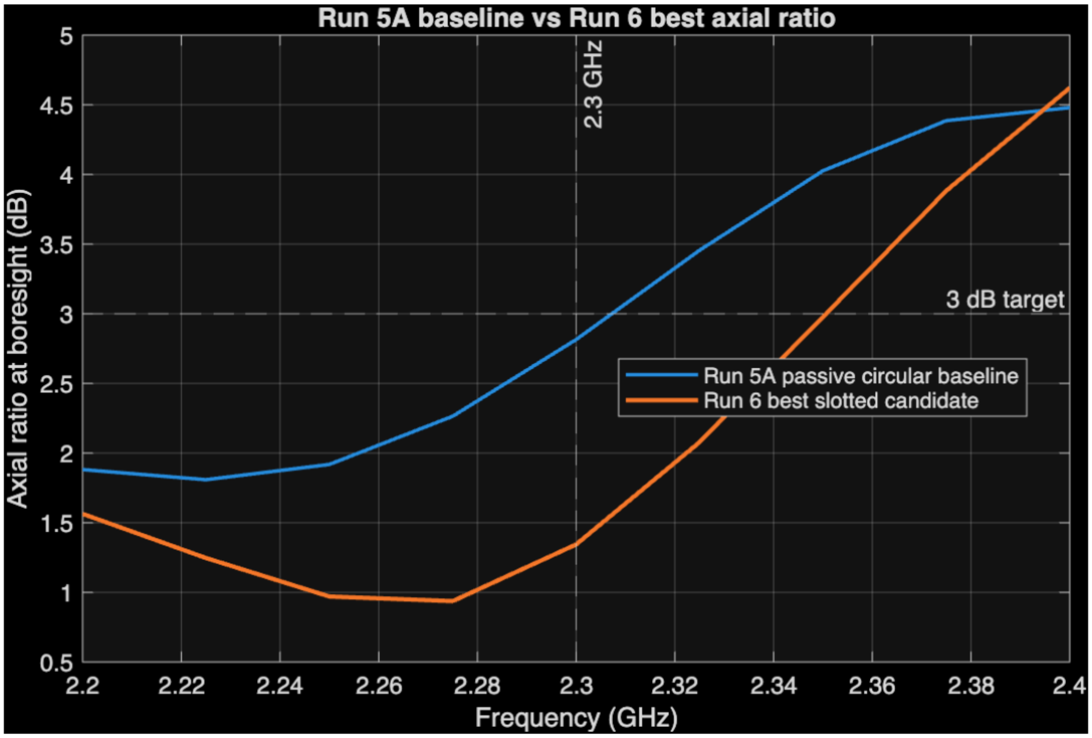

Figure 49. Slot-Based Center Feed Suite - Graph evidence: Baseline vs best axial ratio.

In Figure 50, this plot is a decision plot: it shows whether the new candidate improves active match over the stored baseline across the working band. Read against the section objective, it supports the decision at the end of this block: The slot families produced interesting cases but did not provide a cleaner validated path than passive stacking.

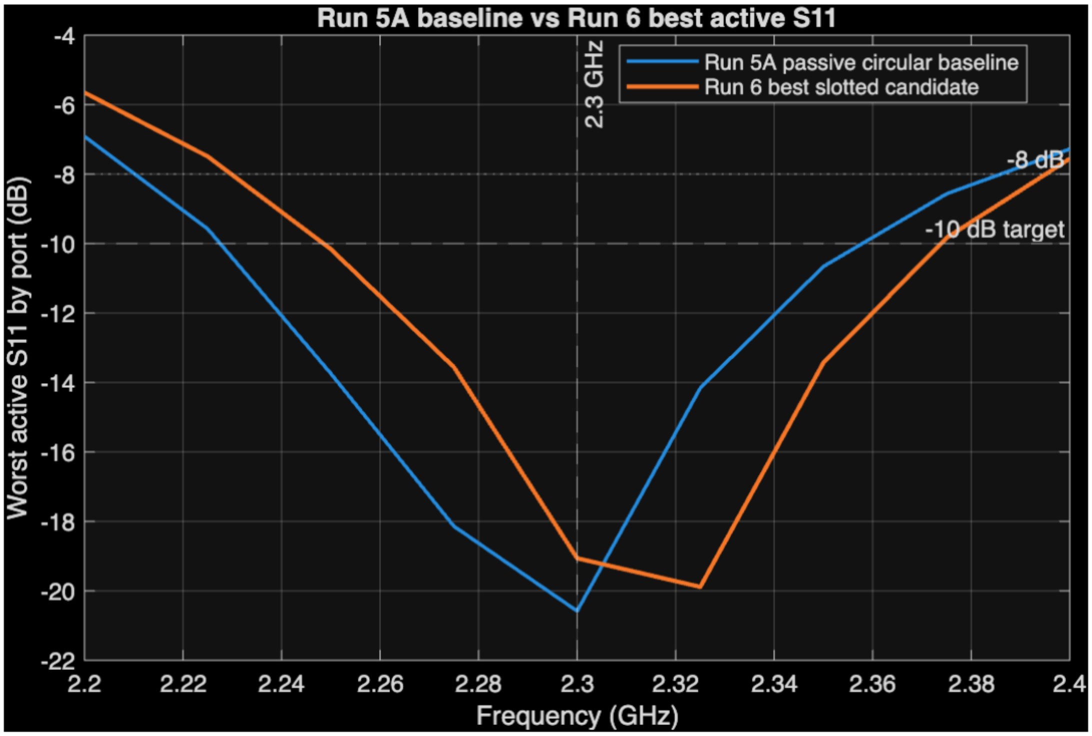

Figure 50. Slot-Based Center Feed Suite - Graph evidence: Baseline vs best s11.

In Figure 51, this slot-suite axial-ratio result shows whether the best slot candidate solved the CP problem as well as the matching problem. Read against the section objective, it supports the decision at the end of this block: The slot families produced interesting cases but did not provide a cleaner validated path than passive stacking.

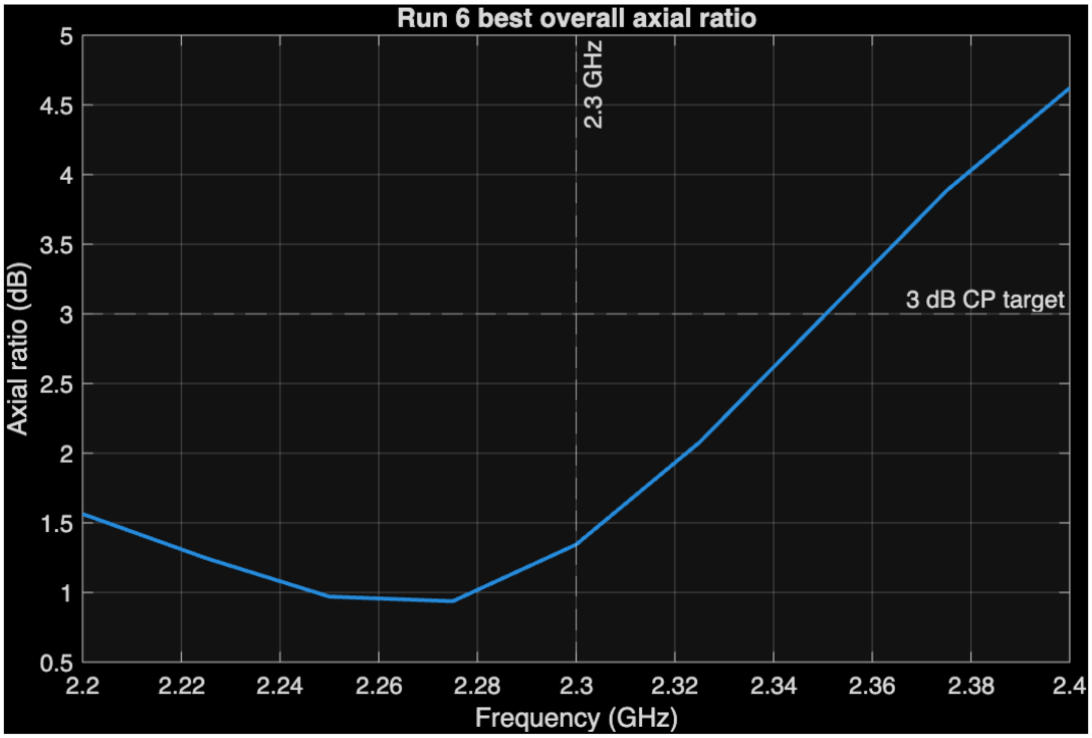

Figure 51. Slot-Based Center Feed Suite - Graph evidence: Best overall axial ratio.

In Figure 52, this slot-suite S11 result shows the best available match behavior from the broader slot search. Read against the section objective, it supports the decision at the end of this block: The slot families produced interesting cases but did not provide a cleaner validated path than passive stacking.

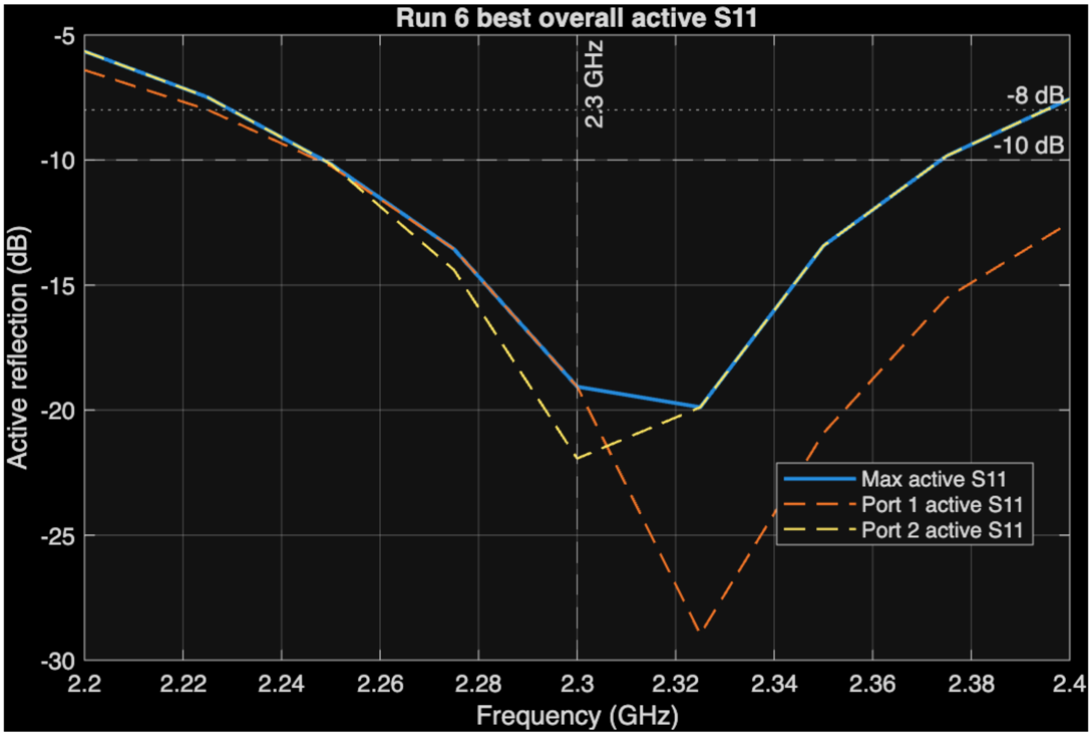

Figure 52. Slot-Based Center Feed Suite - Graph evidence: Best overall s11.

Design decision: The slot families produced interesting cases but did not provide a cleaner validated path than passive stacking. Next step: Perform one more focused driven-layer slot check, then return to passive-stack refinement if it fails.

15. Driven-Layer Slot Final Check

The design question was: Is there a manufacturable driven-layer slot geometry that passes all hard rejects?

We tried it because this was the final challenge to the passive-stack choice: a focused driven-layer slot set with explicit manufacturability screening.

Trace source: Run 6B / driven_layer_slotted_cp_center

Scope

- Generated 48 candidate geometries across rectangular, rounded, circular-ring, curved, and truncated-corner families.
- Ranked active S11, axial ratio, coverage, gain spread, manufacturability, and hard-reject flags.

- Stored every candidate geometry in a gallery for traceability.

Evidence and Interpretation

In Figure 53, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Is there a manufacturable driven-layer slot geometry that passes all hard rejects?

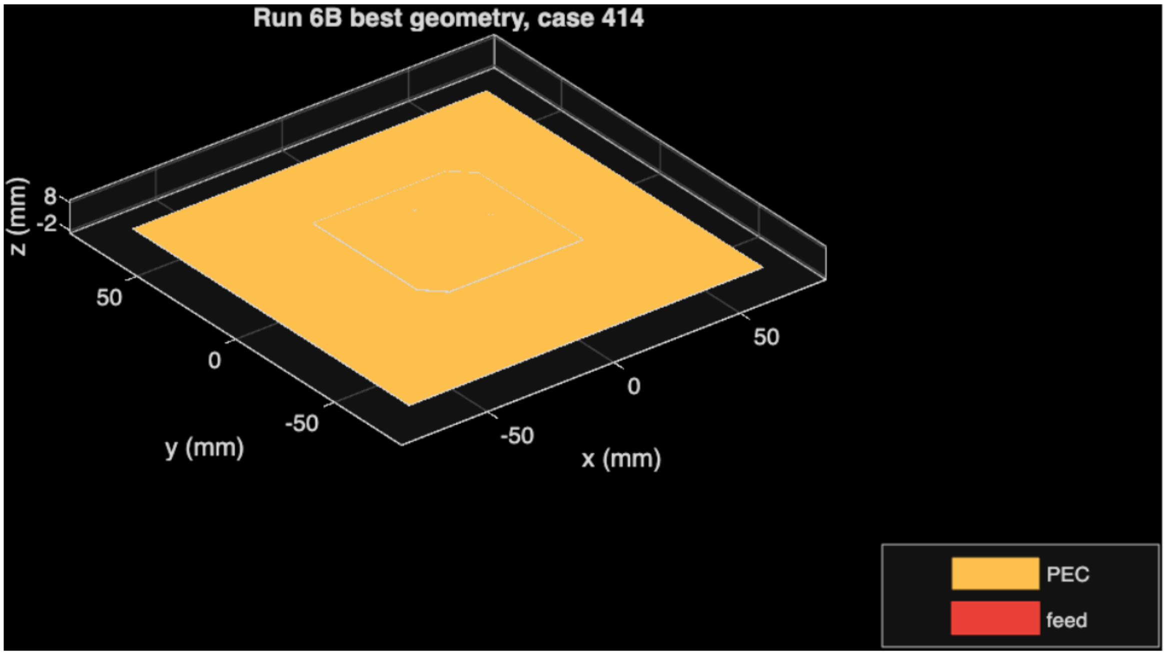

Figure 53. Driven-Layer Slot Final Check - Reference geometry: Best geometry.

In Figure 54, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: No driven-layer slot candidate displaced the passive stack under the hard-reject rules.

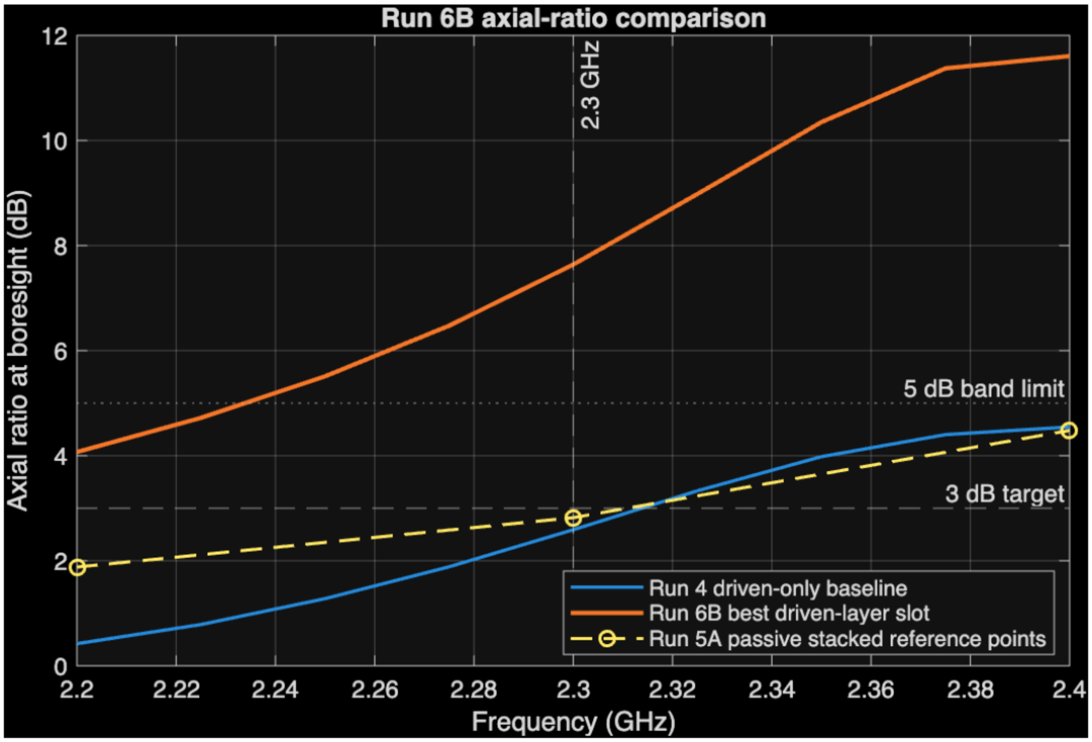

Figure 54. Driven-Layer Slot Final Check - Graph evidence: Baseline vs best axial ratio.

In Figure 55, this plot is a decision plot: it shows whether the new candidate improves active match over the stored baseline across the working band. Read against the section objective, it supports the decision at the end of this block: No driven-layer slot candidate displaced the passive stack under the hard-reject rules.

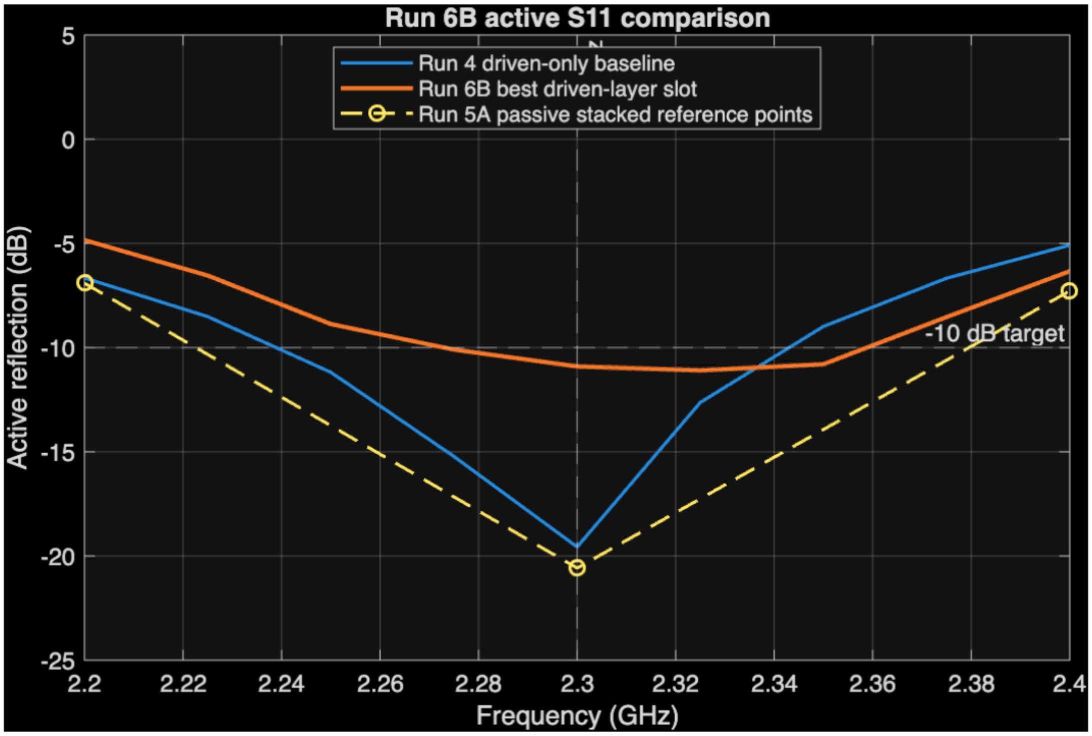

Figure 55. Driven-Layer Slot Final Check - Graph evidence: Baseline vs best s11.

In Figure 56, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: No driven-layer slot candidate displaced the passive stack under the hard-reject rules.

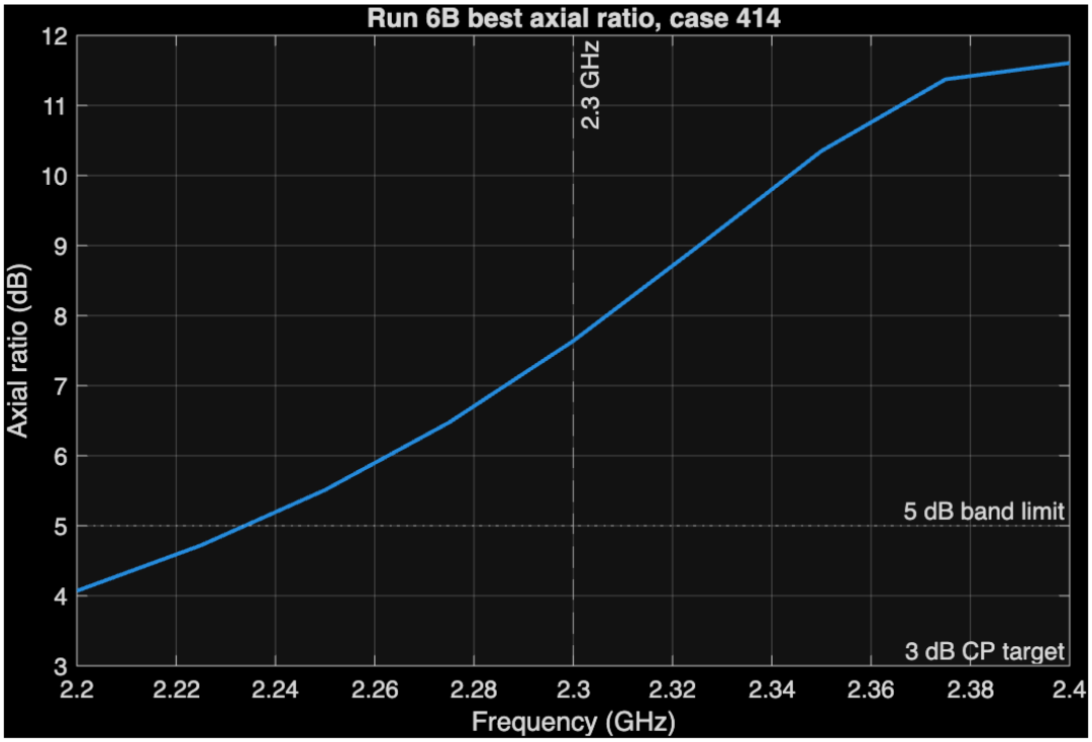

Figure 56. Driven-Layer Slot Final Check - Graph evidence: Best axial ratio.

In Figure 57, this is the final driven-slot check on active match. It was included to make sure the passive-stack choice survived a focused slot challenge. Read against the section objective, it supports the decision at the end of this block: No driven-layer slot candidate displaced the passive stack under the hard-reject rules.

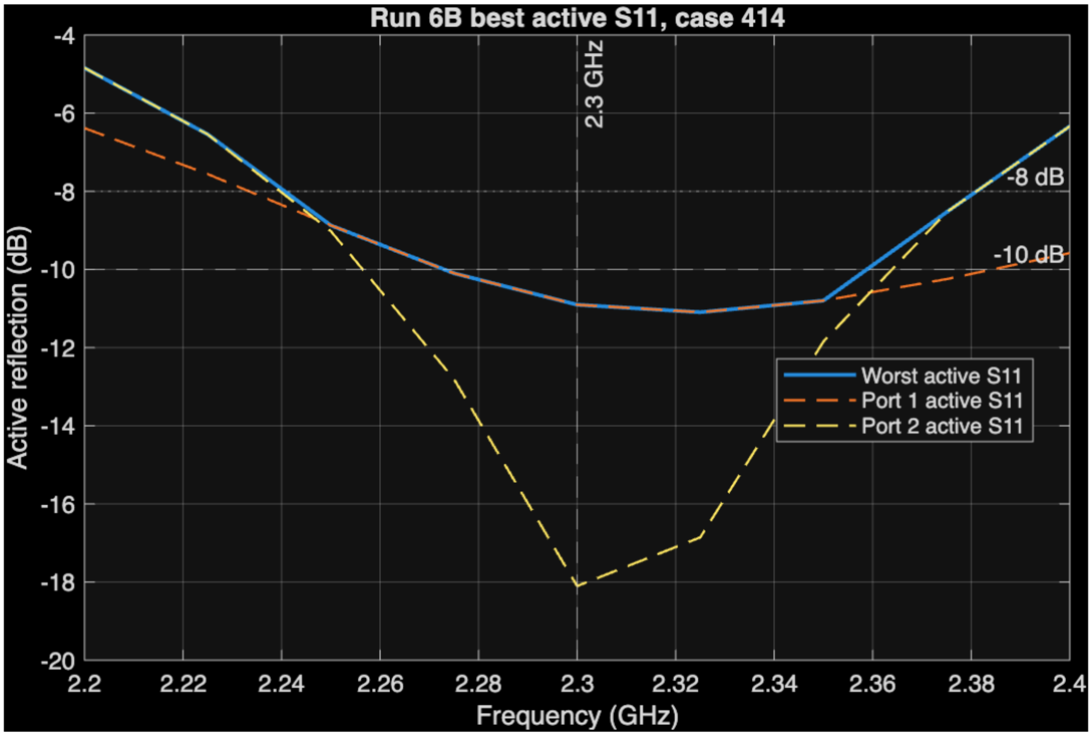

Figure 57. Driven-Layer Slot Final Check - Graph evidence: Best s11.

In Figure 58, this atlas sheet keeps the rejected slot search visible without burying the reader in one candidate per page. The point is breadth: the passive-stack decision was made after these driven-slot geometries were generated and screened. That matters because the rejected geometries are part of the engineering record; they show the design space we deliberately ruled out before returning to the passive stacked center feed.

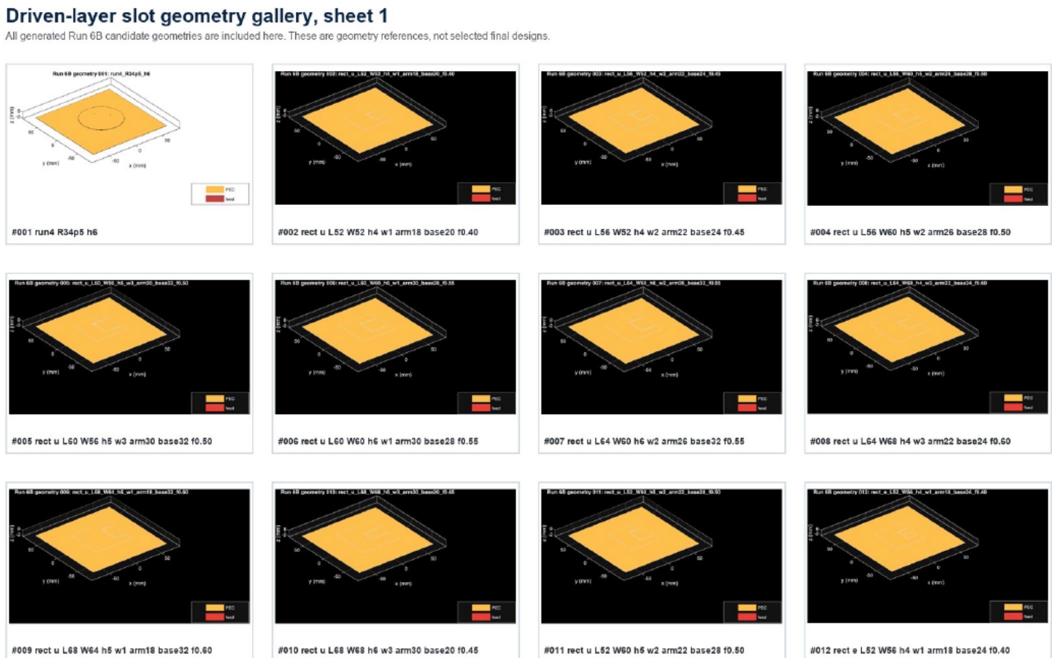

Figure 58. Driven-Layer Slot Final Check - Generated geometry gallery: Geometry gallery sheet 1.

In Figure 59, this atlas sheet keeps the rejected slot search visible without burying the reader in one candidate per page. The point is breadth: the passive-stack decision was made after these driven-slot geometries were generated and screened. That matters because the rejected geometries are part of the engineering record; they show the design space we deliberately ruled out before returning to the passive stacked center feed.

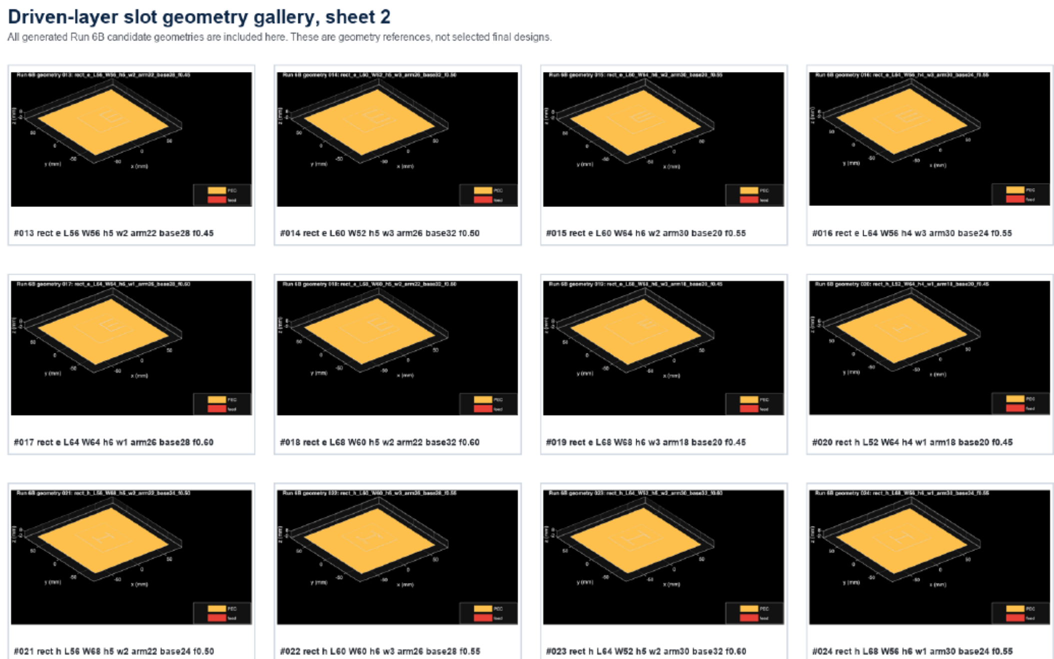

Figure 59. Driven-Layer Slot Final Check - Generated geometry gallery: Geometry gallery sheet 2.

In Figure 60, this atlas sheet keeps the rejected slot search visible without burying the reader in one candidate per page. The point is breadth: the passive-stack decision was made after these driven-slot geometries were generated and screened. That matters because the rejected geometries are part of the engineering record; they show the design space we deliberately ruled out before returning to the passive stacked center feed.

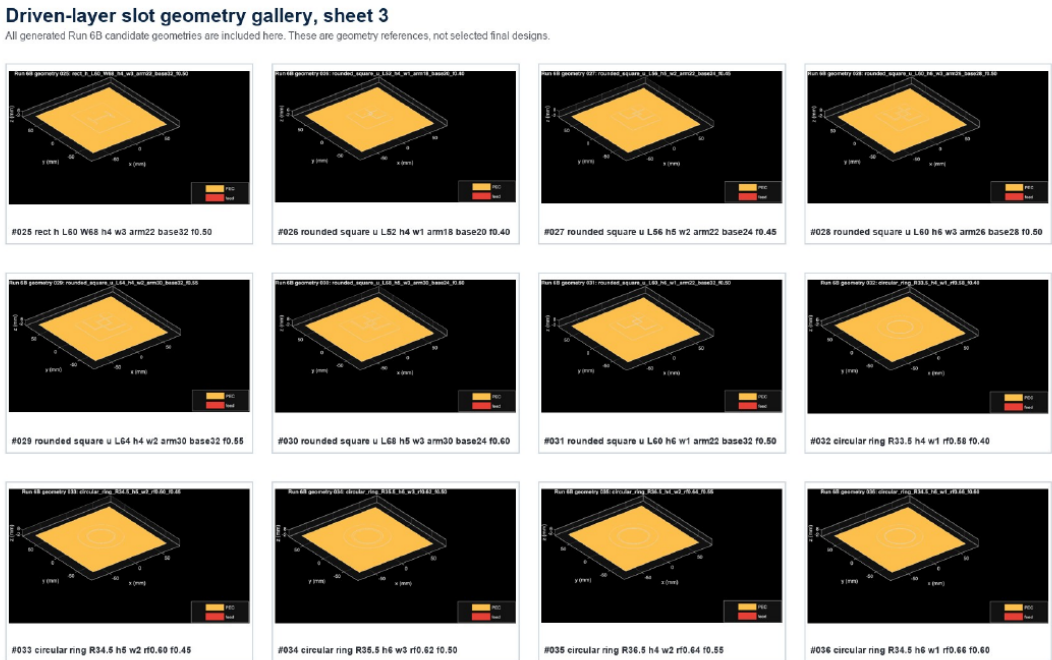

Figure 60. Driven-Layer Slot Final Check - Generated geometry gallery: Geometry gallery sheet 3.

In Figure 61, this atlas sheet keeps the rejected slot search visible without burying the reader in one candidate per page. The point is breadth: the passive-stack decision was made after these driven-slot geometries were generated and screened. That matters because the rejected geometries are part of the engineering record; they show the design space we deliberately ruled out before returning to the passive stacked center feed.

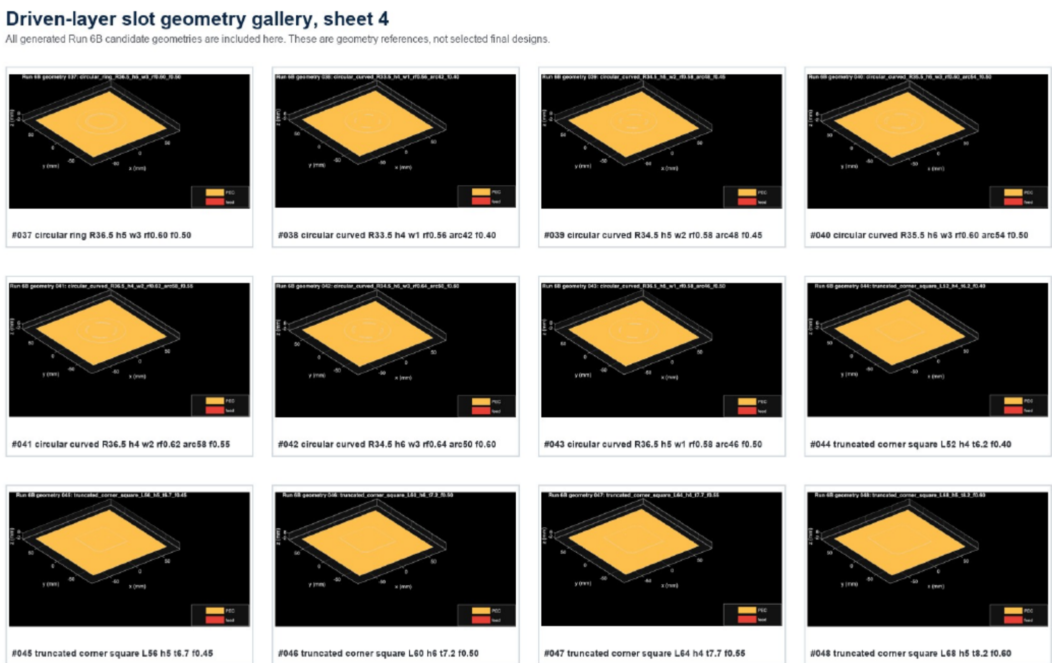

Figure 61. Driven-Layer Slot Final Check - Generated geometry gallery: Geometry gallery sheet 4.

Design decision: No driven-layer slot candidate displaced the passive stack under the hard-reject rules. Next step: Finalize the passive stacked center family and refine it directly in Run7.

16. Center Feed Reference Consolidation

The design question was: What center-feed geometry should be carried into the architecture package?

We tried it because after several competing directions, we needed a direct comparison between driven-only and passive-stack references under consistent scoring.

Trace source: Run 7 / refine_center_feed_references

Scope

- Rebuilt Run4 driven-only and Run5A passive-stack references.
- Ran staged driven-only and passive-stack refinements, then direct-checked top candidates.
- Selected a manufacturable passive-stack case with balanced S11, axial ratio, and gain.

Evidence and Interpretation

In Figure 62, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What center-feed geometry should be carried into the architecture package?

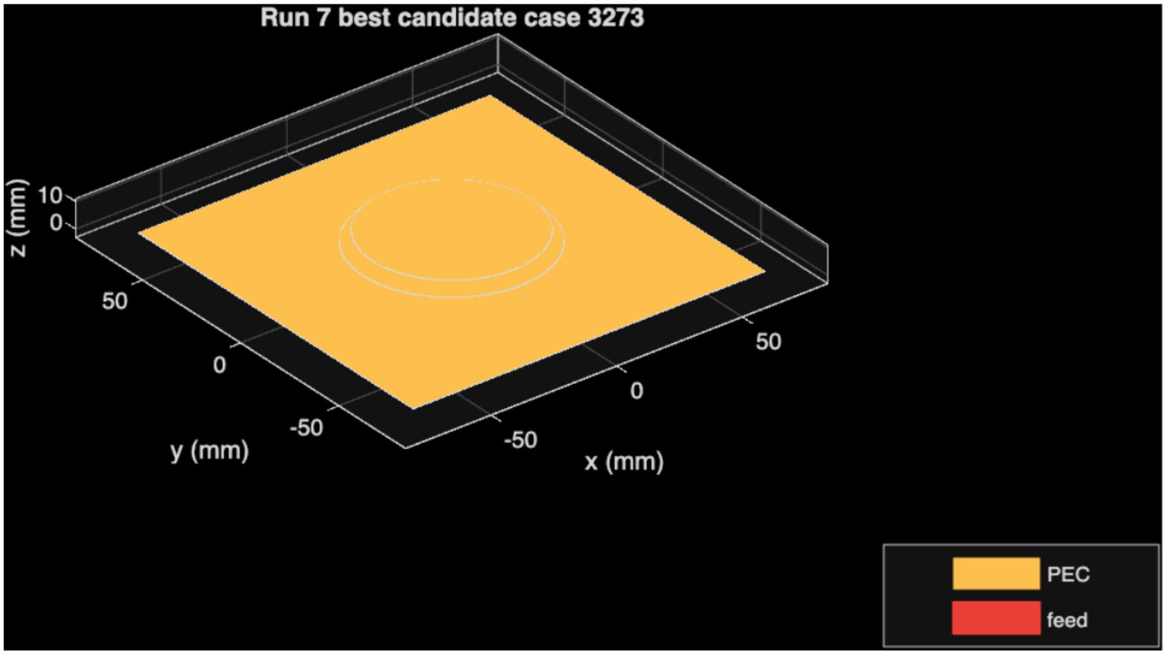

Figure 62. Center Feed Reference Consolidation - Reference geometry: Best geometry.

In Figure 63, this geometry brings the earlier driven-only reference into the same comparison frame as the passive-stack candidate. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What center-feed geometry should be carried into the architecture package?

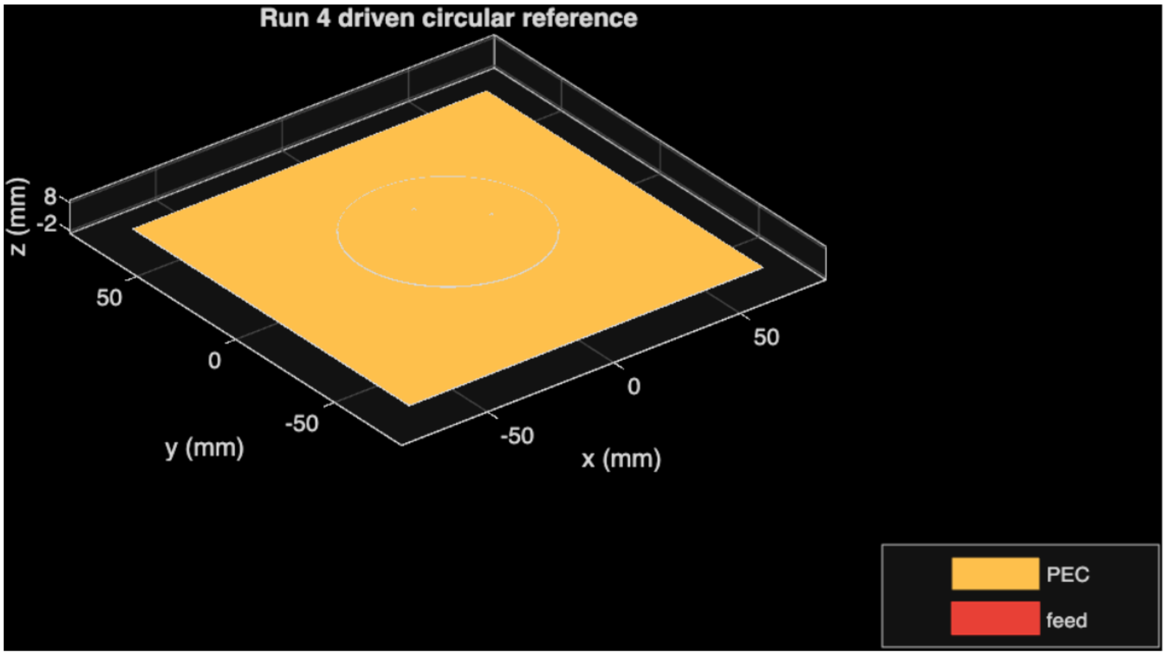

Figure 63. Center Feed Reference Consolidation - Reference geometry: Reference geometry.

In Figure 64, this geometry brings the passive-stack reference into the same comparison frame as the driven-only candidate. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What center-feed geometry should be carried into the architecture package?

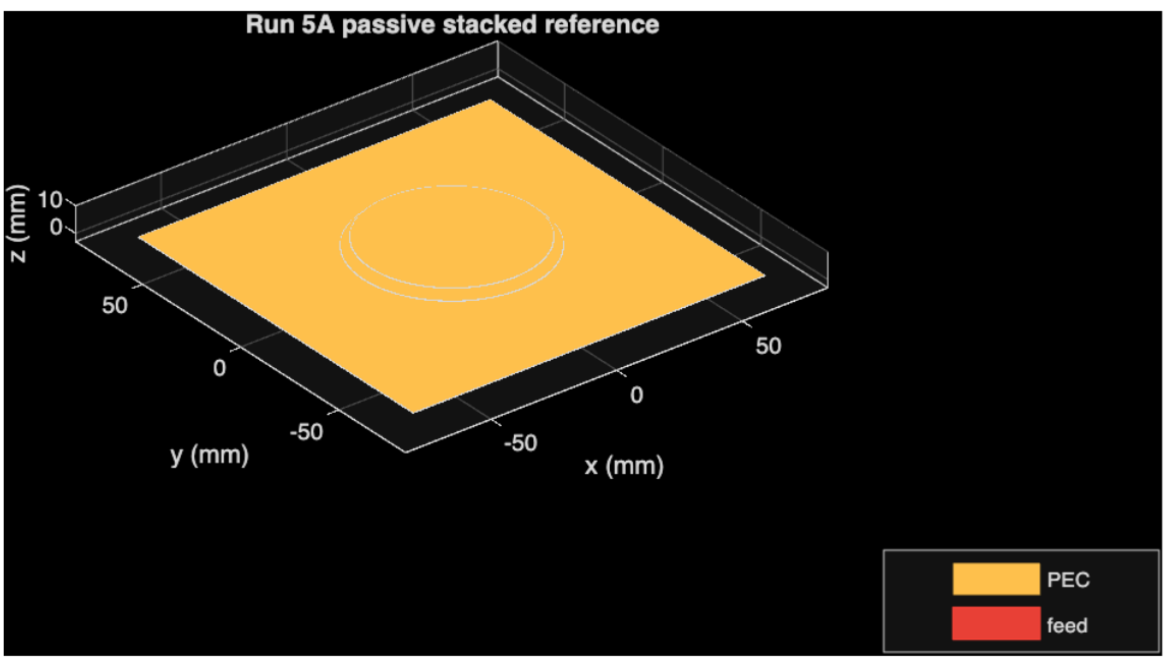

Figure 64. Center Feed Reference Consolidation - Reference geometry: Reference geometry.

In Figure 65, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: Case 3273, passive stacked, lower R35.25 mm, upper R31.50 mm, h4/gap4, became the selected center source.

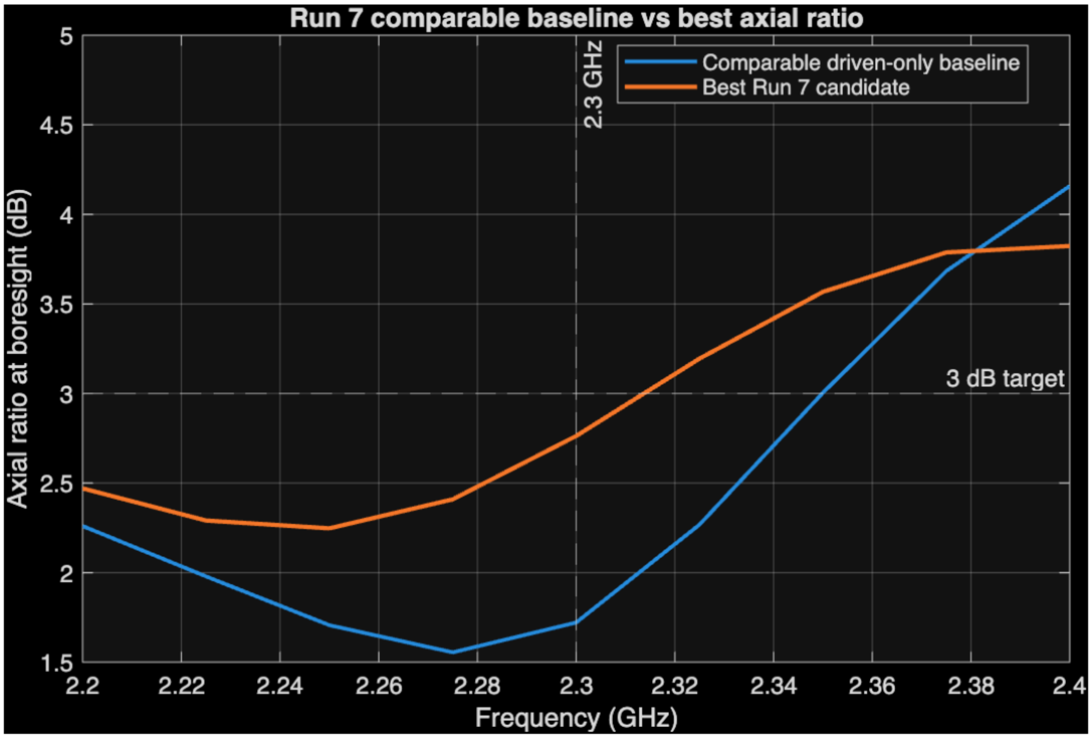

Figure 65. Center Feed Reference Consolidation - Graph evidence: Baseline vs best axial ratio.

In Figure 66, this plot is a decision plot: it shows whether the new candidate improves active match over the stored baseline across the working band. Read against the section objective, it supports the decision at the end of this block: Case 3273, passive stacked, lower R35.25 mm, upper R31.50 mm, h4/gap4, became the selected center source.

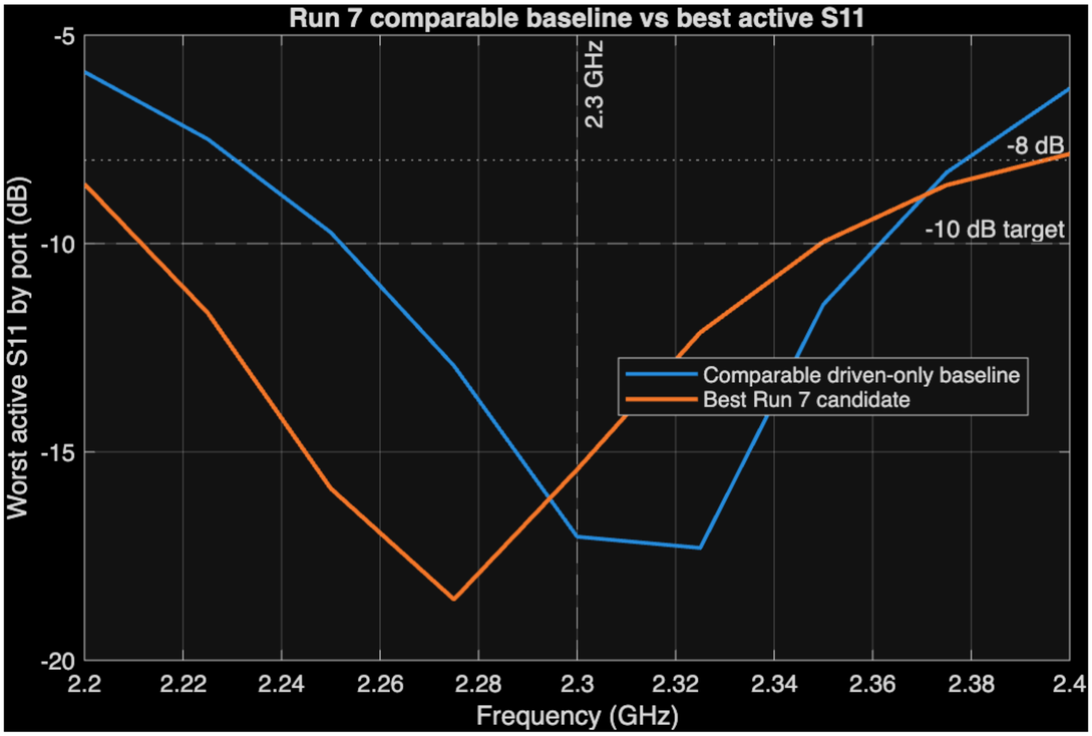

Figure 66. Center Feed Reference Consolidation - Graph evidence: Baseline vs best s11.

In Figure 67, the baseline axial-ratio plot shows what CP behavior had to be improved or at least preserved during refinement. Read against the section objective, it supports the decision at the end of this block: Case 3273, passive stacked, lower R35.25 mm, upper R31.50 mm, h4/gap4, became the selected center source.

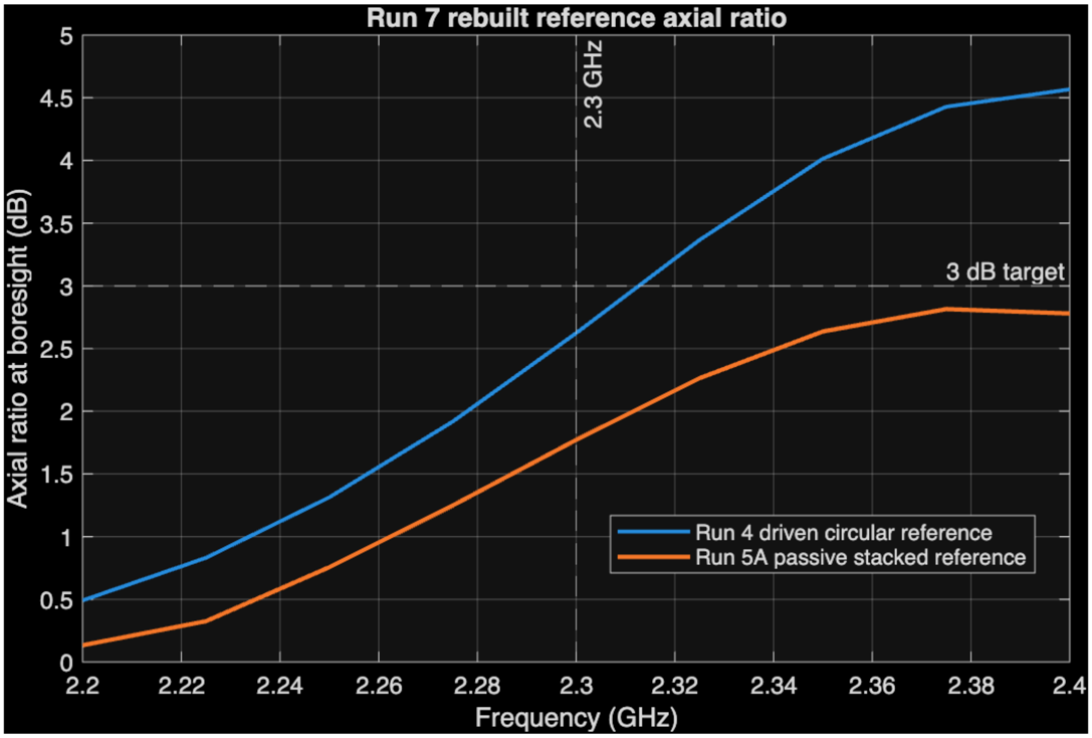

Figure 67. Center Feed Reference Consolidation - Graph evidence: Reference baseline axial ratio.

In Figure 68, the baseline plot keeps the comparison honest by showing the reference response before refinement. Read against the section objective, it supports the decision at the end of this block: Case 3273, passive stacked, lower R35.25 mm, upper R31.50 mm, h4/gap4, became the selected center source.

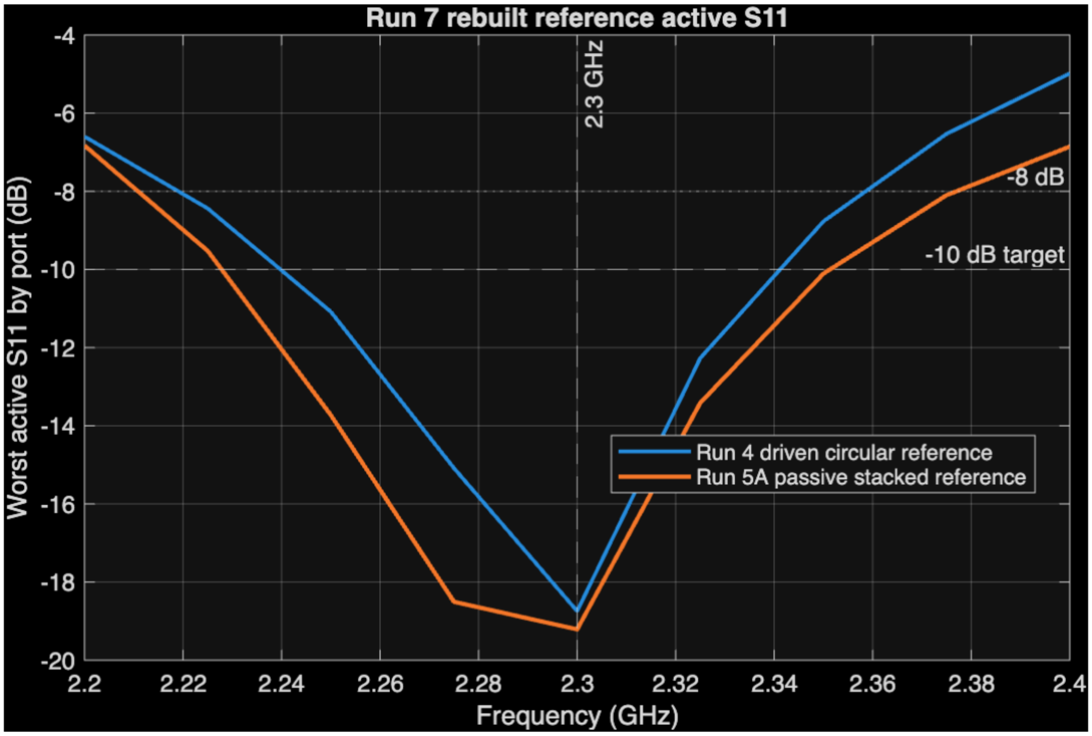

Figure 68. Center Feed Reference Consolidation - Graph evidence: Reference baseline s11.

> **Design decision:** Case 3273, passive stacked, lower R35.25 mm, upper R31.50 mm, h4/gap4, became the selected center source. Next step: Compile the selected center with the selected ecartometer ring in Run8.

17. Architecture Compilation and Polarization Phase Check

The design question was: What combined feedhead architecture should move toward packaging?

We tried it because at this point the best center and best sensor-ring data lived in separate runs; Run8 assembled the traceable architecture.

Trace source: Run 8 / final_architecture

Scope

- Compiled Run7 center-feed curves, Run4 ecartometer S11/coupling, and a six-port coupling view.
- Added a phase-angle check on the selected geometry to understand CP sensitivity.
- Preserved source-file maps and caveats instead of pretending a new full hybrid solve was performed.

Evidence and Interpretation

In Figure 69, this geometry is the reference shape for the architecture compilation and polarization phase check objective; the following plots should be read against this physical layout. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

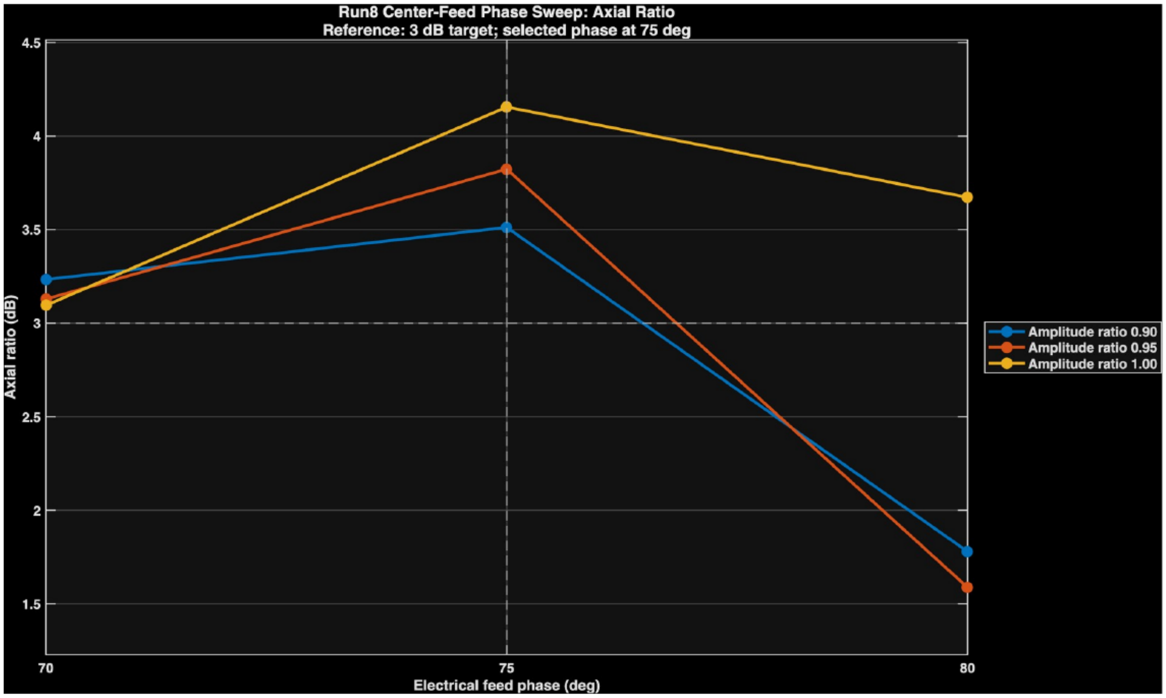

Figure 69. Architecture Compilation and Polarization Phase Check - Reference geometry: Axial ratio vs phase selected geometry.

In Figure 70, this plot shows how well the undesired circular-polarization component is suppressed as phase changes. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

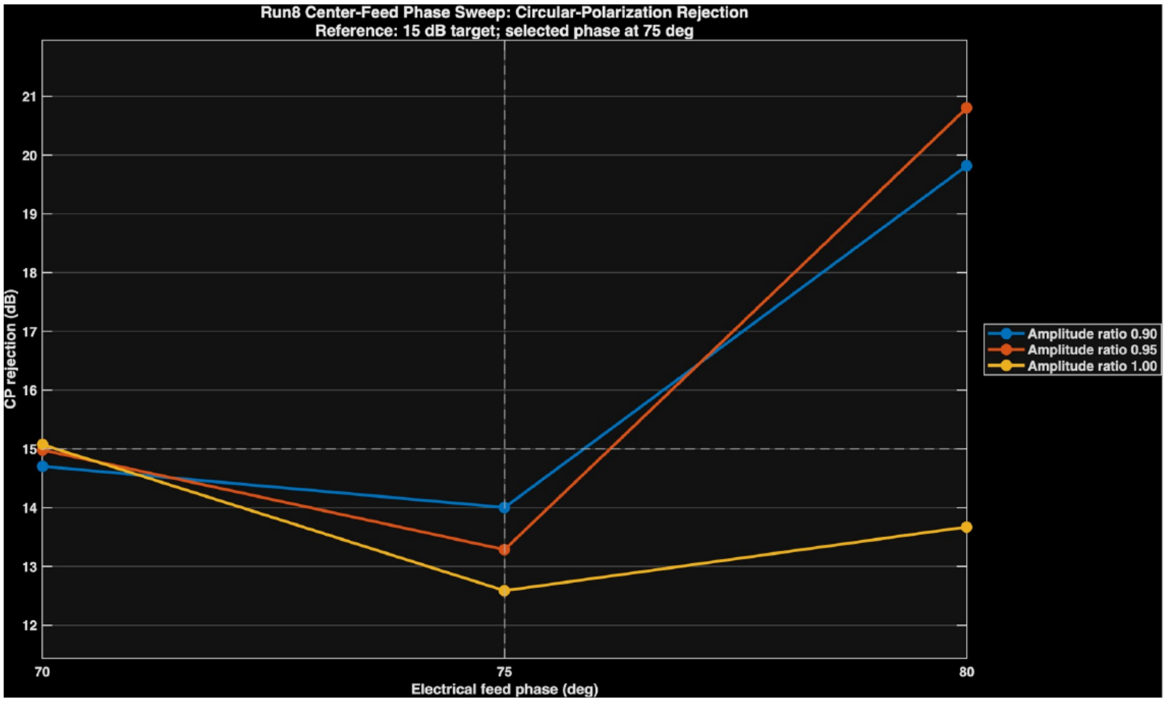

Figure 70. Architecture Compilation and Polarization Phase Check - Reference geometry: Cp rejection vs phase selected geometry.

In Figure 71, the phase sweep tests how sensitive the selected CP behavior is to excitation phase. It informs feed-network and switching tolerances. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

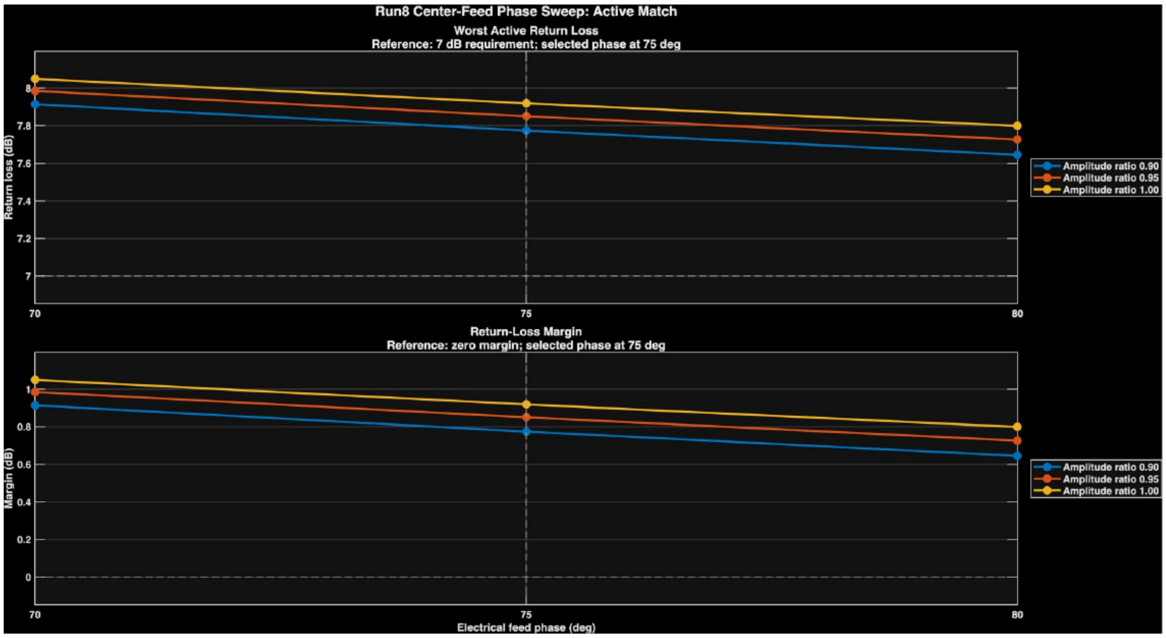

Figure 71. Architecture Compilation and Polarization Phase Check - Reference geometry: Phase sweep active s11 selected geometry.

In Figure 72, the phase sweep tests how sensitive the selected CP behavior is to excitation phase. It informs feed-network and switching tolerances. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

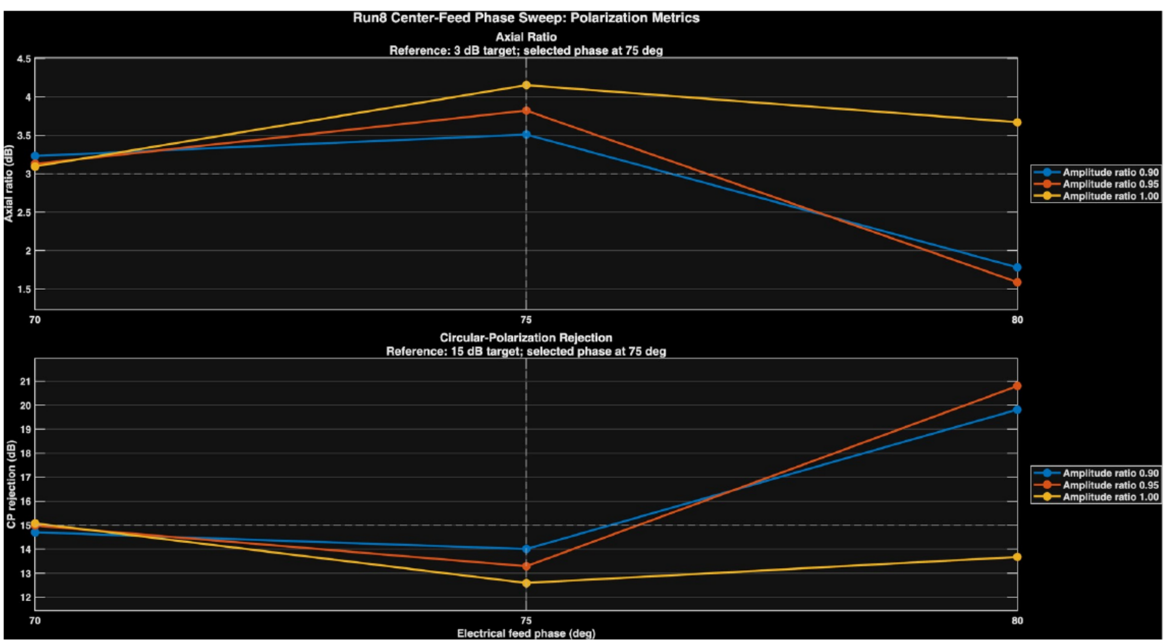

Figure 72. Architecture Compilation and Polarization Phase Check - Reference geometry: Polarization phase sweep selected geometry.

In Figure 73, this is the first architecture-level drawing: selected center feed plus selected sensor ring. It is the bridge from isolated antenna studies to a feedhead. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

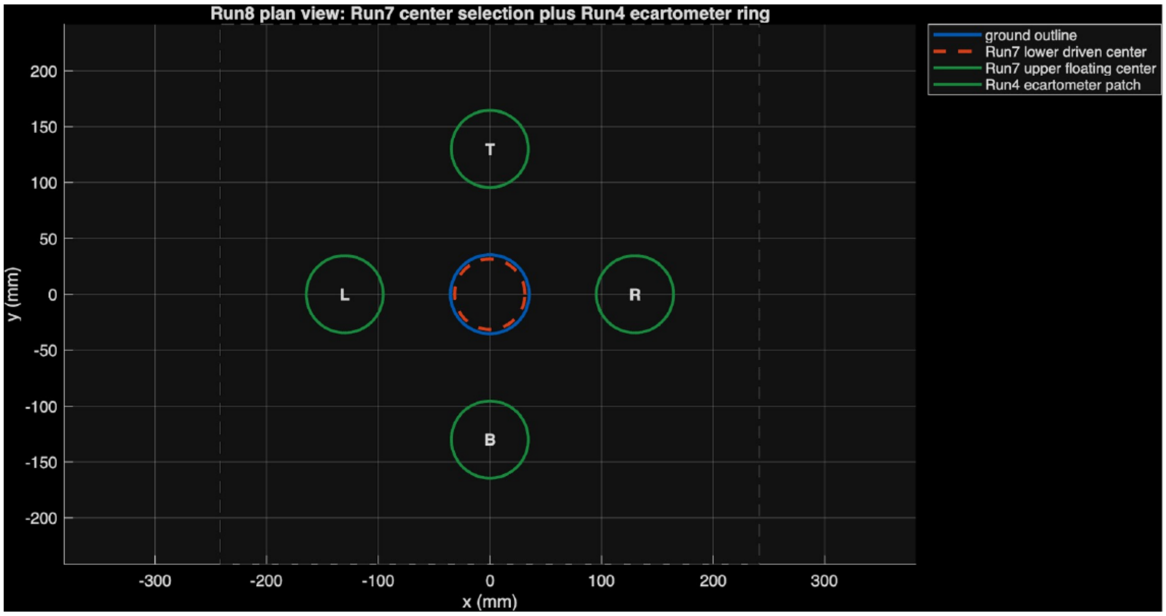

Figure 73. Architecture Compilation and Polarization Phase Check - Reference geometry: Architecture plan view.

In Figure 74, this geometry locks the center-stack source used by the architecture compilation, so the following curves can be traced back to a physical feed. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What combined feedhead architecture should move toward packaging?

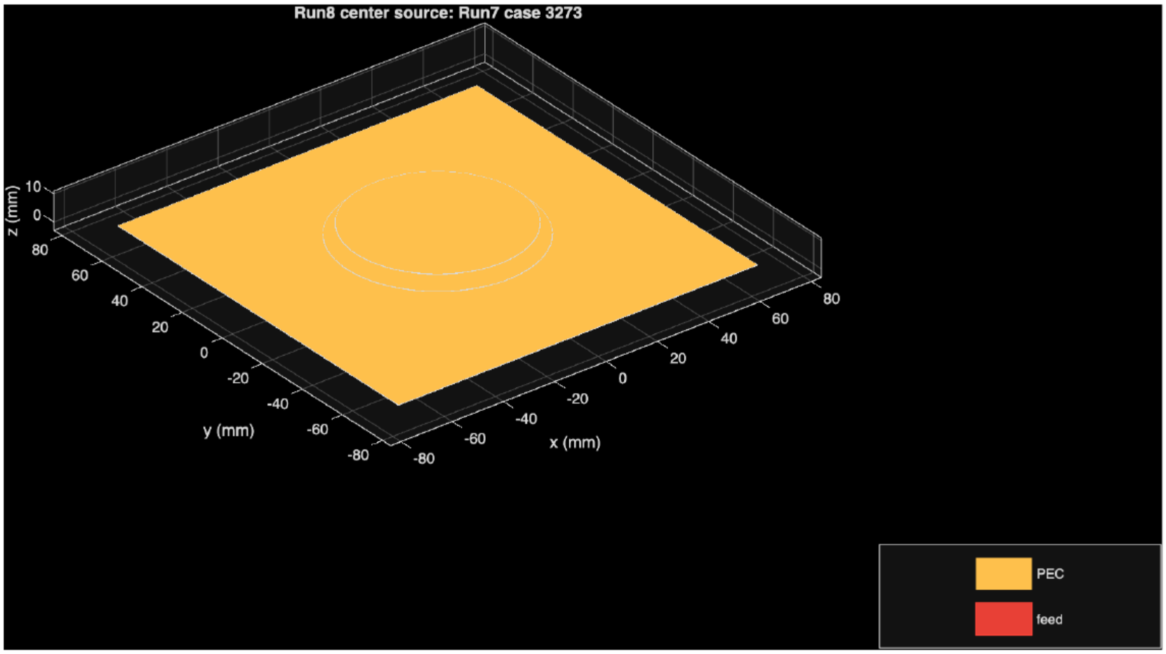

Figure 74. Architecture Compilation and Polarization Phase Check - Reference geometry: Center geometry run7 case3273.

In Figure 75, this plot carries the selected center-feed CP behavior into the architecture package and shows the polarization margin we were protecting. It is included here because it affected the next step in the design path: Use the compiled architecture as the feedhead input for dish/reflector validation.

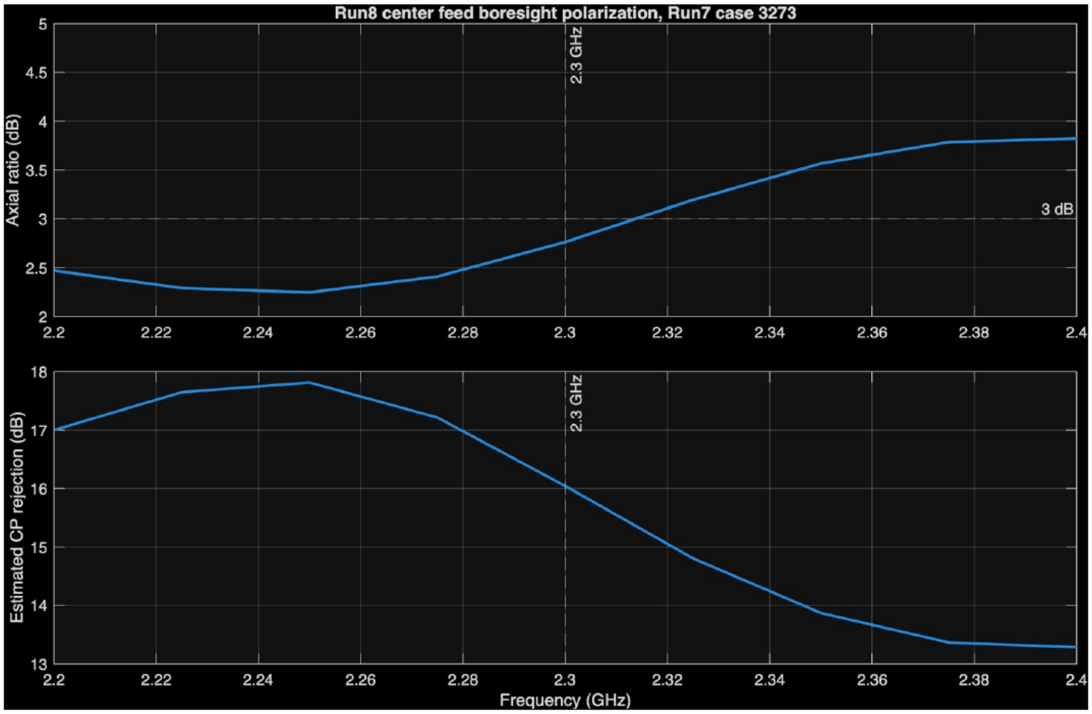

Figure 75. Architecture Compilation and Polarization Phase Check - Generated evidence: Center polarization run7 case3273.

In Figure 76, this plot carries the selected center-feed match into the architecture package without pretending a new hybrid solve was performed. Read against the section objective, it supports the decision at the end of this block: Run8 became the source architecture for reflector work and early fabrication packaging.

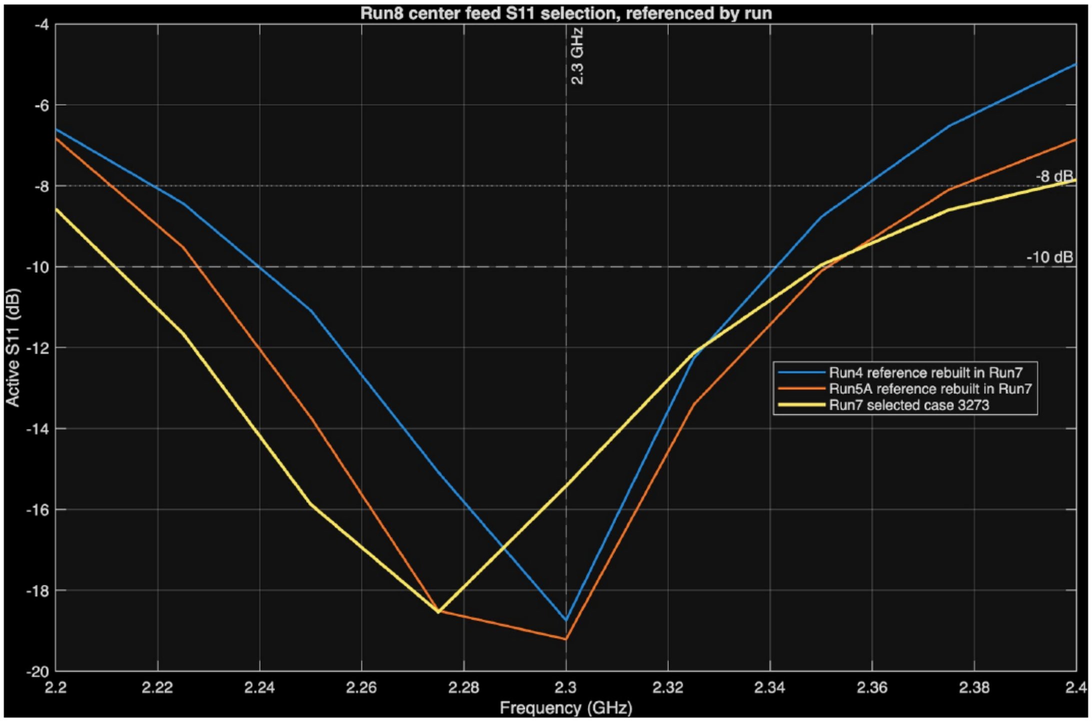

Figure 76. Architecture Compilation and Polarization Phase Check - Graph evidence: Center s11 curves run7.

In Figure 77, the coupling matrix is the sensor-coexistence test: useful tracking sensors need isolation and symmetry, not just individual match. Read against the section objective, it supports the decision at the end of this block: Run8 became the source architecture for reflector work and early fabrication packaging.

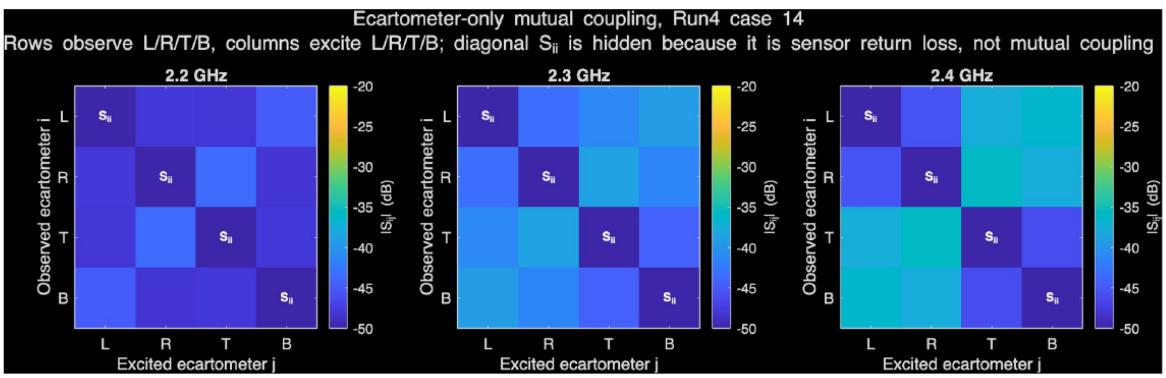

Figure 77. Architecture Compilation and Polarization Phase Check - Graph evidence: Ecartometer mutual coupling matrix run4 case14.

In Figure 78, these plots show the tracking-ring behavior that was carried into the architecture package from the selected ecartometer case. Read against the section objective, it supports the decision at the end of this block: Run8 became the source architecture for reflector work and early fabrication packaging.

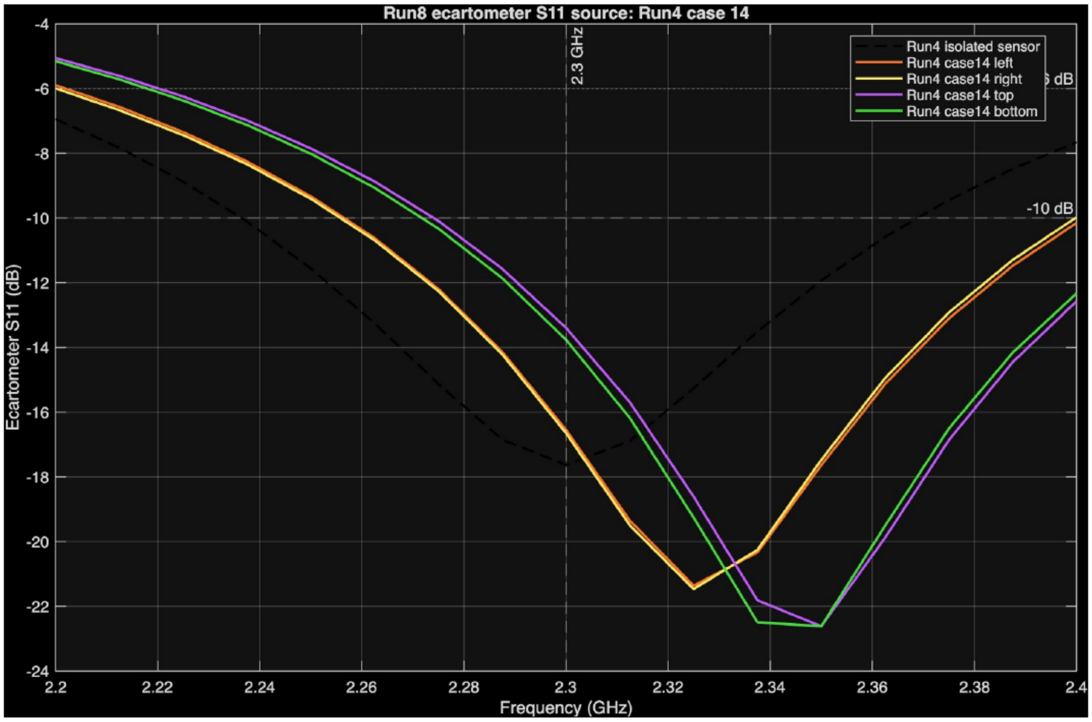

Figure 78. Architecture Compilation and Polarization Phase Check - Graph evidence: Ecartometer s11 curves run4 case14.

In Figure 79, the coupling matrix is the sensor-coexistence test: useful tracking sensors need isolation and symmetry, not just individual match. Read against the section objective, it supports the decision at the end of this block: Run8 became the source architecture for reflector work and early fabrication packaging.

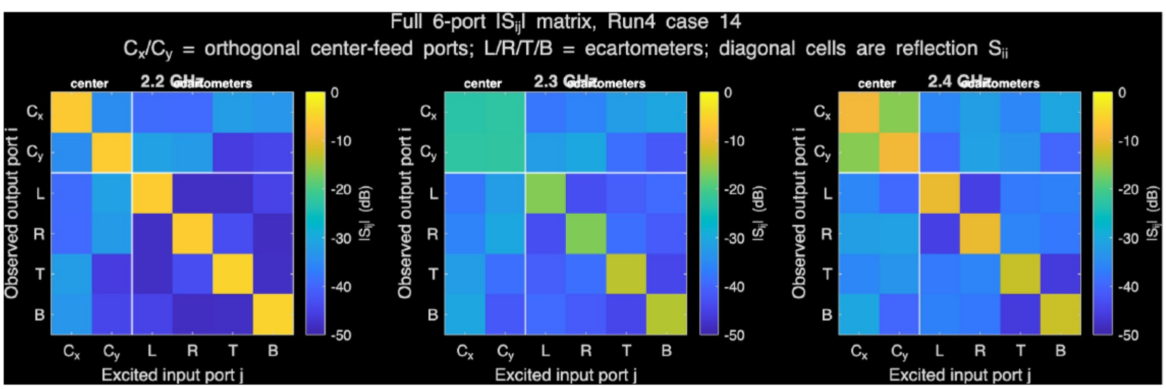

Figure 79. Architecture Compilation and Polarization Phase Check - Graph evidence: Full 6port coupling matrix run4 case14.

Design decision: Run8 became the source architecture for reflector work and early fabrication packaging. Next step: Use the compiled architecture as the feedhead input for dish/reflector validation.

18. School-Dish Reflector Validation

The design question was: Where does the feedhead sit in the real dish geometry, and what reflector performance is plausible?

We tried it because the feedhead only matters as part of a dish system; Run9 connected the selected feed geometry to the school parabola dimensions.

Trace source: Run 9 / parabola_dimensioned_reflector_validation

Scope

- Used Run8 feedhead geometry as source input.
- Swept F/D, focus sensitivity, blockage mask, and ideal parabola cross-sections.
- Kept the model as aperture/blockage/focus estimation, not a full reflector EM solve.

Evidence and Interpretation

In Figure 80, the blockage mask translates the feedhead footprint into reflector loss. It is where mechanical size starts affecting gain. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Where does the feedhead sit in the real dish geometry, and what reflector performance is plausible?

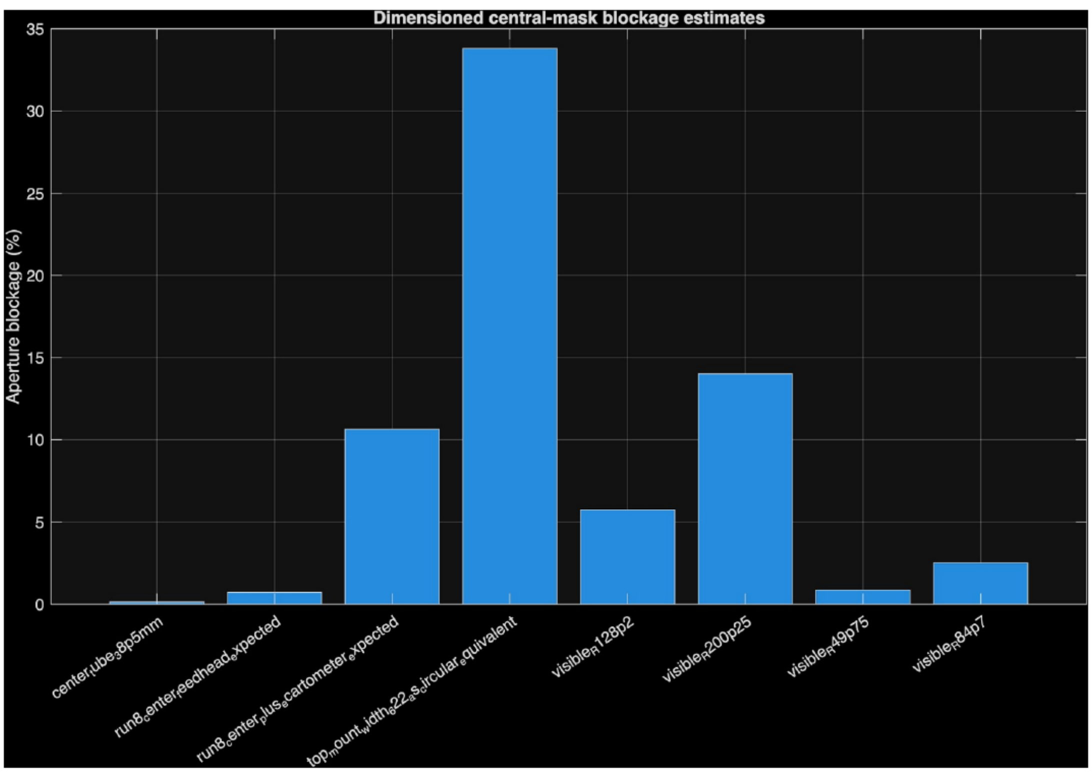

Figure 80. School-Dish Reflector Validation - Reference geometry: Blockage mask summary.

In Figure 81, this figure shows which feedhead geometry was used as the reflector input, keeping the dish estimates tied to the selected architecture. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Where does the feedhead sit in the real dish geometry, and what reflector performance is plausible?

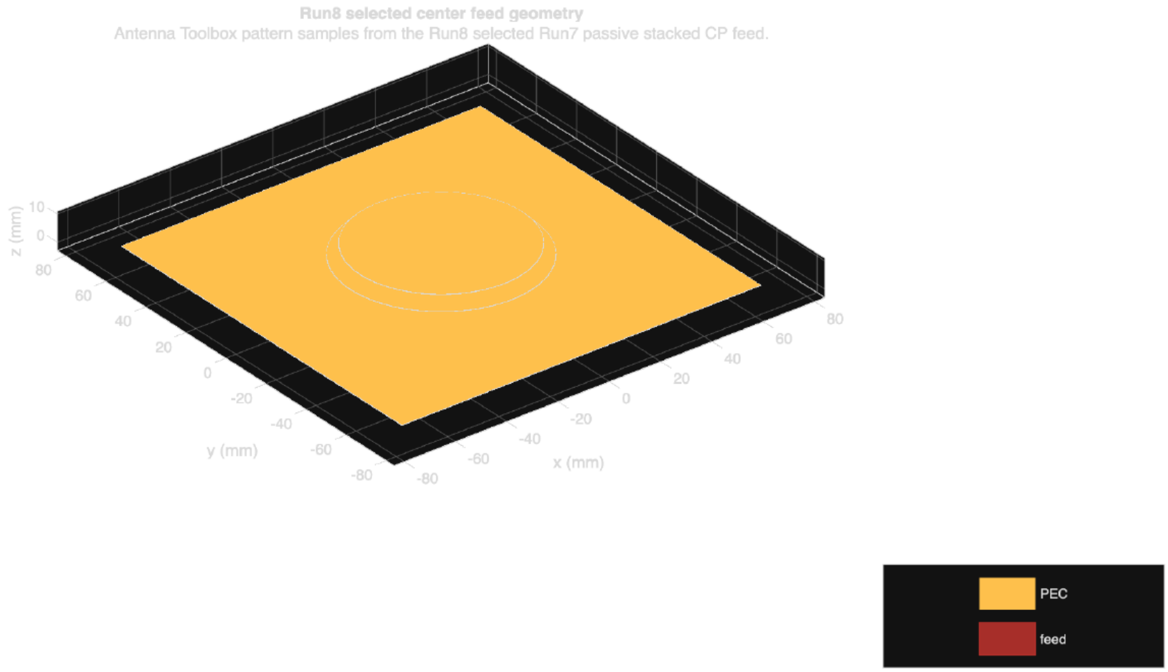

Figure 81. School-Dish Reflector Validation - Reference geometry: Source run8 center feed geometry.

In Figure 82, this geometry anchors the reflector work in the real dish dimensions rather than an abstract feed-only antenna. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Where does the feedhead sit in the real dish geometry, and what reflector performance is plausible?

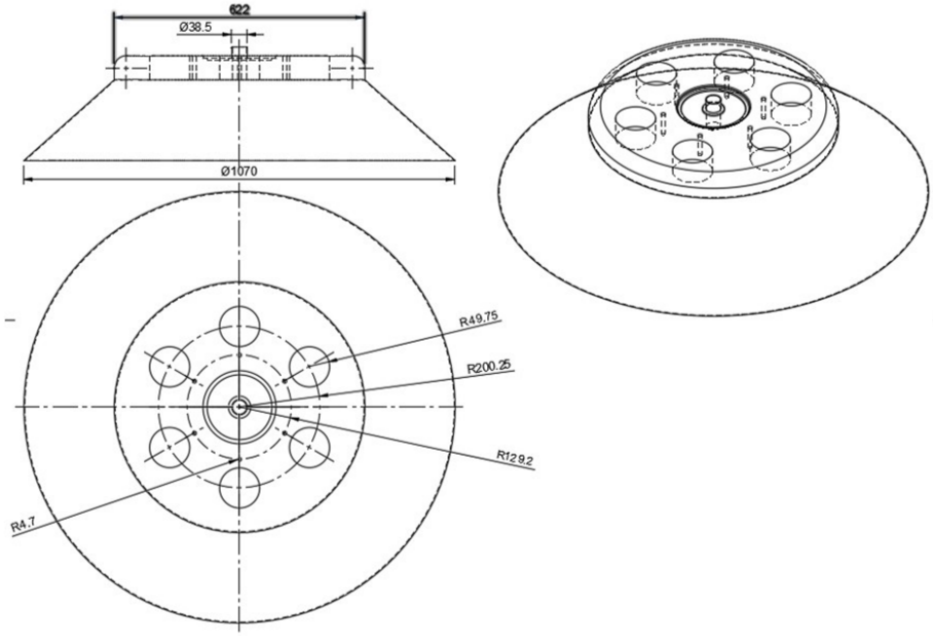

Figure 82. School-Dish Reflector Validation - Reference geometry: Source parabola dimensions.

In Figure 83, this plot compares F/D choices, gain, and beamwidth so the reflector decision is not based on gain alone. Read against the section objective, it supports the decision at the end of this block: F/D 0.45 and focus 481.5 mm became the school-dish baseline.

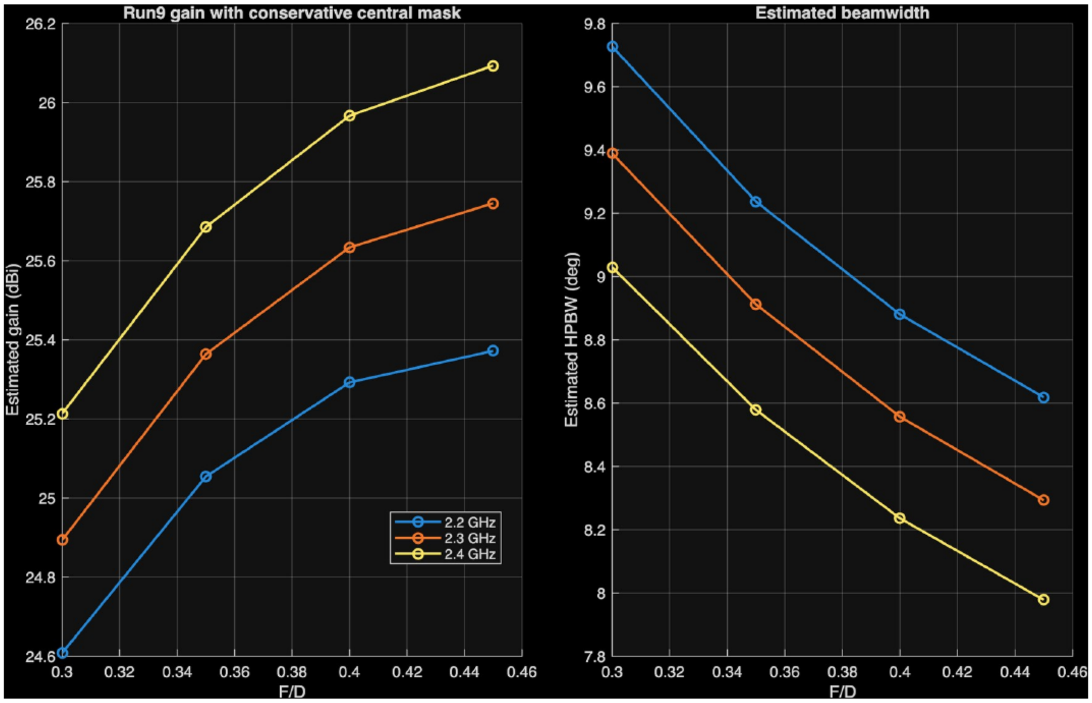

Figure 83. School-Dish Reflector Validation - Graph evidence: Fd gain beamwidth summary.

In Figure 84, this plot shows how much gain is lost when the feed moves away from focus. It sets the practical mounting tolerance. Read against the section objective, it supports the decision at the end of this block: F/D 0.45 and focus 481.5 mm became the school-dish baseline.

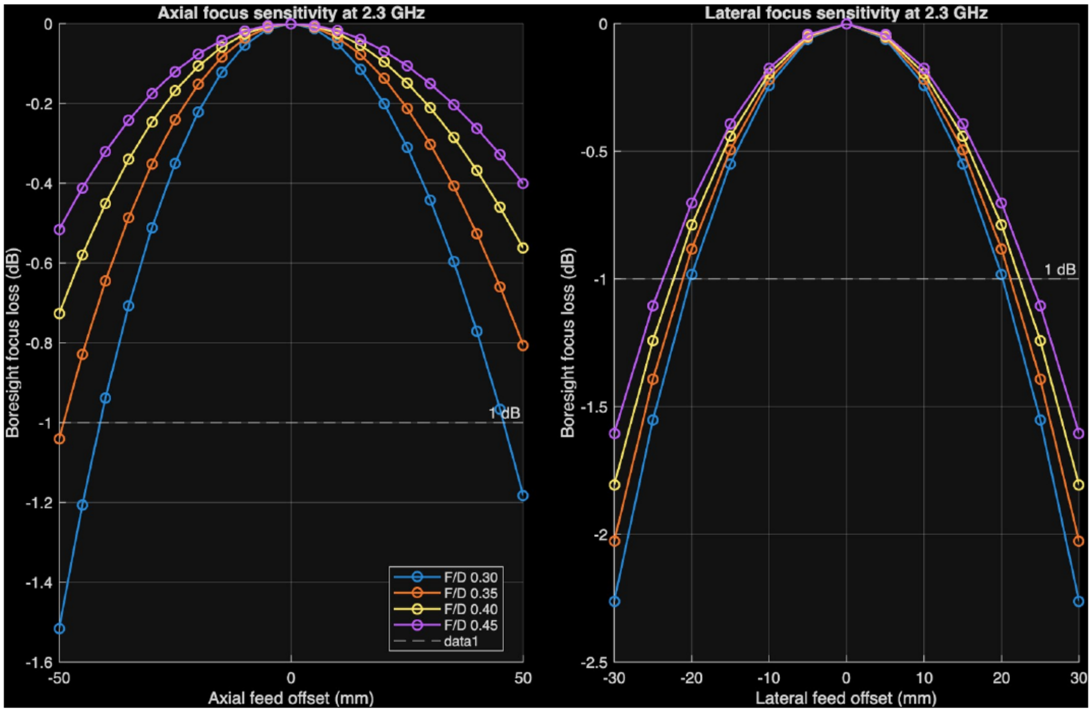

Figure 84. School-Dish Reflector Validation - Graph evidence: Focus sensitivity 2.3GHz.

In Figure 85, the parabola cross-sections convert F/D into actual depth and focus geometry for fabrication and mounting. It is included here because it affected the next step in the design path: Apply the full feedhead mask to the school parabola and then optimize reflector diameter/F/D tradeoffs.

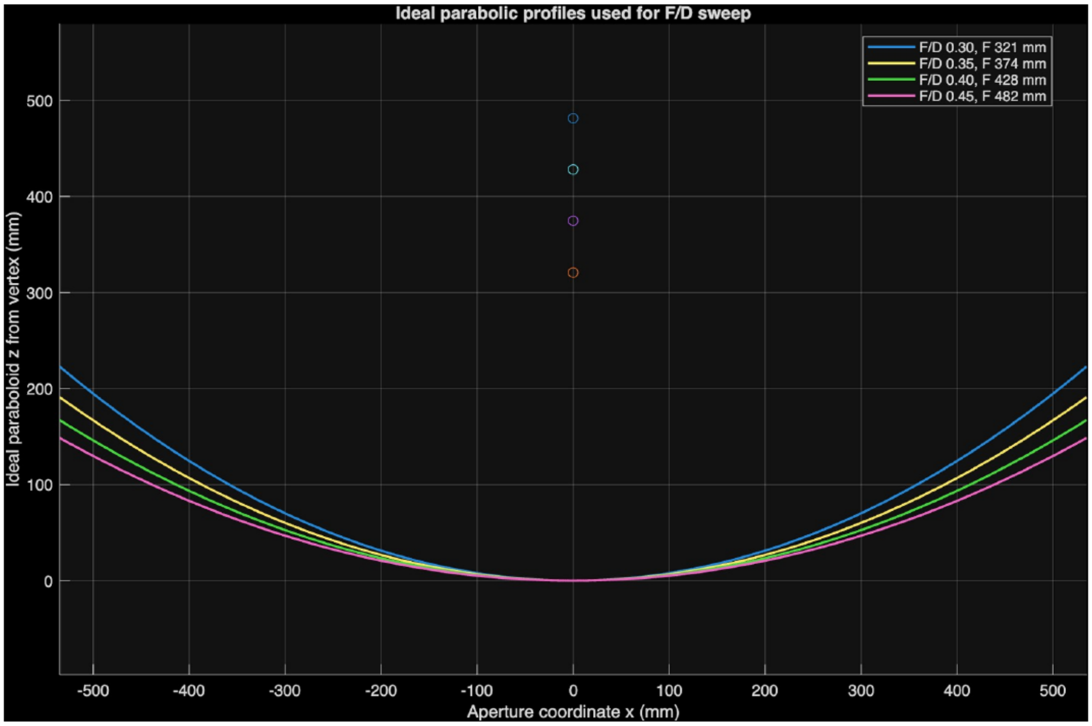

Figure 85. School-Dish Reflector Validation - Generated evidence: Ideal parabola cross sections.

Design decision: F/D 0.45 and focus 481.5 mm became the school-dish baseline. Next step: Apply the full feedhead mask to the school parabola and then optimize reflector diameter/F/D tradeoffs.

19. Full Feedhead on School Parabola

The design question was: What happens when the full feedhead mask, sensors, blockage, focus, and ecartometer response are applied to the school dish?

We tried it because run9 used simplified feed references; Run10A applied the full feedhead footprint and tracking model to the school-dish baseline.

Trace source: Run 10A / full_feedhead_school_parabola

Scope

- Estimated reflector gain, blockage, axial/lateral focus tolerance, and azimuth/elevation error behavior.
- Used Run8 center-feed and ecartometer summaries as staged inputs.
- Tracked warnings where symmetry spread increased outside the selected sensor distance.

Evidence and Interpretation

In Figure 86, this is the full feedhead mask used in the school-dish estimate. It connects the RF feed package to physical blockage and focal-plane behavior. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What happens when the full feedhead mask, sensors, blockage, focus, and ecartometer response are applied to the school dish?

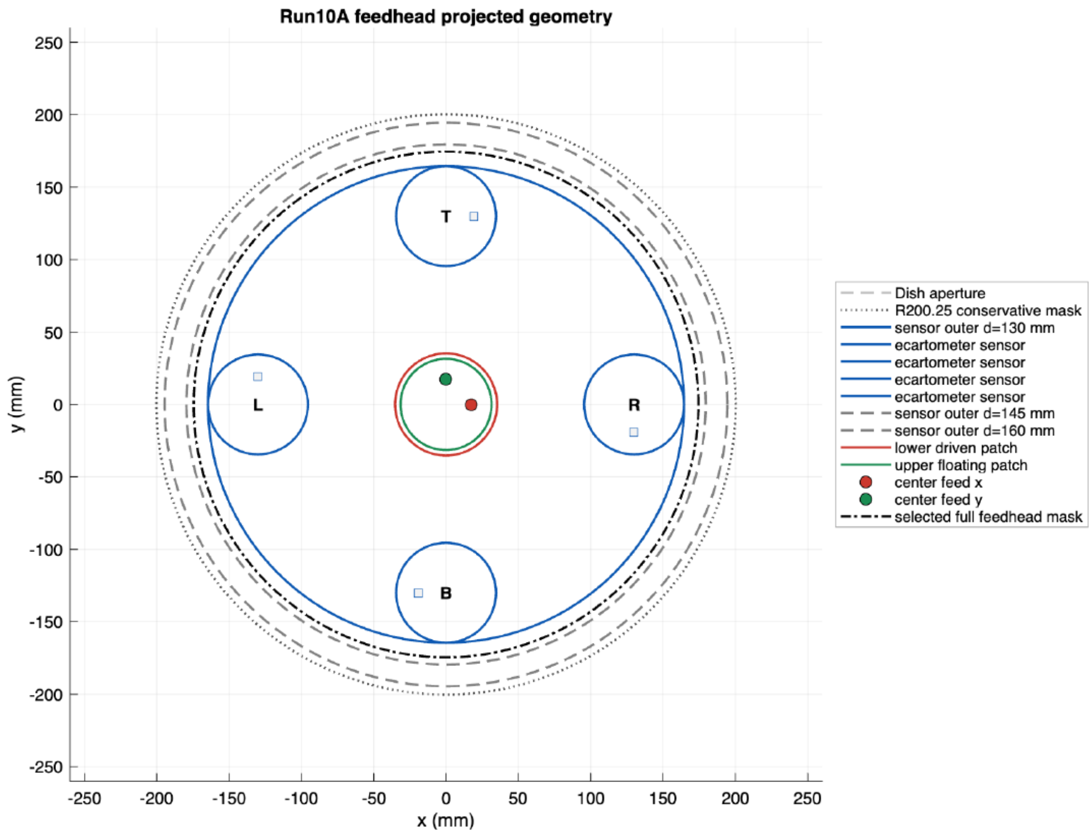

Figure 86. Full Feedhead on School Parabola - Reference geometry: Feedhead geometry.

In Figure 87, this azimuth delta plot is part of the ecartometer tracking readout: it shows the differential signal available for pointing correction. Read against the section objective, it supports the decision at the end of this block: The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45.

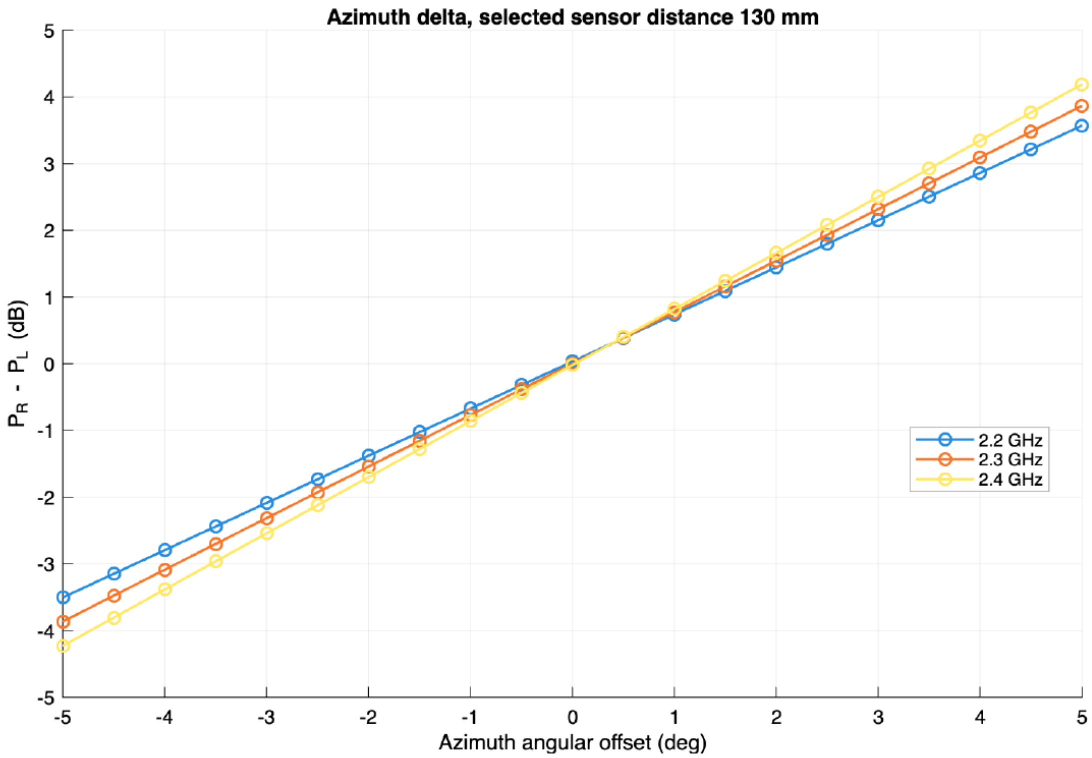

Figure 87. Full Feedhead on School Parabola - Graph evidence: Az delta dB vs angle.

In Figure 88, this azimuth error curve converts differential signal behavior into estimated pointing error. It is included here because it affected the next step in the design path: Search larger reflector geometries to see what performance is available beyond the school dish.

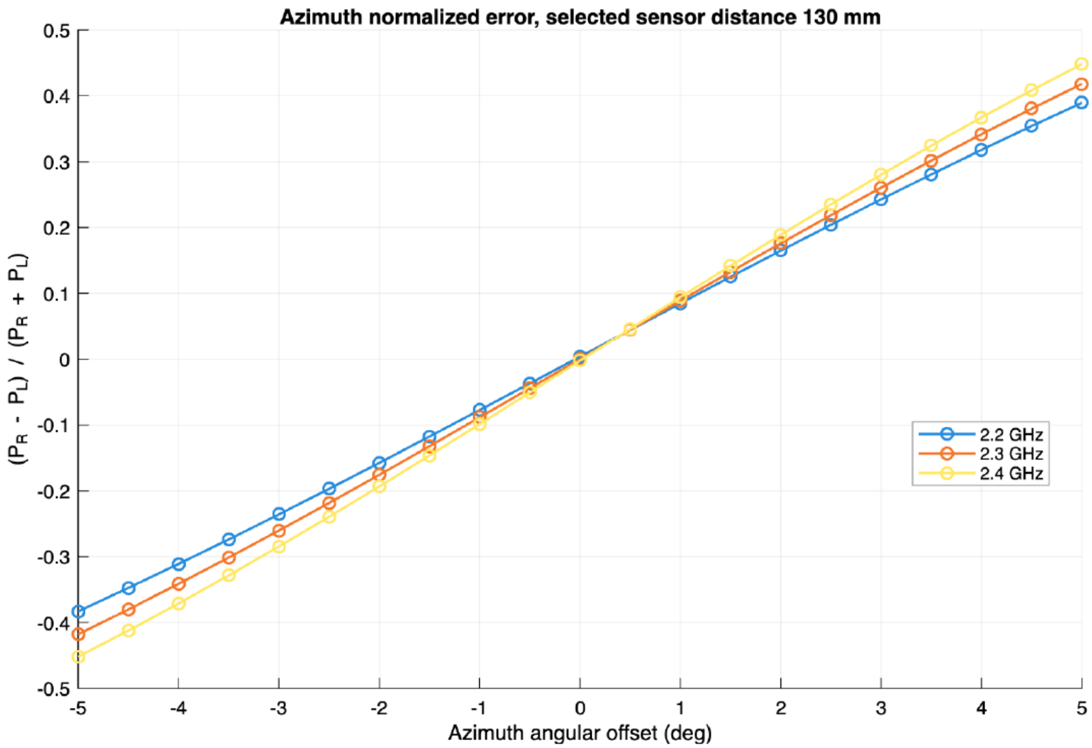

Figure 88. Full Feedhead on School Parabola - Generated evidence: Az error vs angle.

In Figure 89, this comparison shows how much the full feedhead footprint costs relative to simpler blockage assumptions. Read against the section objective, it supports the decision at the end of this block: The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45.

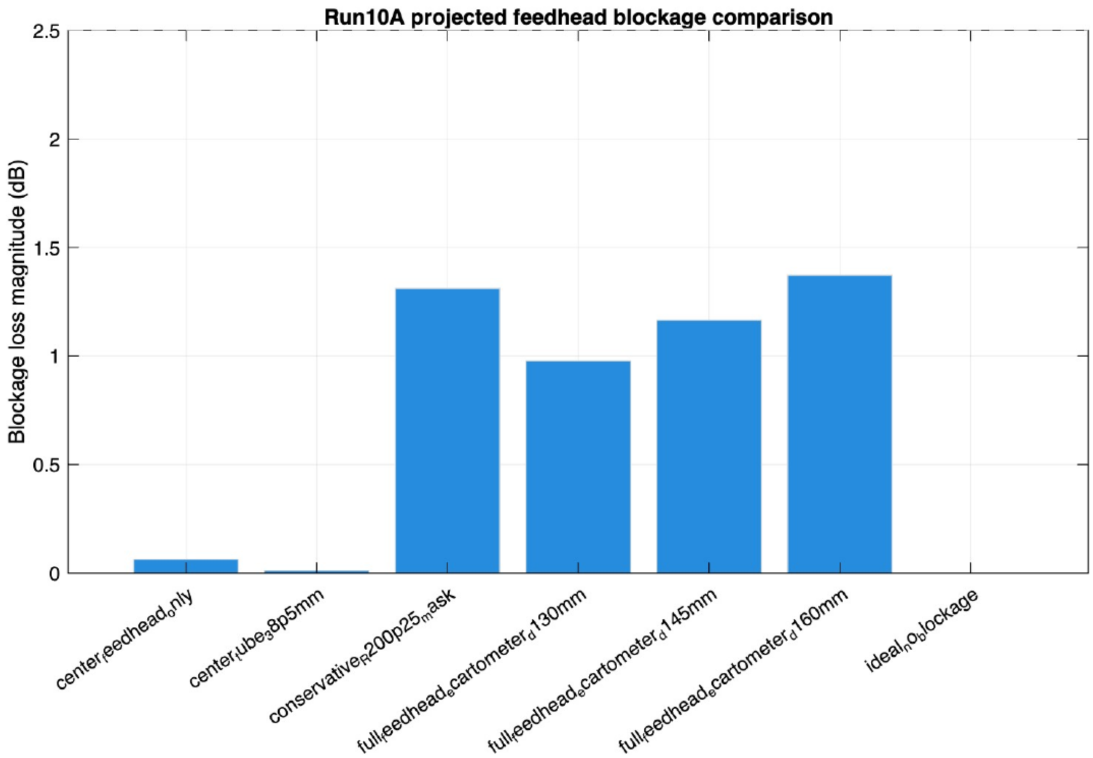

Figure 89. Full Feedhead on School Parabola - Graph evidence: Blockage comparison.

In Figure 90, this elevation delta plot is the companion tracking readout for the orthogonal sensor pair. Read against the section objective, it supports the decision at the end of this block: The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45.

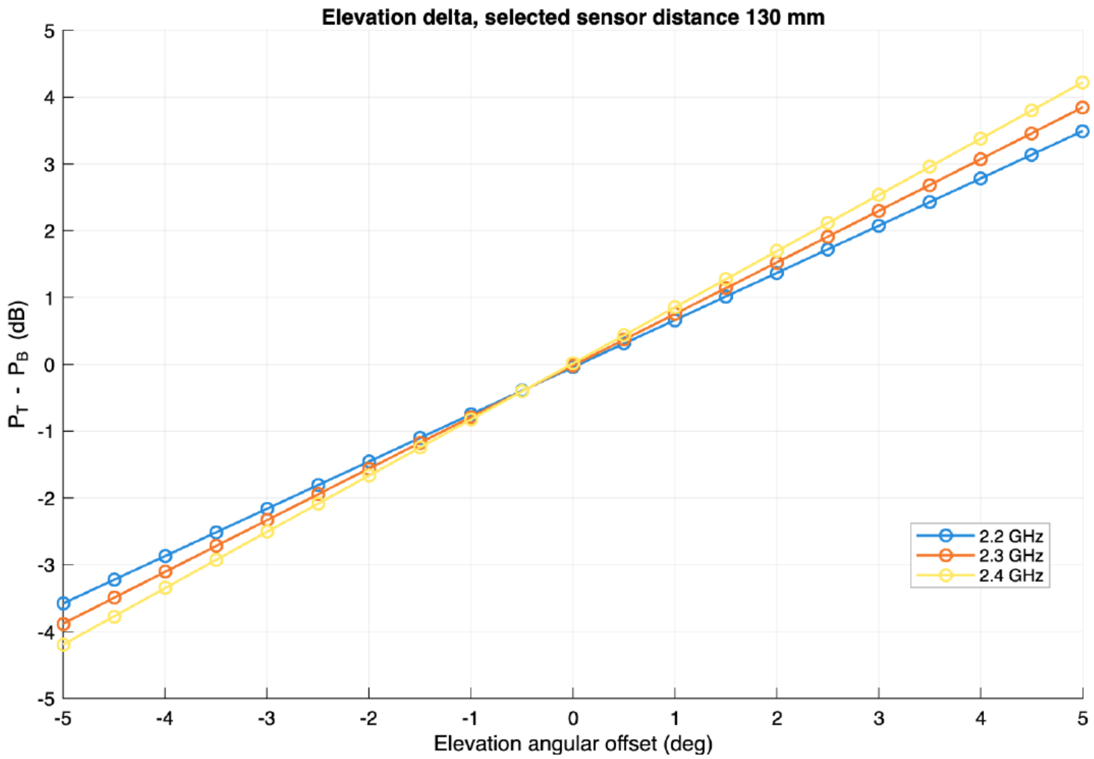

Figure 90. Full Feedhead on School Parabola - Graph evidence: El delta dB vs angle.

In Figure 91, this elevation error curve converts the elevation-channel response into estimated pointing error. It is included here because it affected the next step in the design path: Search larger reflector geometries to see what performance is available beyond the school dish.

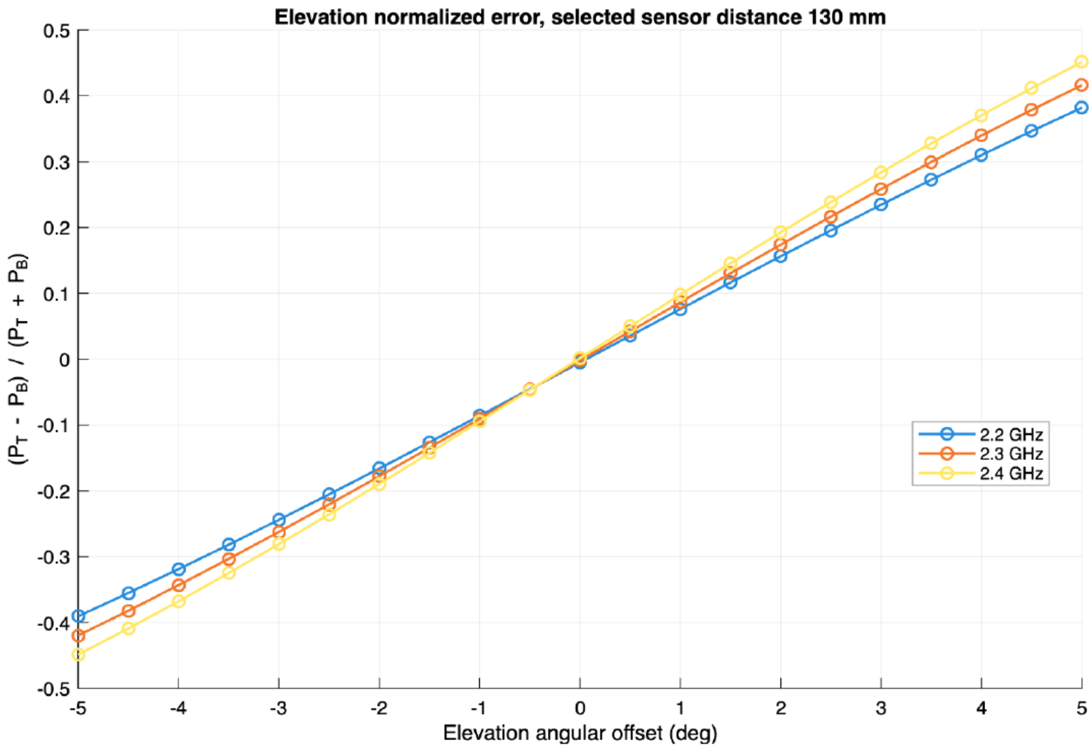

Figure 91. Full Feedhead on School Parabola - Generated evidence: El error vs angle.

In Figure 92, the axial-offset curve shows how forgiving the reflector is to feed placement along the boresight. Read against the section objective, it supports the decision at the end of this block: The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45.

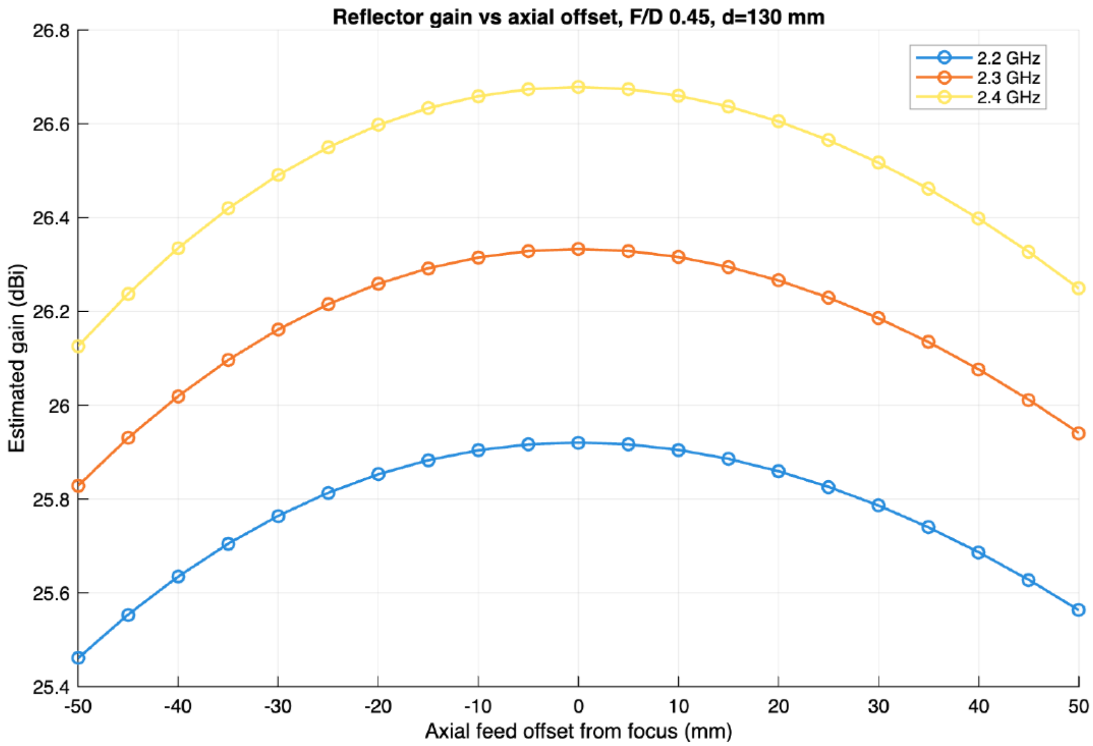

Figure 92. Full Feedhead on School Parabola - Graph evidence: Reflector gain vs axial offset.

In Figure 93, the lateral-offset curve shows how forgiving the reflector is to sideways feed placement, which is usually the tighter mechanical tolerance. Read against the section objective, it supports the decision at the end of this block: The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45.

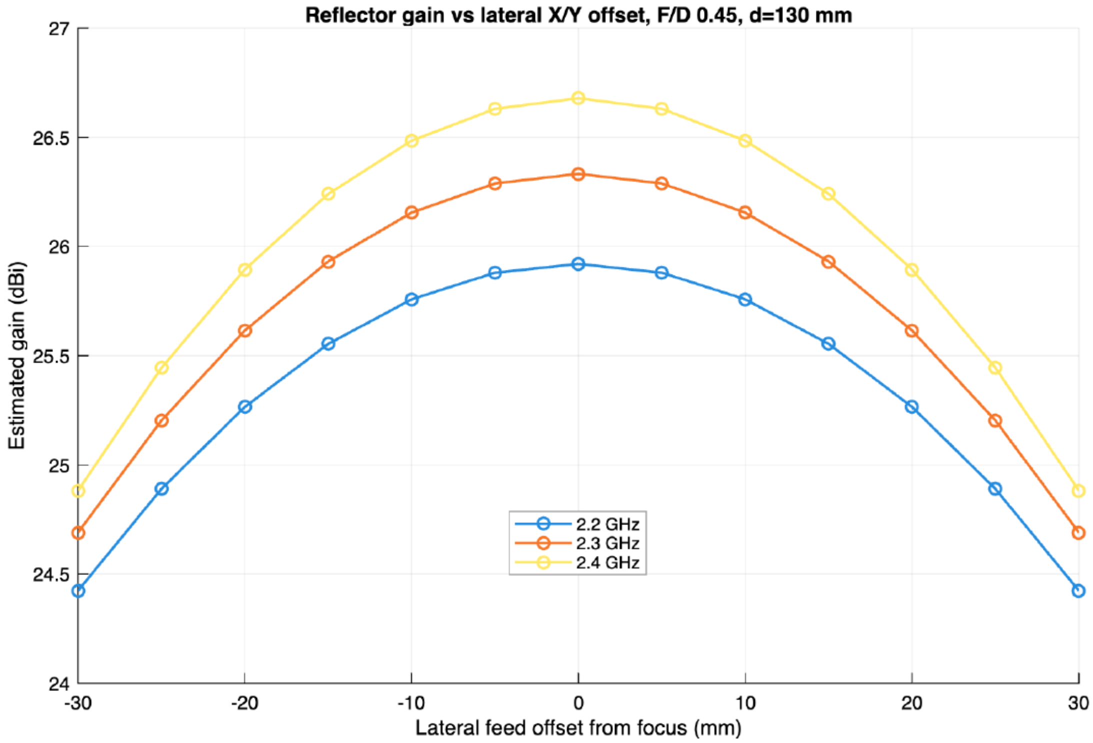

Figure 93. Full Feedhead on School Parabola - Graph evidence: Reflector gain vs lateral offset.

> **Design decision:** The school dish remained usable as a validation baseline, with estimated 26.33 dBi gain at 2.3 GHz and F/D 0.45. Next step: Search larger reflector geometries to see what performance is available beyond the school dish.

20. Reflector Geometry Optimization

The design question was: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

We tried it because the school dish was available, but not necessarily optimal. This sweep quantified the trade space.

Trace source: Run 10B / optimize_parabola_geometry_for_feedhead

Scope

- Swept reflector diameter and F/D with center-only, full-feedhead, and conservative blockage masks.
- Compared gain, blockage loss, axial/lateral tolerance, feed pattern references, and ecartometer slope.

- Ranked top candidates under full feedhead/ecartometer mask assumptions.

Evidence and Interpretation

In Figure 94, this plot prevents a naive maximum-diameter answer by showing blockage loss across the reflector trade space. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

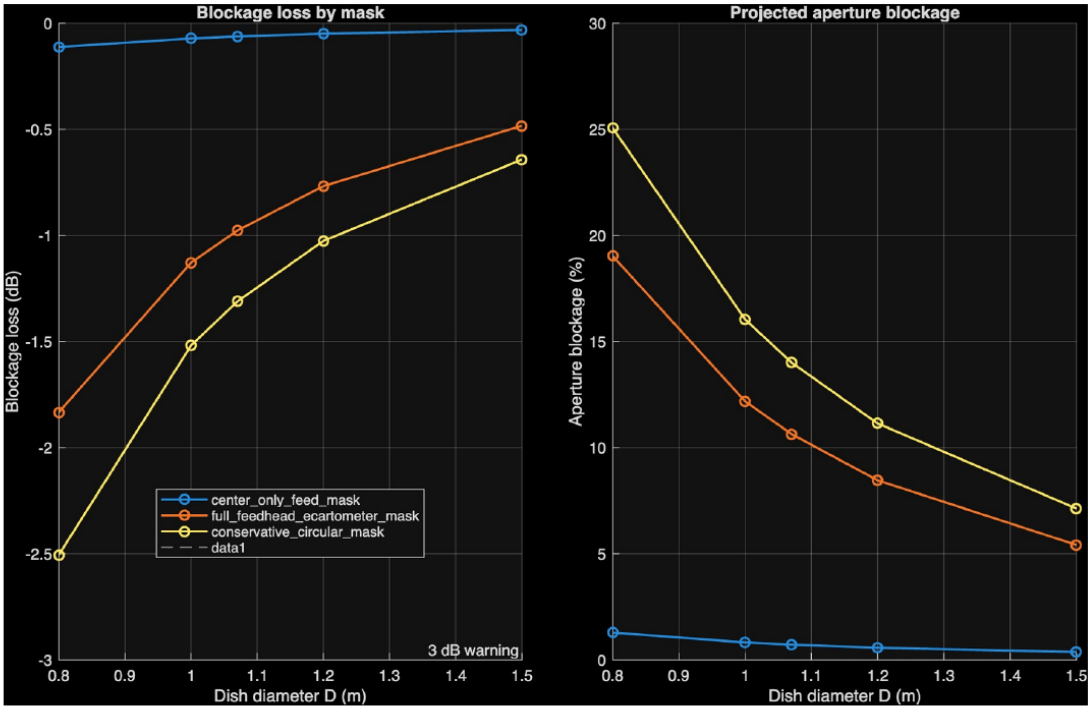

Figure 94. Reflector Geometry Optimization - Reference geometry: Blockage loss vs diameter fd.

In Figure 95, this plot keeps tracking performance in the reflector optimization instead of optimizing only payload gain. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

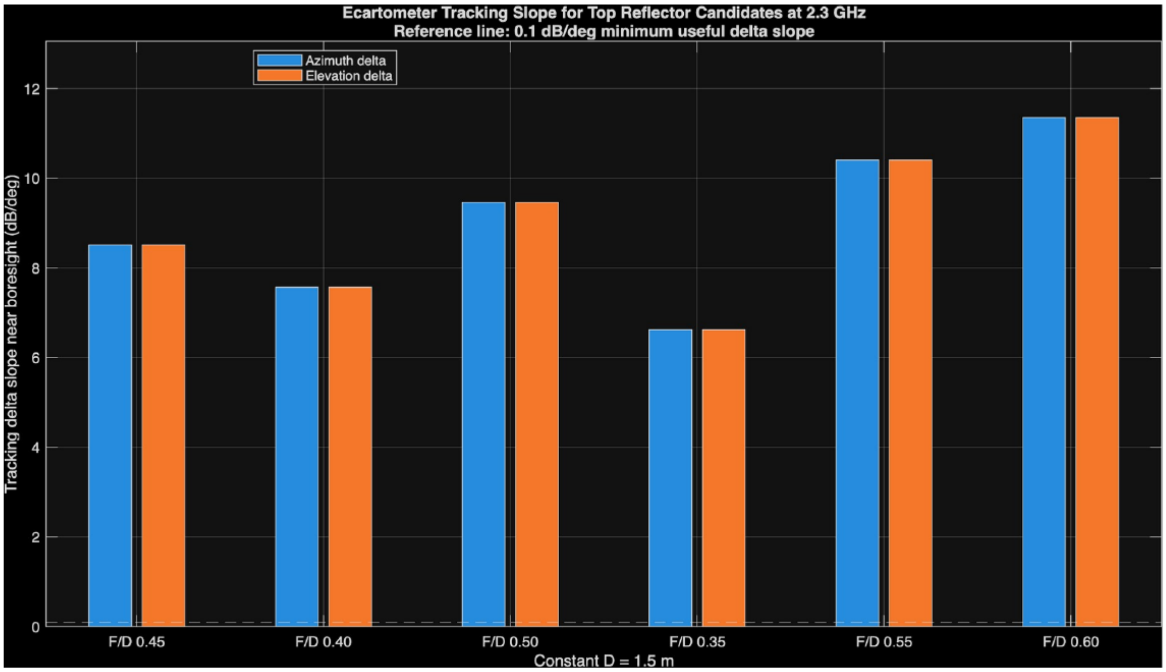

Figure 95. Reflector Geometry Optimization - Reference geometry: Ecartometer slope top candidates.

In Figure 96, this axial-ratio reference checks that reflector optimization is not detached from the selected CP feed quality. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

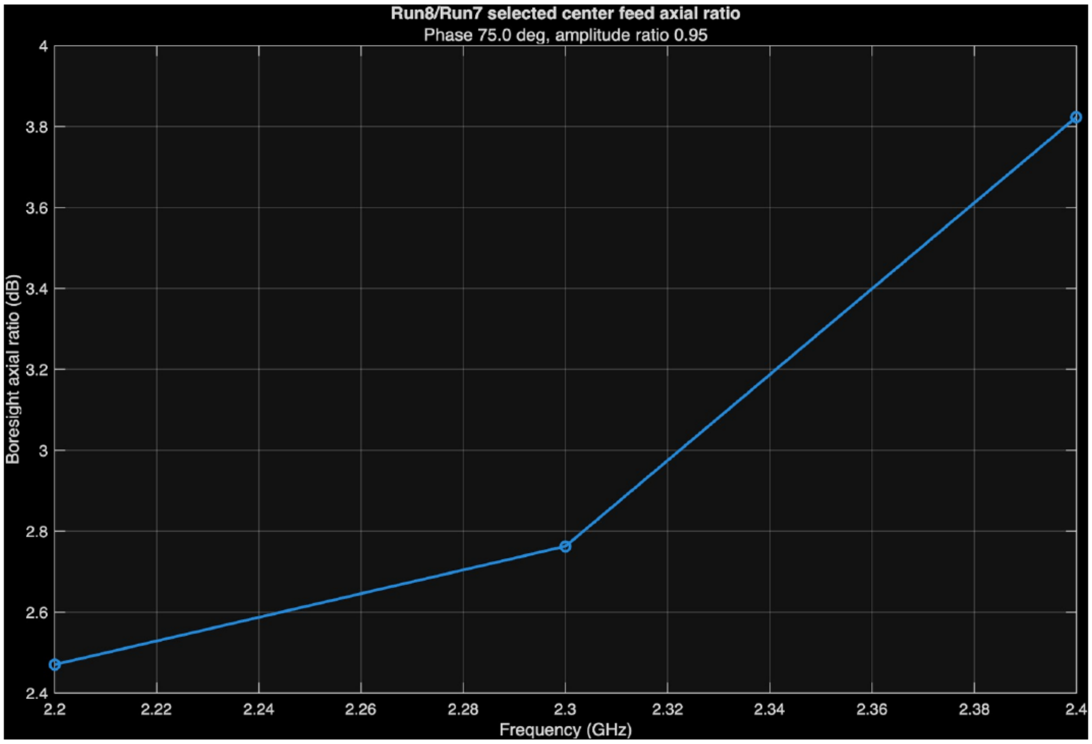

Figure 96. Reflector Geometry Optimization - Reference geometry: Feed pattern axial ratio reference.

In Figure 97, this feed-pattern reference anchors the reflector estimate to the selected feed behavior. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

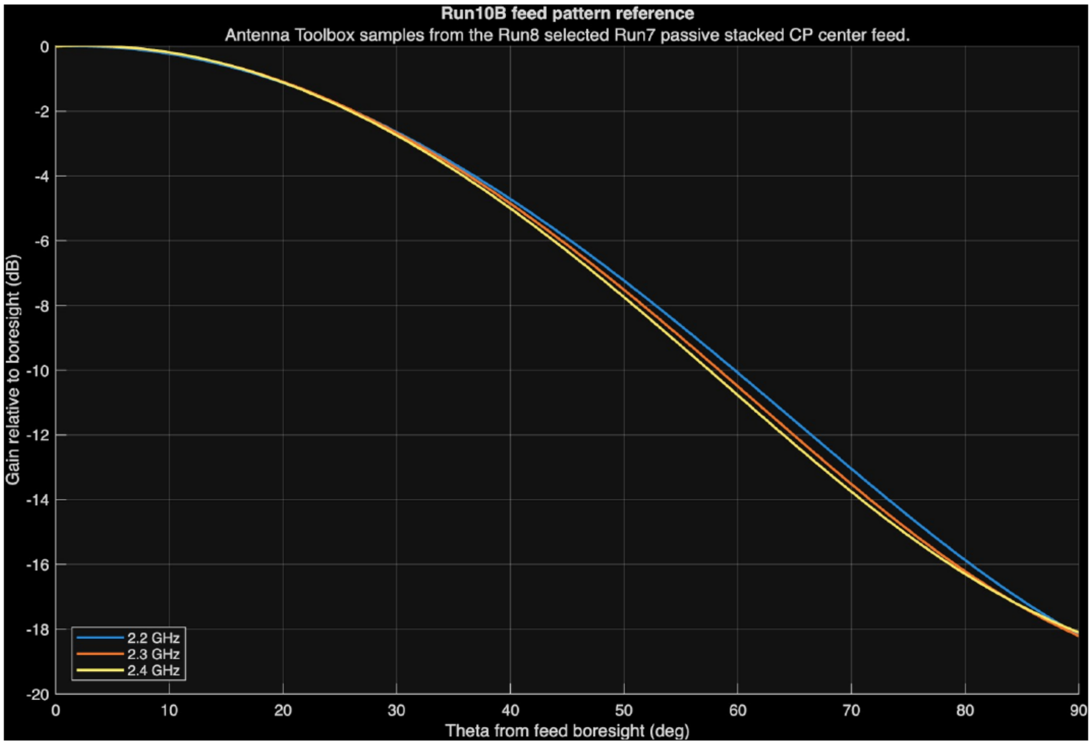

Figure 97. Reflector Geometry Optimization - Reference geometry: Feed pattern gain reference.

In Figure 98, this plot ranks reflector candidates by axial tolerance, which matters for a buildable focus mount. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

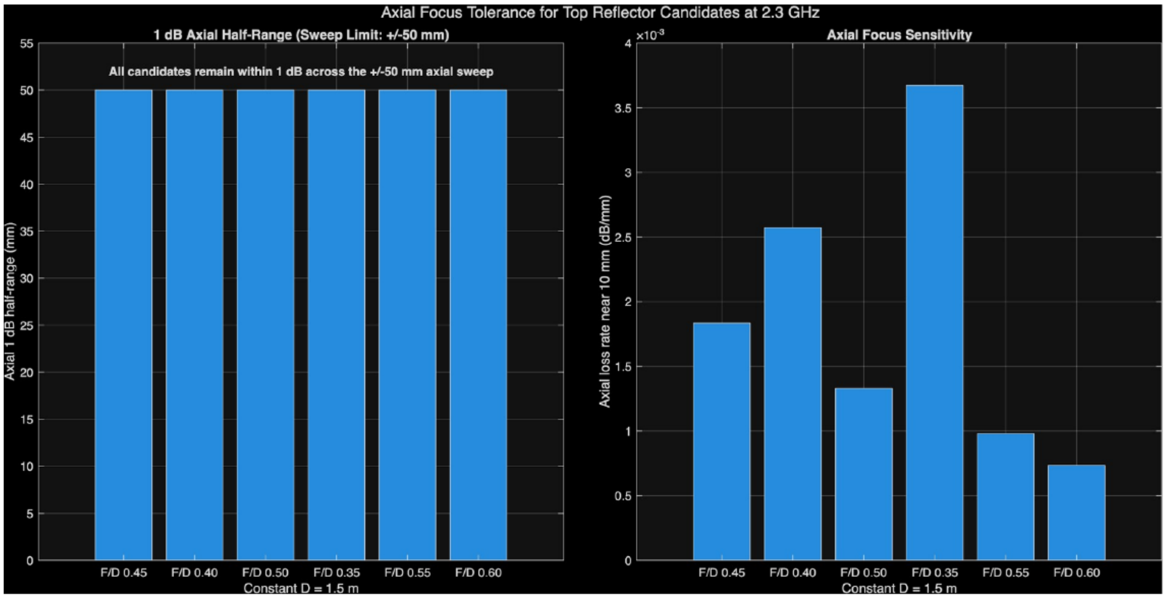

Figure 98. Reflector Geometry Optimization - Reference geometry: Focus tolerance top candidates.

In Figure 99, this optimization plot shows why a larger reflector option was considered: diameter and F/D trade directly against gain and packaging. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

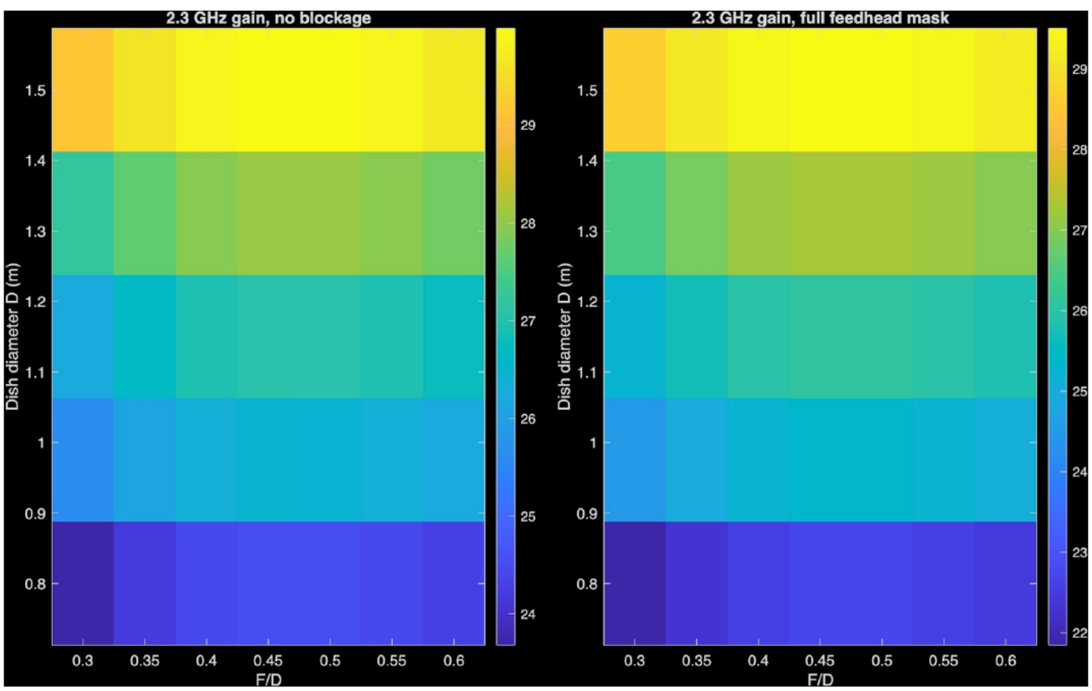

Figure 99. Reflector Geometry Optimization - Reference geometry: Gain vs diameter fd.

In Figure 100, this plot ranks reflector candidates by lateral tolerance, which matters for field alignment and structural deflection. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: If reflector size is not fixed, which diameter and F/D best balance gain, blockage, focus tolerance, and ecartometer slope?

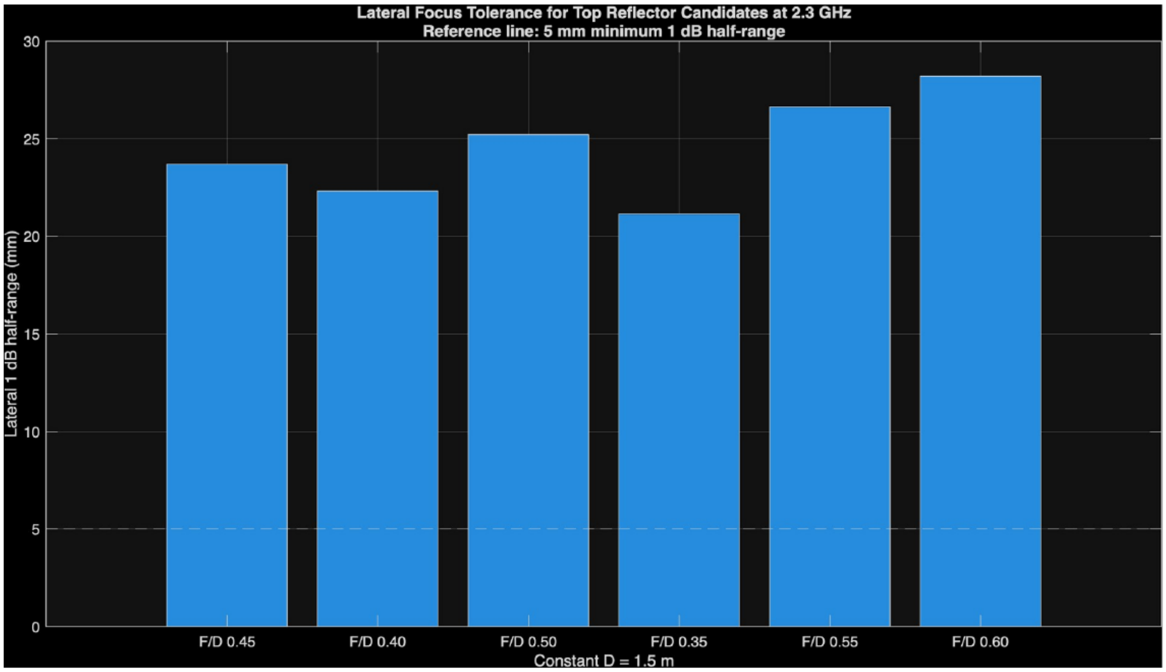

Figure 100. Reflector Geometry Optimization - Reference geometry: Lateral tolerance top candidates.

Design decision: D1.50 m at F/D 0.45 ranked highest under the full feedhead mask, with estimated 29.51 dBi at 2.3 GHz. Next step: Keep the school dish as practical baseline and the 1.50 m F/D 0.45 reflector as the optimized option.

21. Center-Feed RHCP/LHCP Sense Check

The design question was: Can the selected center feed switch circular-polarization sense cleanly by changing feed phase or mirrored feed point?

We tried it because selectable CP became a hardware requirement, so we needed to understand which feed choices actually changed RHCP/LHCP dominance.

Trace source: Run 11 / center_feed_rhcp_lhcp_sense_check

Scope

- Compared nominal and mirrored feed geometry.
- Evaluated active S11, axial ratio, CP purity, boresight RHCP/LHCP, and pattern cuts.
- Explicitly excluded ecartometer sensors and the dish reflector from this center-only sense check.

Evidence and Interpretation

In Figure 101, the mirrored-feed geometry tests whether physical feed selection can change CP sense more effectively than phase sign alone. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the selected center feed switch circular-polarization sense cleanly by changing feed phase or mirrored feed point?

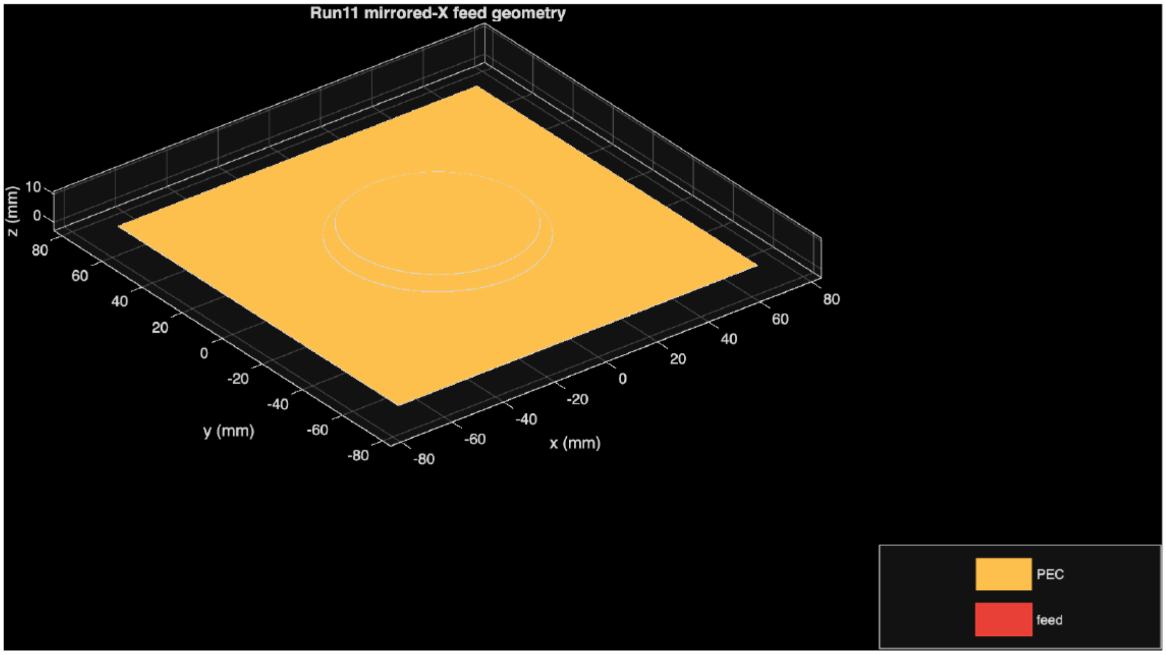

Figure 101. Center-Feed RHCP/LHCP Sense Check - Reference geometry: Geometry mirrored.

In Figure 102, the nominal-feed geometry is the first CP-sense reference: it shows the feed placement before mirrored-feed testing. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the selected center feed switch circular-polarization sense cleanly by changing feed phase or mirrored feed point?

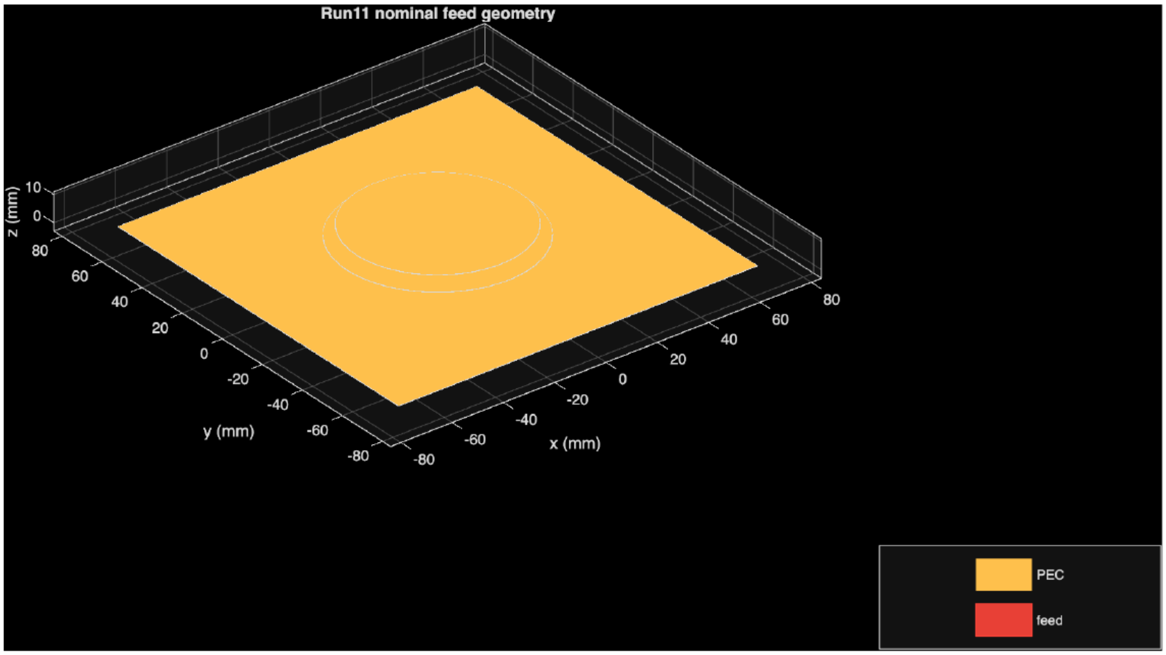

Figure 102. Center-Feed RHCP/LHCP Sense Check - Reference geometry: Geometry nominal.

In Figure 103, this case comparison checks that a CP switching option does not create an unacceptable input match penalty. Read against the section objective, it supports the decision at the end of this block: Phase sign alone was not enough; mirroring one feed point changed sense more usefully.

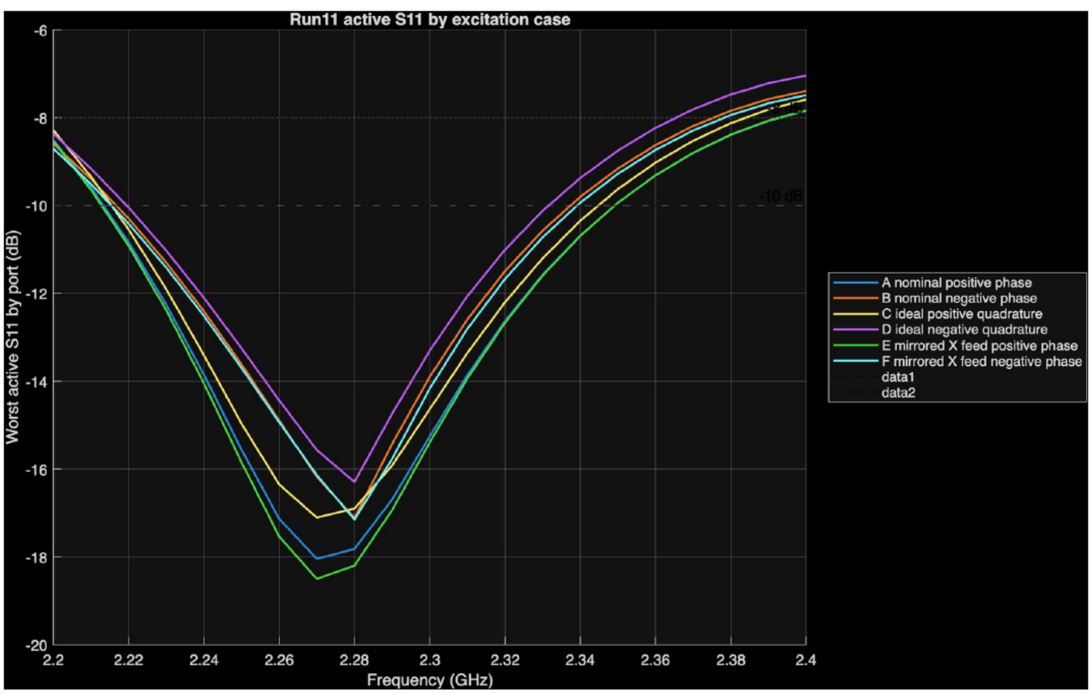

Figure 103. Center-Feed RHCP/LHCP Sense Check - Graph evidence: Active s11 by case.

In Figure 104, this case comparison checks that the switching option still preserves usable CP quality. Read against the section objective, it supports the decision at the end of this block: Phase sign alone was not enough; mirroring one feed point changed sense more usefully.

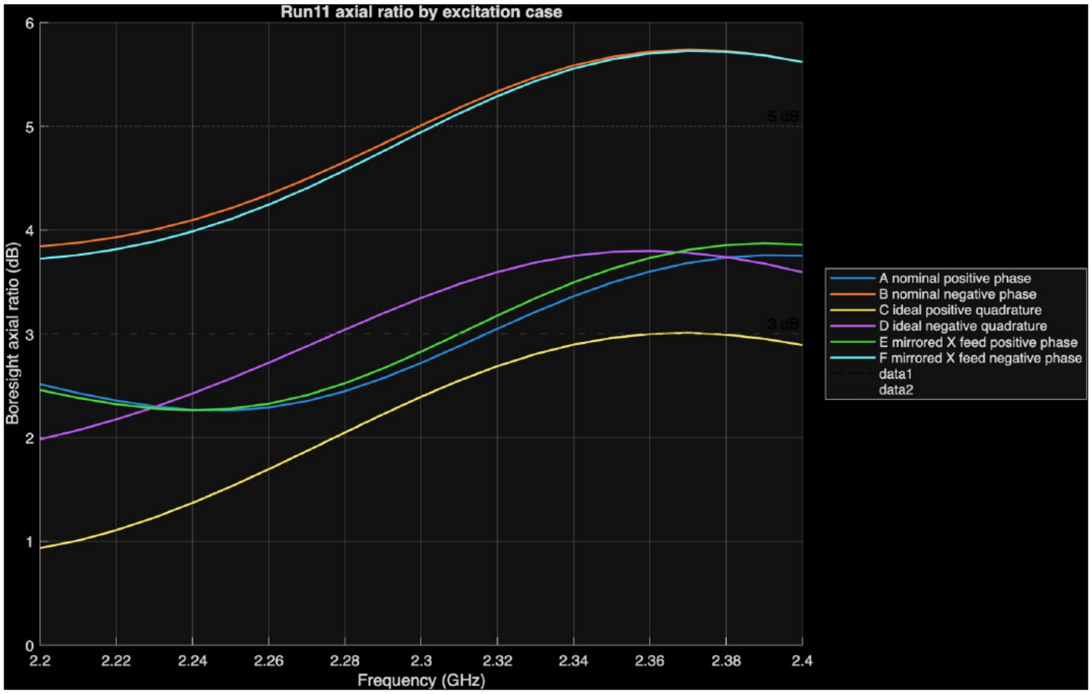

Figure 104. Center-Feed RHCP/LHCP Sense Check - Graph evidence: Axial ratio by case.

In Figure 105, the pattern cut shows whether the selected CP case remains plausible away from boresight instead of only at one sampled angle. Read against the section objective, it supports the decision at the end of this block: Phase sign alone was not enough; mirroring one feed point changed sense more usefully.

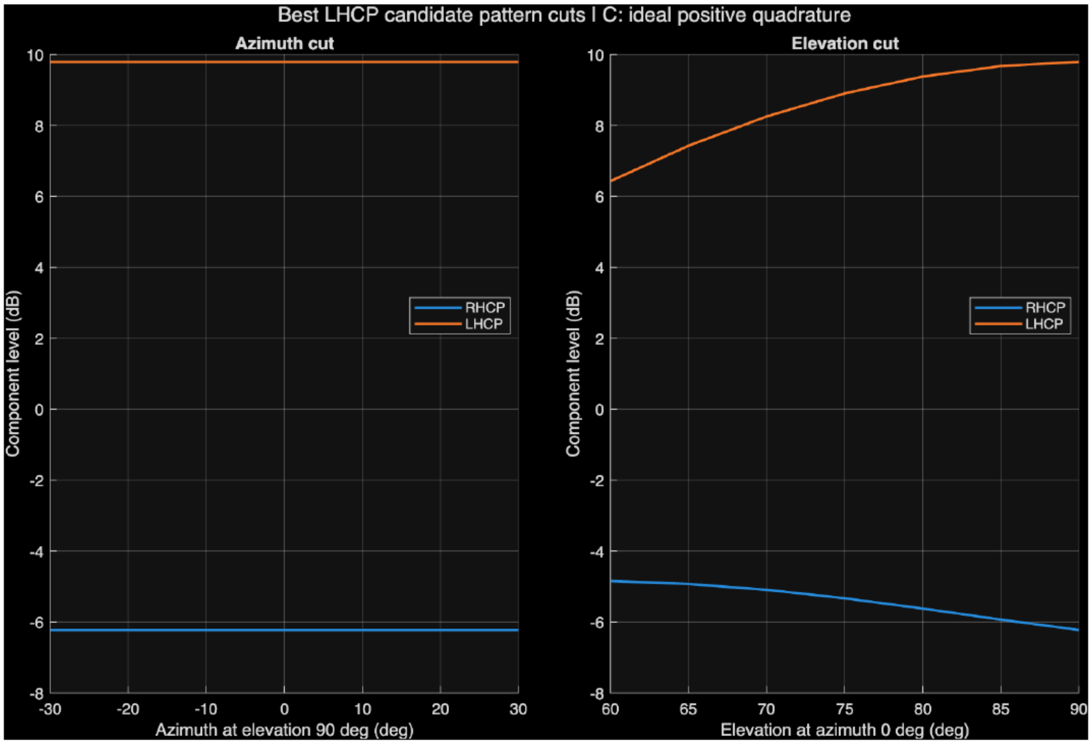

Figure 105. Center-Feed RHCP/LHCP Sense Check - Graph evidence: Best lhcp pattern cut.

In Figure 106, the pattern cut shows whether the selected CP case remains plausible away from boresight instead of only at one sampled angle. Read against the section objective, it supports the decision at the end of this block: Phase sign alone was not enough; mirroring one feed point changed sense more usefully.

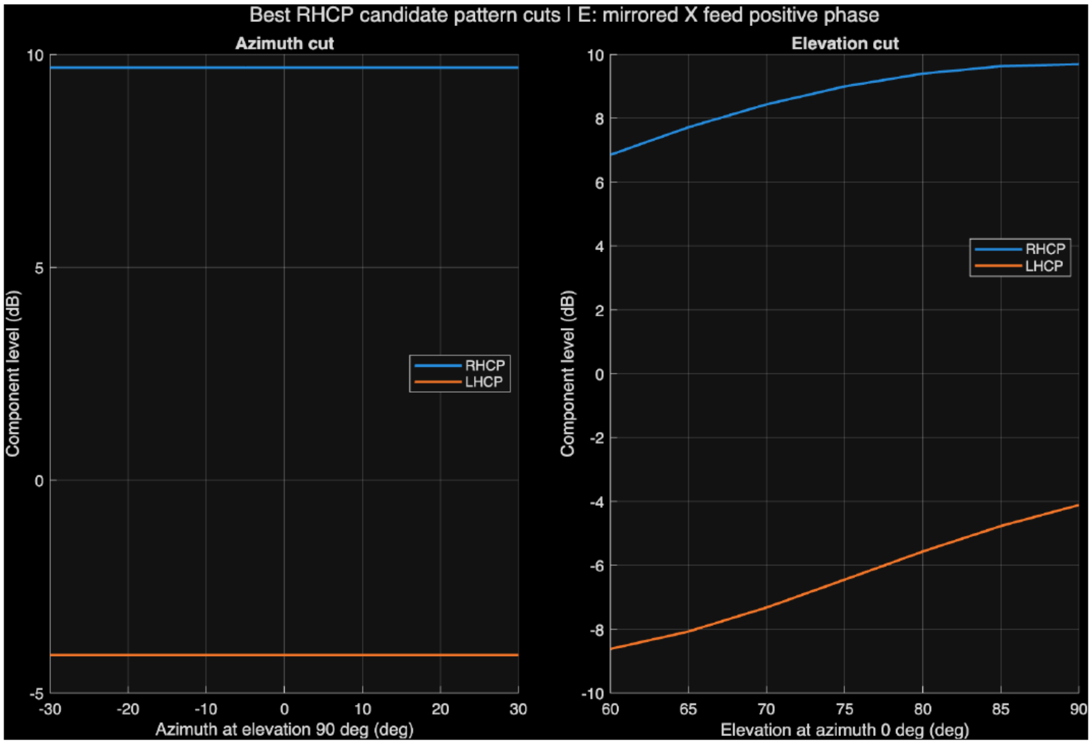

Figure 106. Center-Feed RHCP/LHCP Sense Check - Graph evidence: Best rhcp pattern cut.

In Figure 107, this plot tracks desired-versus-undesired CP content over frequency, which is more informative than a single handedness label. Read against the section objective, it supports the decision at the end of this block: Phase sign alone was not enough; mirroring one feed point changed sense more usefully.

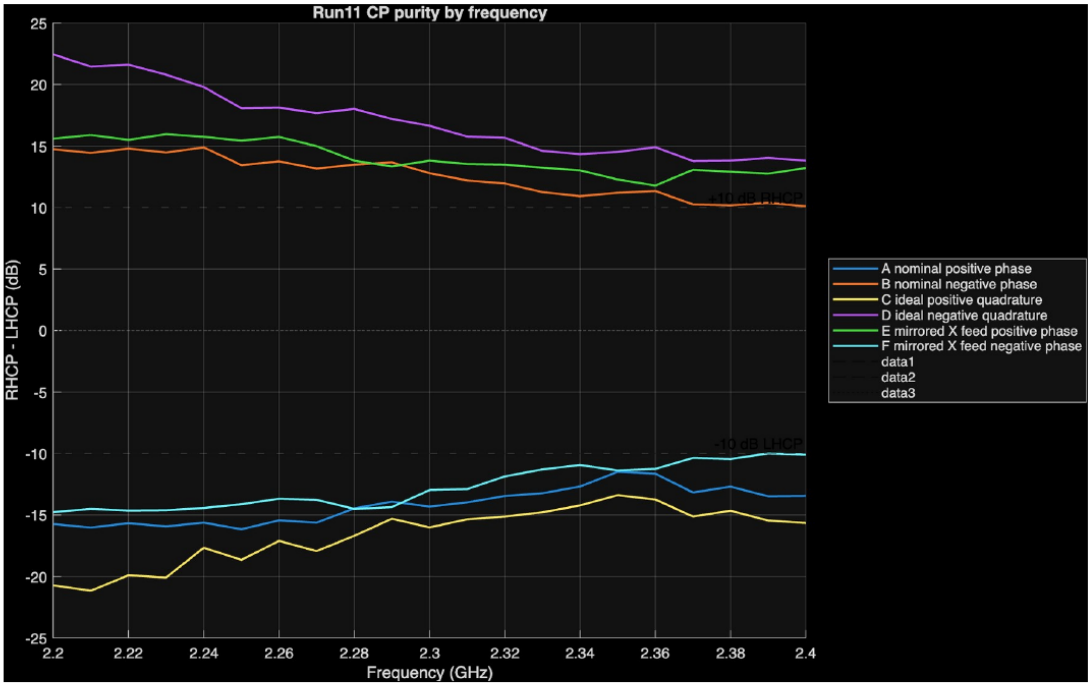

Figure 107. Center-Feed RHCP/LHCP Sense Check - Graph evidence: Cp purity by frequency.

In Figure 108, this boresight comparison is the handedness evidence: it shows which cases favor RHCP or LHCP in the feed-only model. It is included here because it affected the next step in the design path: Find a minimal feed-point set that supports both senses without fabricating a full four-feed implementation.

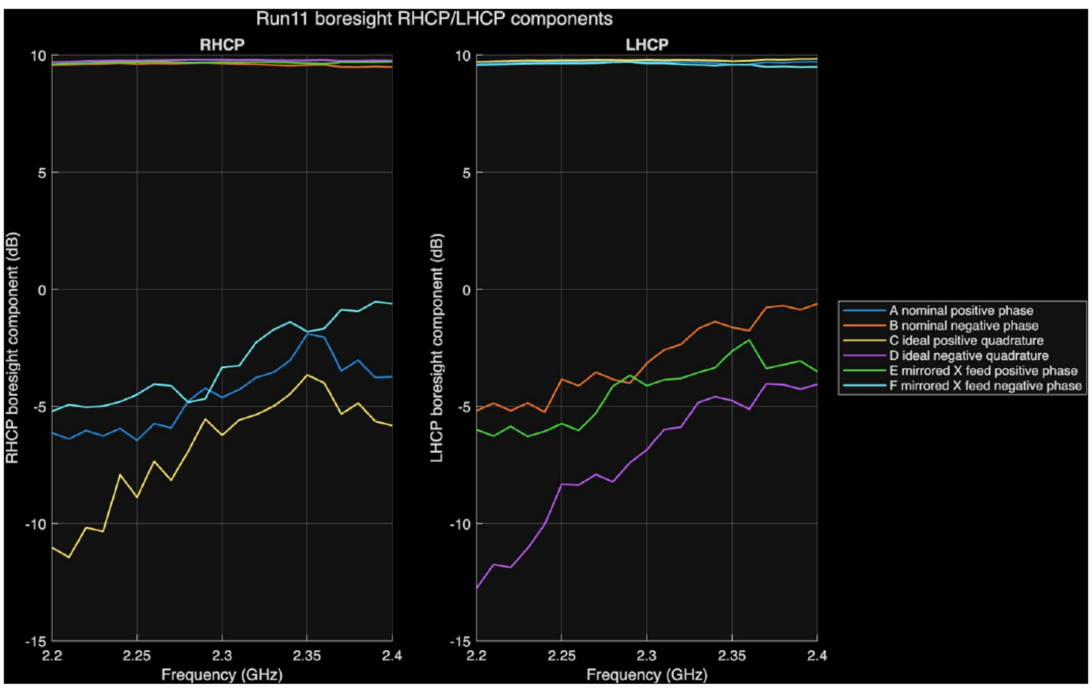

Figure 108. Center-Feed RHCP/LHCP Sense Check - Generated evidence: Rhcp lhcp boresight by case.

Design decision: Phase sign alone was not enough; mirroring one feed point changed sense more usefully. Next step: Find a minimal feed-point set that supports both senses without fabricating a full four-feed implementation.

22. Minimal Three-Feed Selectable CP

The design question was: What is the smallest physical feed-point set that supports both CP senses?

We tried it because four possible center-feed points are mechanically possible, but fewer feed points simplify switching, routing, and assembly.

Trace source: Run 11B / three_feed_selectable_cp_check

Scope

- Tested +X, -X, -Y, and optional +Y feed-point combinations.
- Compared active S11, axial ratio, and RHCP/LHCP dominance by feed pair.
- Selected a common-feed hardware strategy rather than a purely simulation-optimal case.

Evidence and Interpretation

In Figure 109, this feed-point drawing turns the CP-sense result into hardware: common feed, switched feed, and optional backup point. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: What is the smallest physical feed-point set that supports both CP senses?

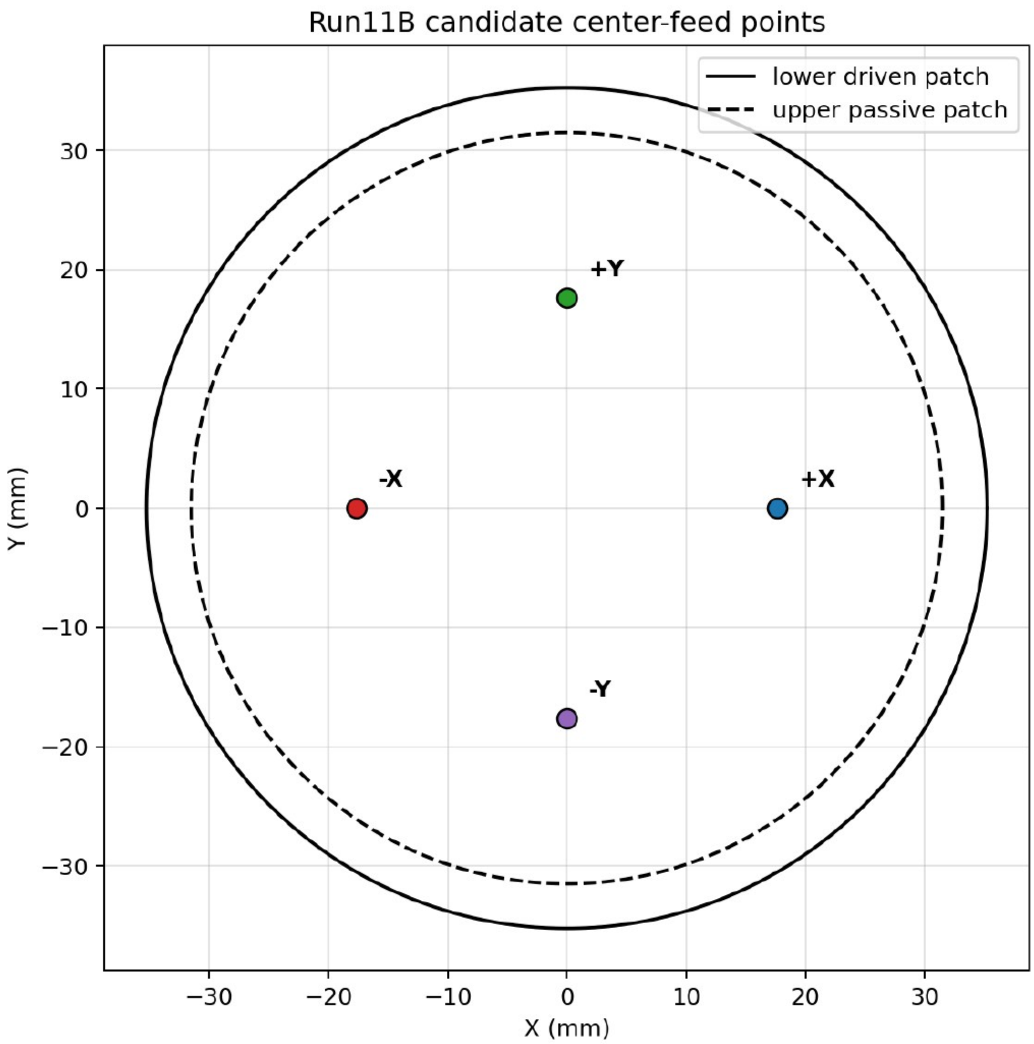

Figure 109. Minimal Three-Feed Selectable CP - Reference geometry: Geometry feed points.

In Figure 110, this case comparison checks that a CP switching option does not create an unacceptable input match penalty. Read against the section objective, it supports the decision at the end of this block: Fabricate +X, -Y, and -X. Use -Y common; switch +X/-X with amp 0.85 and +90 deg phase for opposite senses in the feed-only model.

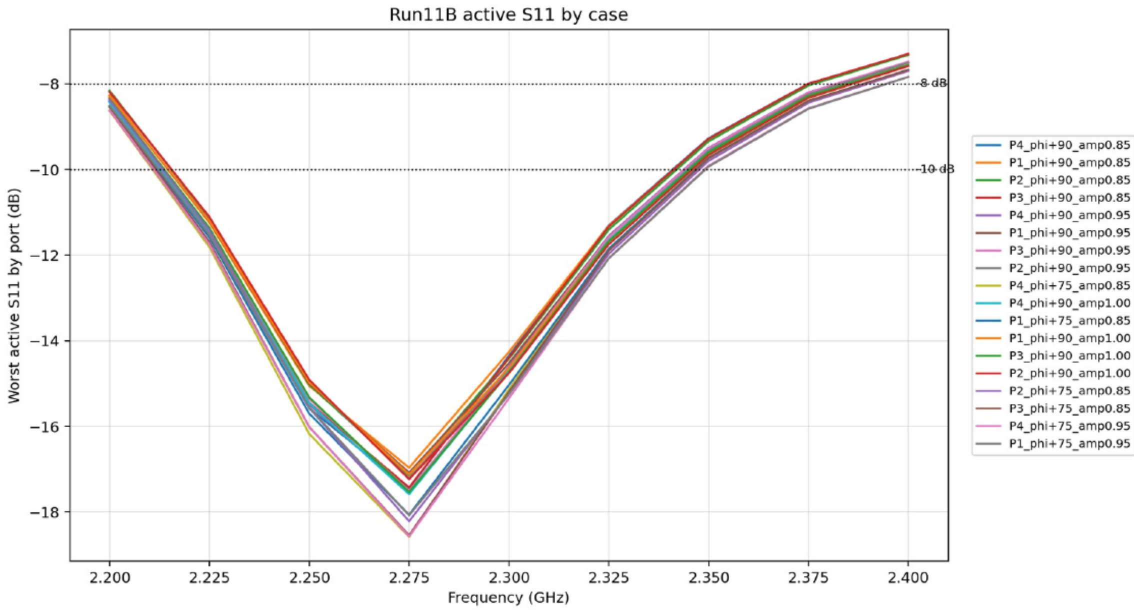

Figure 110. Minimal Three-Feed Selectable CP - Graph evidence: Active s11 by case.

In Figure 111, this case comparison checks that the switching option still preserves usable CP quality. Read against the section objective, it supports the decision at the end of this block: Fabricate +X, -Y, and -X. Use -Y common; switch +X/-X with amp 0.85 and +90 deg phase for opposite senses in the feed-only model.

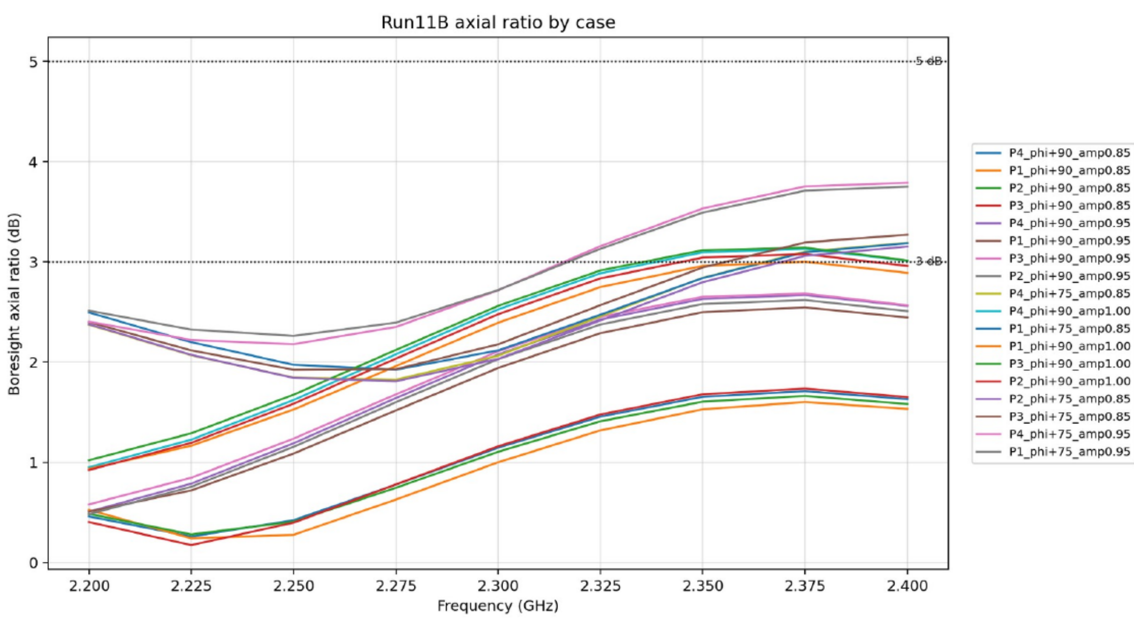

Figure 111. Minimal Three-Feed Selectable CP - Graph evidence: Axial ratio by case.

In Figure 112, this dominance plot is the practical CP-switching result that justified the three-feed hardware option. Read against the section objective, it supports the decision at the end of this block: Fabricate +X, -Y, and -X. Use -Y common; switch +X/-X with amp 0.85 and +90 deg phase for opposite senses in the feed-only model.

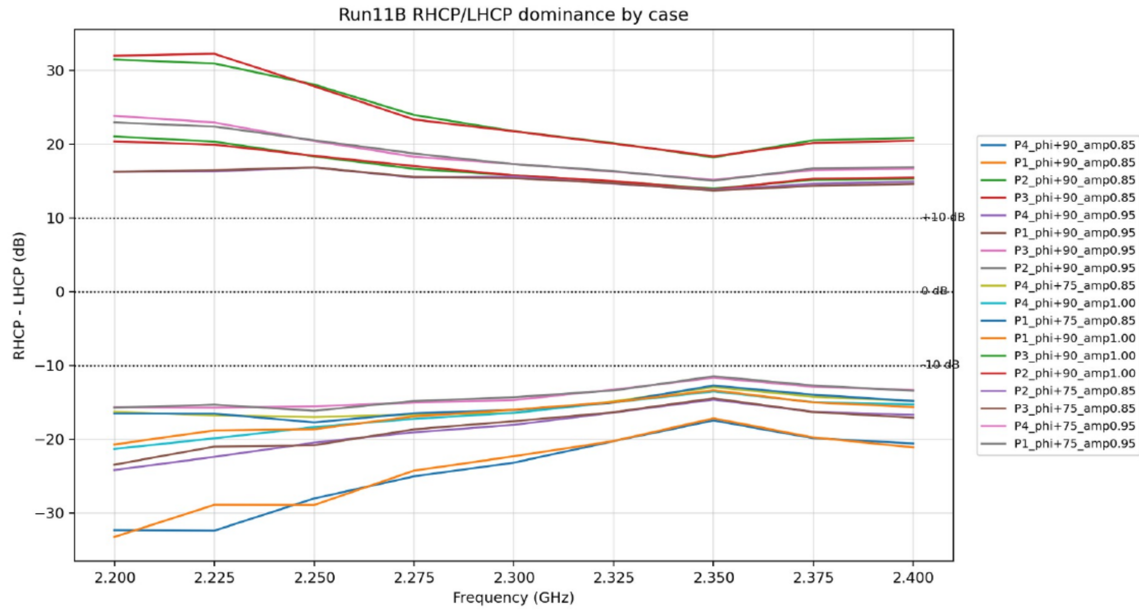

Figure 112. Minimal Three-Feed Selectable CP - Graph evidence: Rhcp lhcp dominance by case.

Design decision: Fabricate +X, -Y, and -X. Use -Y common; switch +X/-X with amp 0.85 and +90 deg phase for opposite senses in the feed-only model. Next step: Carry the three-feed selectable CP pattern into the final height and CAD package work.

23. Final Height Distribution and Manufacturing Check

The design question was: Can the selected passive stack be improved by moving the air gaps and tightening dimensions for fabrication?

We tried it because run7 selected the family; Run12 converted that family into final V5A dimensions and manufacturing constraints.

Trace source: Run 12 / height_distribution_sweep_center_feed

Scope

- Validated the Run7 baseline, swept height distributions, locally trimmed radius/excitation, and checked manufacturing variants.
- Compared baseline versus selected best S11 and axial ratio.
- Generated heatmaps for worst active S11, axial-ratio limit, and -10 dB coverage.

Evidence and Interpretation

In Figure 113, this is the geometry reference for the plots that follow. It fixes the physical interpretation of the S11, coupling, and axial-ratio curves on this page. That matters before the graphs because it fixes the physical layout behind the objective, so the plots answer the question: Can the selected passive stack be improved by moving the air gaps and tightening dimensions for fabrication?

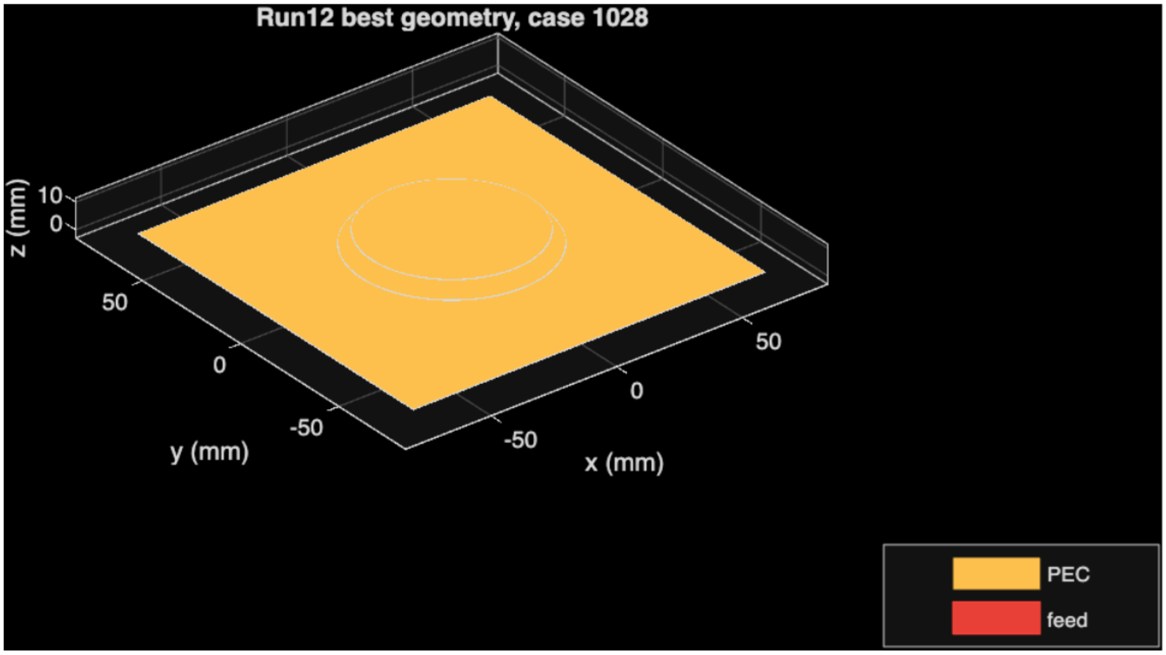

Figure 113. Final Height Distribution and Manufacturing Check - Reference geometry: Best geometry.

In Figure 114, this generated evidence is included for traceability because it affected the final height distribution and manufacturing check decision path. It is included here because it affected the next step in the design path: Transfer the final dimensions into V4/V5A CAD, RF windows, feed pins, and ISO-style drawing sheets.

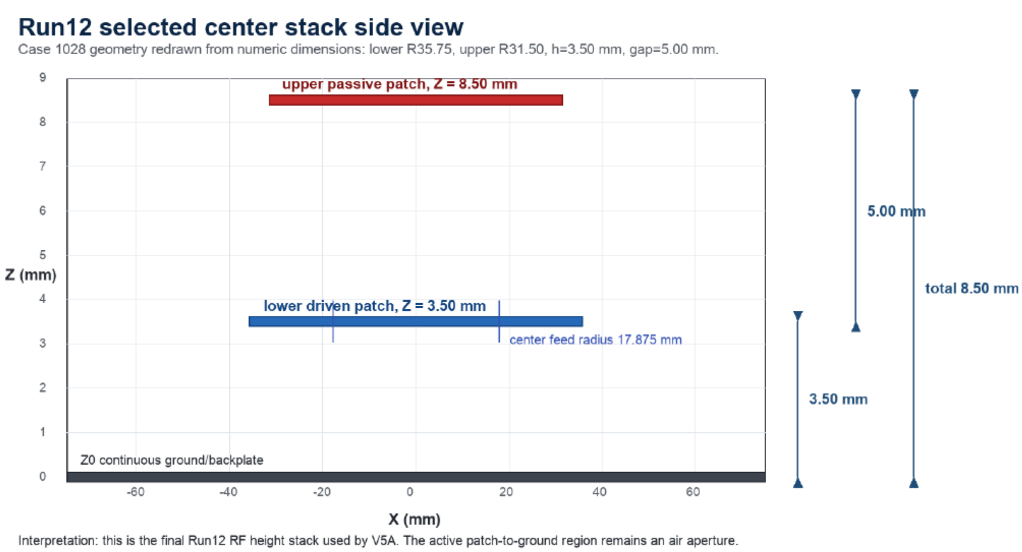

Figure 114. Final Height Distribution and Manufacturing Check - Generated evidence: Clean run12 stack side view.

In Figure 115, the axial-ratio plot is the CP quality check. It separates geometries that merely radiate from geometries that can plausibly support circular polarization. Read against the section objective, it supports the decision at the end of this block: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap.

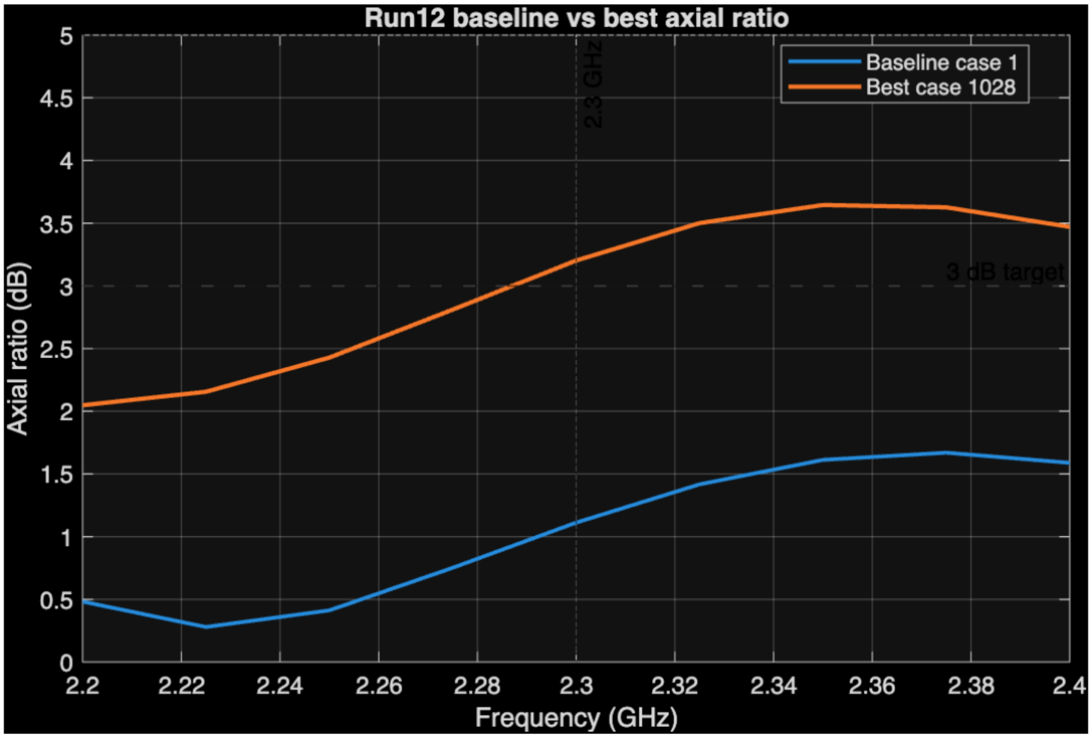

Figure 115. Final Height Distribution and Manufacturing Check - Graph evidence: Baseline vs best axial ratio.

In Figure 116, this plot is a decision plot: it shows whether the new candidate improves active match over the stored baseline across the working band. Read against the section objective, it supports the decision at the end of this block: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap.

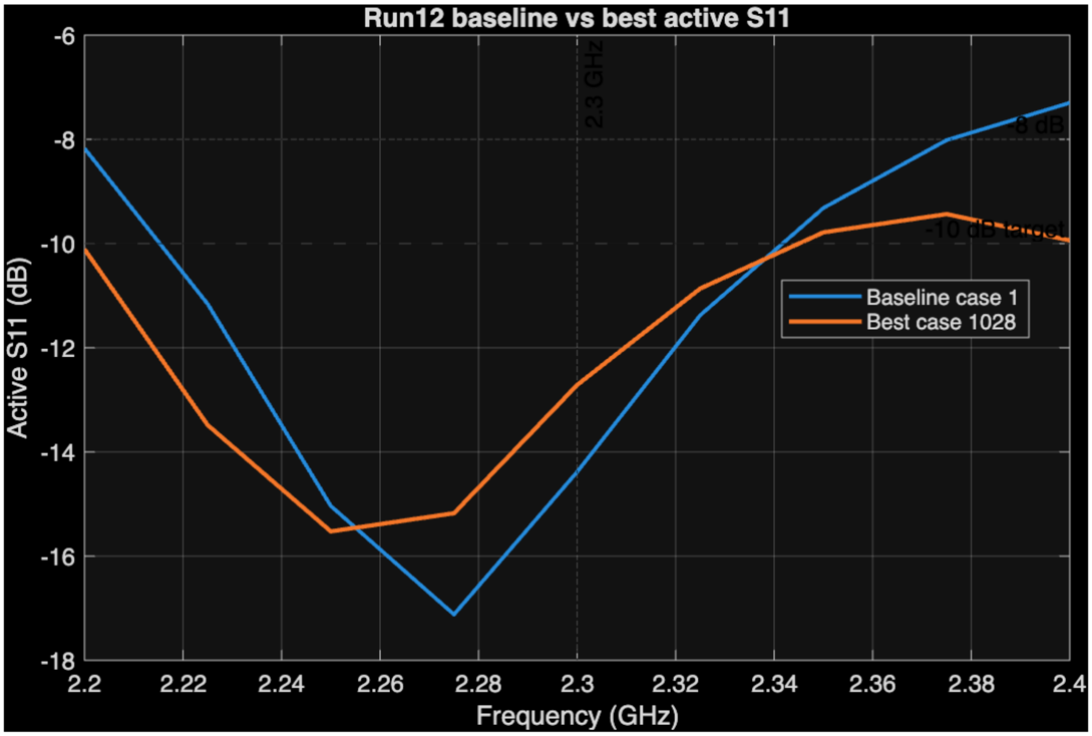

Figure 116. Final Height Distribution and Manufacturing Check - Graph evidence: Baseline vs best s11.

In Figure 117, this heatmap shows the -10 dB coverage tradeoff that pushed the final stack toward the selected air-gap distribution. Read against the section objective, it supports the decision at the end of this block: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap.

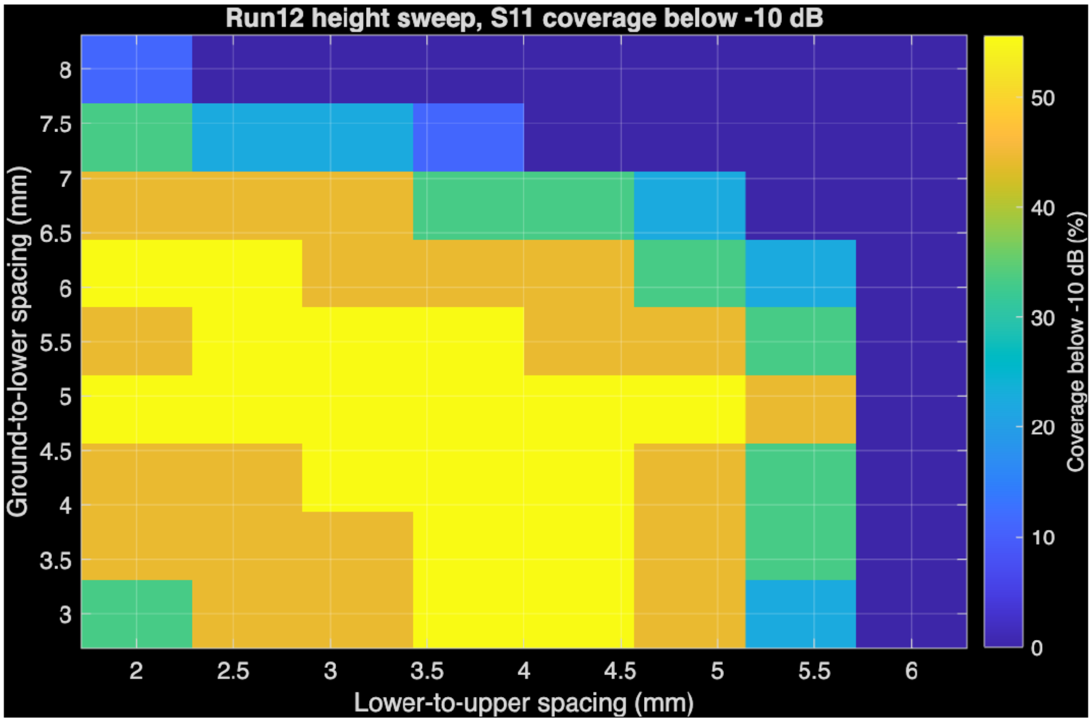

Figure 117. Final Height Distribution and Manufacturing Check - Graph evidence: Height heatmap coverage minus10.

In Figure 118, this heatmap shows how the same height choices affect axial-ratio risk, so the final height is not selected from S11 alone. Read against the section objective, it supports the decision at the end of this block: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap.

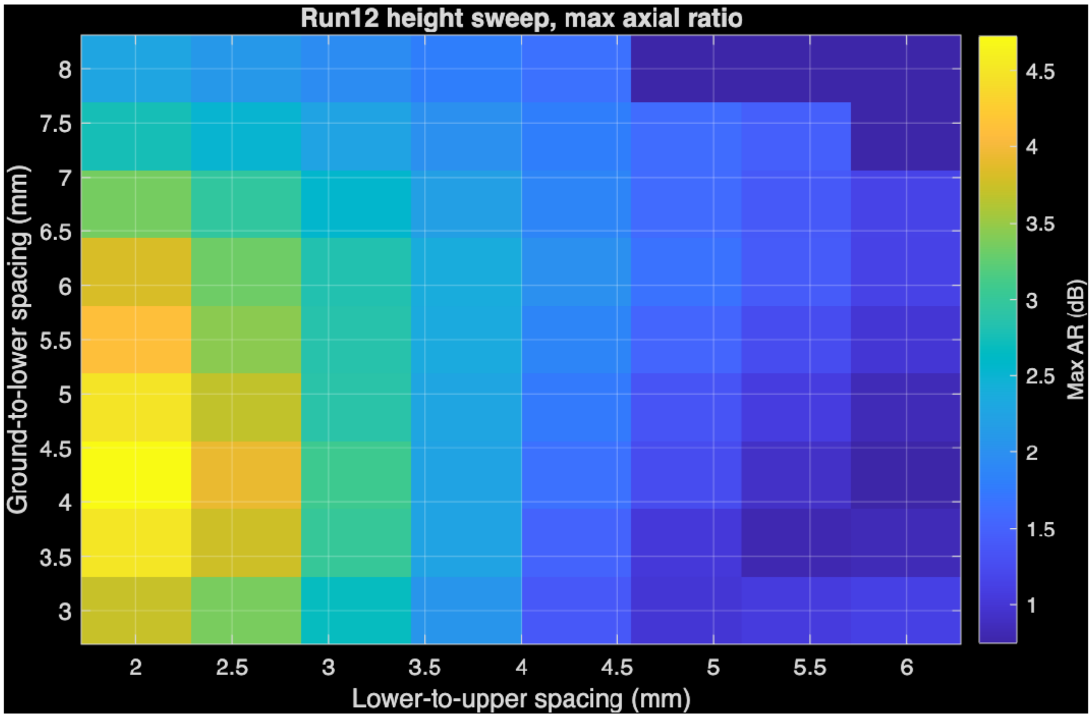

Figure 118. Final Height Distribution and Manufacturing Check - Graph evidence: Height heatmap max ar.

In Figure 119, this plot is why the simple patch could not be frozen: the center frequency can look good while the 2.2-2.4 GHz band edges remain weak. Read against the section objective, it supports the decision at the end of this block: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap.

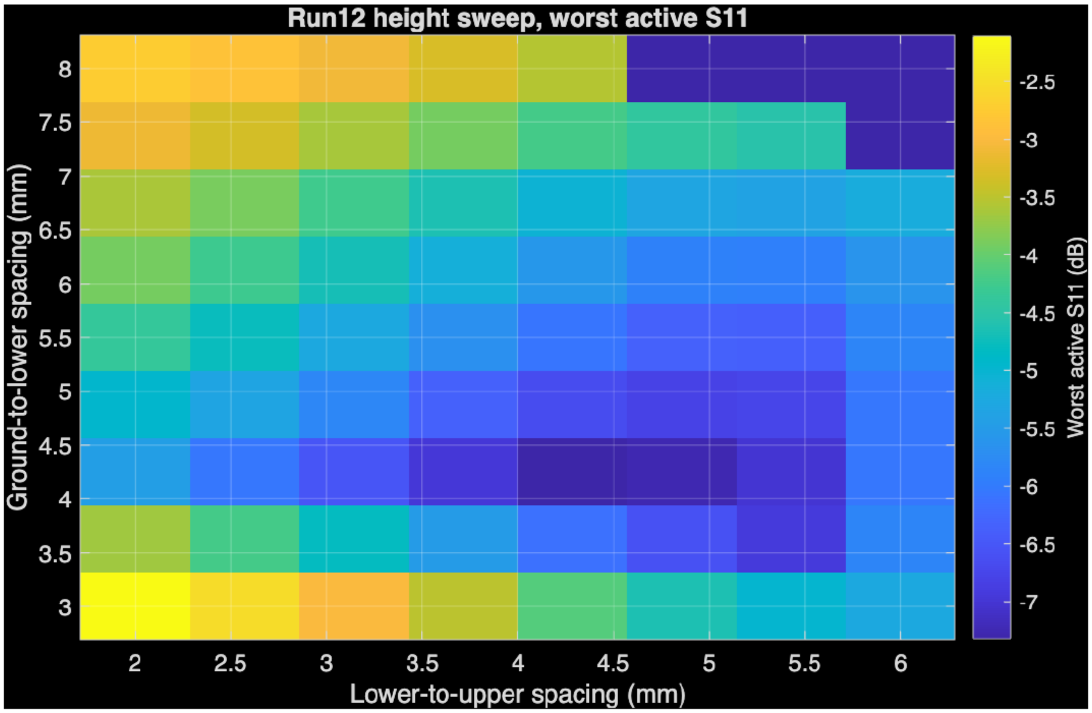

Figure 119. Final Height Distribution and Manufacturing Check - Graph evidence: Height heatmap worst s11.

Design decision: Selected final center: lower R35.75 mm, upper R31.50 mm, 3.50 mm ground-to-lower gap, 5.00 mm lower-to-upper gap. Next step: Transfer the final dimensions into V4/V5A CAD, RF windows, feed pins, and ISO-style drawing sheets.

Final Selection: What the Story Produced

Only after the patch, CP, sensor-ring, stack, slot, reflector, and selectable-CP checks does the report arrive at the build geometry. This is the endpoint of the evidence chain, not the starting assumption.

Final item	Selected value	Source
Center lower driven patch	R35.75 mm, copper Z=3.50 mm	Final height-distribution sweep -> V4/V5A
Center upper passive patch	R31.50 mm, copper Z=8.50 mm	Final height-distribution sweep -> V4/V5A
Center feed radius	17.875 mm, three-feed selectable CP plus optional backup point exported	Selectable-CP check and final stack
Ecartometer sensors	R34.50 mm at R130.0 mm, copper Z=3.50 mm	Selected ecartometer symmetry case, carried into architecture package
Carrier RF windows	Center R50 mm; sensor windows R45 mm	Manufacturing notes -> V5A
School-dish focus reference	F/D 0.45, focus 481.5 mm from vertex	School-dish reflector validation
Optimized reflector option	D1.50 m, F/D 0.45, F=675 mm, estimated gain 29.51 dBi	Reflector geometry optimization

From Coordinates to Blueprints

The blueprint packages convert the selected RF coordinates into buildable objects. V3 created a modular tile layout from the selected architecture. V4 moved the package to the final height and radius selection and added the windowed carrier/backplate, spacer collars, feed-through markers, cupola/radome keepout references, and tile-specific outputs. V5A is the latest integrated wideband feedhead package.

V5A contains the final passive stacked CP circular center feed plus four ecartometer sensor tiles. Major CAD objects are exported as separate bodies/components and as individual part files. The carrier and tile substrate frames are windowed so active patch-to-ground fields remain air gaps.

Blueprint asset	Purpose
assembly/V5A_integrated_feedhead_for_fusion.step	Primary Fusion import entry point with separate major components.
drawings/V5A_full_assembly_top_view.pdf/svg	Top-view mechanical/RF placement drawing.
drawings/V5A_side_stack_dimensions.pdf/svg	Z-height and spacer stack drawing.
drawings/V5A_feed_port_map.pdf/svg	Center and ecartometer feed-point reference.
drawings/V5A_ground_backplate_layout.pdf/svg	Backplate, mounting, and reference geometry.
drawings/V5A_rf_keepouts.pdf/svg	RF windows and no-intrusion areas.
kicad/*.kicad_pcb	Board templates for center lower, center upper, ground/backplate, and sensor tiles.
dimensions/*.csv	Machine-readable coordinates for feed points, windows, mounting holes, pins, and manifests.

Blueprint handling: The report embeds ISO-style technical drawing sheets generated from the V5A dimension CSVs: border, title block, first-angle projection note, line-weighted views, centerlines, extension/dimension lines, item balloons, and coordinate tables. The native STEP/DXF/KiCad/PDF exports remain the manufacturing authority.

ISO-style technical drawing sheet - V5A selected RF geometry

VIEW A - TOP ASSEMBLY, +Z TOWARD DISH/RADOME

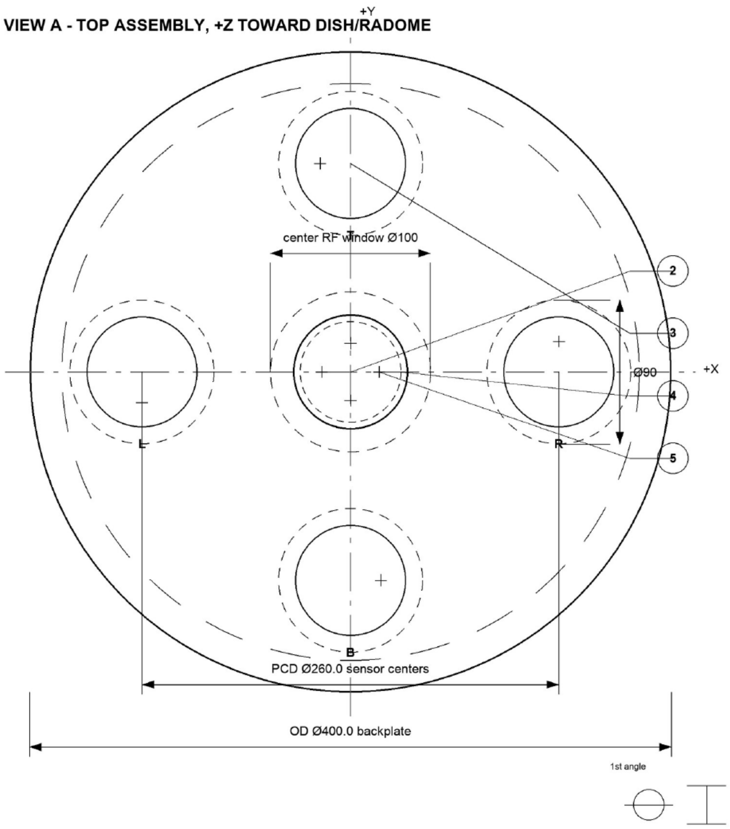

Item	Description	Primary size	Qty
1	Ground/backplate	Ø400.0	1
2	Center lower driven patch	Ø71.50	1
3	Ecartometer sensor patch	Ø69.00	4
4	Upper passive patch	Ø63.00	1
5	Feed probe points	see sheet 3	8
6	Carrier RF windows	Ø100/Ø90	5

Symbol	Meaning	Ref.
continuous thick	copper / physical component edge	ISO 128
dash-dot	centerline or datum axis	ISO 128
dashed	RF keepout / hidden reference	ISO 128
Ø	diameter dimension	ISO 129

NOTES

1. Do not move RF tile coordinates in Fusion.
2. Keep carrier material outside all dashed RF windows.
3. D0 is the feed-stack center; Z0 is the ground/backplate top face.
4. Drawing summarizes V5A dimension CSVs; CAD/DXF/STEP remain manufacturing authority.

V5A TOP ASSEMBLY

DWG No.	**V5A-ISO-001**	Rev	**A**
Scale	**1:3**	Units	**mm**
Sheet	**1/3**	Projection	**ISO 1st angle**

Figure 120. ISO-style sheet 1: V5A top assembly and primary dimensions.

ISO-style technical drawing sheet - V5A selected RF geometry

SECTION A-A - CENTER STACK AND SENSOR PLANE

Vertical scale enlarged for readability. Dimensions remain true numeric values in mm.

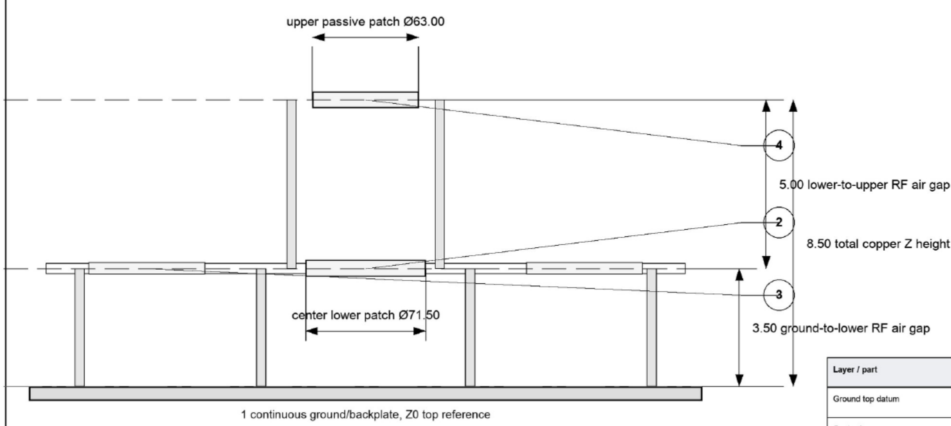

Layer / part	Z mm	Critical size	Manufacturing note
Ground top datum	0.00	Ø400.0	continuous conductive reference
Center lower copper	3.50	Ø71.50	driven feed tile
Sensor copper	3.50	Ø69.00	4x ecartometer tiles
Upper passive copper	8.50	Ø63.00	floating, centered
RF air window	0-3.50	Ø100/Ø90	no carrier material

BUILD RULE: active RF windows remain air; dielectric fill requires a new EM solve.

TOLERANCE TARGETS

Matched spacer height sets: target ±0.10 mm prototype.

1st angle

V5A STACK SECTION

DWG No.	**V5A-ISO-002**	Rev	**A**
Scale	**X 1:2 / Z 10:1**	Units	**mm**
Sheet	**2/3**	Projection	**ISO 1st angle**

Figure 121. ISO-style sheet 2: V5A stack section and Z-height dimensions.

ISO-style technical drawing sheet - V5A selected RF geometry

DETAIL C - CENTER SELECTABLE CP FEED POINTS

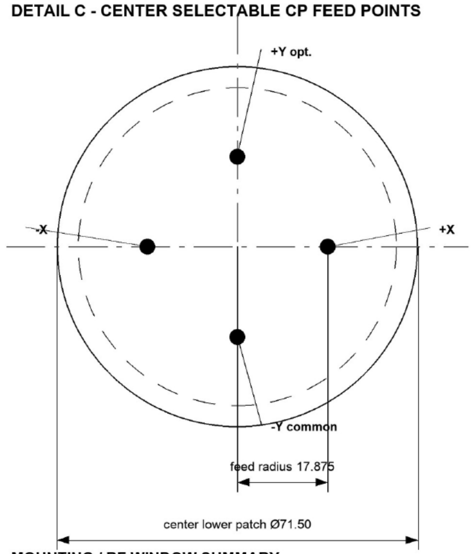

FEED COORDINATES, GLOBAL DATUM D0/Z0

Port	X	Y	Z	Function
C_X_POS	17.875	0	3.5	selectable cp switched
C_X_NEG	-17.875	0	3.5	selectable cp switched
C_Y_NEG	0	-17.875	3.5	selectable cp common
C_Y_POS	0	17.875	3.5	optional test backup
ECART_R	130	18.975	3.5	ecartometer tangential f
ECART_L	-130	-18.975	3.5	ecartometer tangential f
ECART_T	-18.975	130	3.5	ecartometer tangential f
ECART_B	18.975	-130	3.5	ecartometer tangential f

SOURCE: V5A dimension CSVs and integrated STEP.

DETAIL D - SENSOR FEED LOCAL ORIENTATION

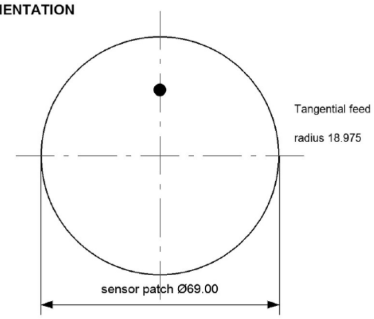

MOUNTING / RF WINDOW SUMMARY

Feature	Size / count	Note
RF windows	center Ø100; 4x Ø90	carrier remains open
Clamp holes	20x Ø3.2	outside active copper
Locating pins	10x Ø3.0	asymmetric pattern
Sensor centers	R130 from D0	L/R/T/B ring

1st angle

V5A FEED / MOUNTING DATA

DWG No.	**V5A-ISO-003**	Rev	**A**
Scale	**various**	Units	**mm**
Sheet	**3/3**	Projection	**ISO 1st angle**

Figure 122. ISO-style sheet 3: V5A feed coordinates and mounting data.

Build Sequence Carried by V5A

- Prepare the continuous ground plane/backplate.
- Install the windowed carrier and confirm open RF windows under all patches.
- Install 3.5 mm ground-to-lower spacers.
- Install the center lower driven tile and four sensor tiles.
- Use the feed-driver pin layout STEP/STP and feed_points_global.csv to place and check feed pins.
- Install 5.0 mm center lower-to-upper spacers.
- Install the center upper passive tile into the lower driven tile interface pattern.
- Run VNA testing before and after radome/cupola installation, then mount at dish focus.

Source Traceability and Caveats

The report follows the files currently present in the workspace. The critical engineering distinction is between electromagnetic simulation outputs, architecture compilations, reflector estimates, and CAD/fabrication translations.

Category	Source files used
Project overview	satcom-feed-sim/README.md
Run8 architecture	out/Run8/final_architecture/run8_architecture_report.md and associated PNG/CSV files
Run9 reflector validation	out/Run9/parabola_dimensioned_reflector_validation/ run9_parabola_dimensioned_reflector_validation_report.md
Run10 school and optimized reflector	out/Run10/Run10A/.../recommendations.txt and out/Run10/Run10B/.../recommended_parabola_geometry.txt
Run11/Run11B CP checks	out/Run11/.../metrics.txt and out/Run11B/.../metrics.txt
Run12 final center geometry	out/Run12/height_distribution_sweep_center_feed/metrics.txt and manufacturing_height_notes.txt
V5A package	out/FabPackage/V5_fusion_clean_components/ V5A_integrated_wideband_feedhead/README.md, FAB_REPORT.md, validation_status.txt, drawings, previews, and dimensions CSVs

Explicit Caveats

- Run8 combines the best stored center-feed source and the best stored ecartometer source. It is not a newly solved exact hybrid model.
- Run9 and Run10 reflector results are aperture-efficiency, blockage, focus, and focal-plane models. They should be validated by reflector-inclusive simulation or measurement before final external performance claims.
- Feed-only RHCP/LHCP sense can invert through a single metallic reflector. The final user-facing handedness labels should be confirmed with the reflector installed.
- V5A validates geometry and component separation; it does not replace VNA measurement, connector stack-up verification, dielectric re-solve, or radome/cupola retesting.