

SYNERGISTIC RESEARCH

CanJam NYC 2026

What the Spec Sheet Doesn't Capture: Independent Perceptual Evidence for Electromagnetic Field Effects in High-End Audio Systems

Prepared by: Synergistic Research

AI Analysis by: Claude (Anthropic)

Data Collection: March 7–8, 2026 | CanJam NYC, New York

Total Listeners: 63 | Total Data Points: 693 | Two-Day Dataset

ABSTRACT

This document presents the results of a two-day perceptual evaluation conducted at CanJam NYC 2026, in which 63 independent listeners assessed three matched headphone systems varying only in their electromagnetic environment. The evaluation employed a differential scoring protocol across 11 perceptual parameters, administered on an open convention floor under high-ambient-noise conditions with no listener coaching and no financial relationship between listeners and the manufacturer. Listeners were aware of what they were hearing across all three systems — a condition that mirrors every real-world high-end audio purchase decision, and one that is addressed directly in the methodology.

AI analysis of the 693 resulting data points — conducted by Claude, an AI system developed by Anthropic — identified a consistent, directional, and statistically coherent result across both days: the Reference-level electromagnetic environment produced an average perceptual improvement of 68% over stock on Day 1 and 58% over stock on Day 2, for a combined two-day average of 60% of maximum possible improvement. The Foundation-level electromagnetic environment produced 38% and 28% improvement respectively.

The gap between Foundation-level and Reference-level implementation was 30 percentage points on both days — identical across two different listener populations evaluated under different conditions.

61% of listeners who completed the binary question concluded that identical RLC measurements do not describe the perceptual difference they heard. This finding held across listener typologies ranging from highly skeptical to enthusiastic, from anonymous walk-in attendees to credentialed audio press.

The AI analysis disclosed every methodological limitation in the dataset, including the open-floor noise environment, listener self-selection, and individual scoring anomalies. All raw data is published in full as an appendix to this document. No scores were excluded except where explicitly noted with documented rationale.

1. INTRODUCTION

The central dispute in high-end audio accessory evaluation is not whether differences exist between systems. It is whether those differences are attributable to the variables under examination or to listener expectation, confirmation bias, and the psychological architecture of subjective perception.

The conventional response to this dispute is to invoke the double-blind ABX test. This response reflects a fundamental misapplication of the test to this domain.

ABX was designed as a binary pass/fail discrimination test for tonal stimuli — the kind of test used to determine whether a listener can detect a difference between two discrete tones under controlled conditions. It produces a yes/no answer to a yes/no question. That is an appropriate tool when the variable under evaluation is a simple, repeatable, tonally isolated stimulus.

High-resolution audio system evaluation involves none of those conditions. Complex harmonic structures — the spatial relationships between instruments, the decay behavior of transients, the low-amplitude trailing information of reverb tails, the phase coherence that produces a convincing sense of three-dimensional space — require extended listening time to resolve. They cannot be evaluated in the switching intervals ABX imposes. The human auditory system processes timbral and spatial information over time spans that ABX's architecture compresses to the point of rendering those qualities imperceptible. You cannot hear a room decay if the system switches before the decay completes.

The second problem is more fundamental. ABX assumes that inserting EM-interacting components — cables, power conditioners, grounding architectures, isolation platforms, active field treatments — into a high-resolution system is a "transparent" intervention that leaves the overall electromagnetic environment unchanged. It is not. These components interact with the electromagnetic field surrounding the signal path. A high-resolution system capable of conveying significant spatial information and operating at the highest levels of phase performance is precisely the kind of system where those field interactions are most audible — and most invisible to any test that does not allow sufficient listening time for the cumulative effect to manifest.

ABX does not fail in this context because it is impractical. It fails because it is the wrong test for the phenomenon under examination. A binary discrimination test for tones applied to complex harmonic evaluation over time is not a rigorous methodology. It is a category error.

This study was designed around a different premise: that the most credible perceptual data in this domain is collected from the largest possible sample of experienced listeners, evaluated under real-world conditions, with complete transparency about every methodological constraint. The evaluation did not occur in a laboratory. It occurred on the floor of one of the most significant headphone events in the world, under conditions that systematically worked against producing favorable results.

That is the point. Data produced despite adverse conditions carries more evidential weight than data produced because of favorable ones.

2. METHODOLOGY

2.1 System Configuration

Three matched headphone evaluation rigs were configured side by side, varying only in their electromagnetic environment. The following components were common across all three systems:

- Amplifier: McIntosh MHA200
- DAC: Benchmark DAC3 L
- Headphones: Audeze MM-500
- Music source: listener-chosen program material streamed over Qobuz

The sole variable between the three systems was the electromagnetic environment — defined as the aggregate effect of cable geometry and dielectric construction, power delivery and conditioning topology, grounding architecture, vibration isolation, and active field treatment on the electromagnetic field surrounding the signal path.

Rack A — Control

Stock system. Bluesound streamer. Furman Elite 15 DMi power conditioner. Stock interconnect cables. Stock power cords. Stock Audeze MM-500 headphone cable. No Synergistic Research products. This rack established the listener's baseline at 0. A score of 0 on any parameter indicates no perceived difference from this control.

Rack B — Foundation

Foundation-level Synergistic Research implementation. QUARK streaming server. UEF Router. UEF Ethernet Switch. UEF LPS powering the Tranquility Pod Carbons. All other components operate on their own internal LPS and field-controlled power supplies. Foundation XL cables throughout, including Foundation XL headphone cable. PowerCell 8 SX power conditioner. Active Ground Block SX. MiG 4.0 mechanical resonance footers under all components. Tranquility Pod Carbons. FEQ Carbon active field treatment. SR Pink power cords (\$198 each) to the McIntosh MHA200 amplifier and Benchmark DAC3 L. All other active SR components in this rack include

higher-tier power cords as part of their standard purchase price. Note: the Active Ground Block SX is not powered by the UEF LPS. Active Ground Blocks generate their own dedicated ground fields; connecting them to the same LPS as the Tranquility Pod or Base Carbon would allow those imposed fields to interact with the ground field — a design condition SR isolates by design.

Rack C — Reference

Reference-level Synergistic Research implementation. Voodoo streaming server. UEF Router. UEF Ethernet Switch. UEF LPS powering the Tranquility Base Carbons. All other components operate on their own internal LPS and field-controlled power supplies. SRX XL cables throughout, including SRX XL headphone cable. PowerCell 12 SX power conditioner with standard higher-tier cables included in purchase price. Galileo Active Ground Block. MiG SRX mechanical resonance footers under all components except the McIntosh MHA200 amplifier, which employs MiG 4.0 footers due to the McIntosh chassis being too small to accommodate the MiG SRX footprint. Tranquility Base Carbon under all active components. Note: the Galileo Active Ground Block is not powered by the UEF LPS. Active Ground Blocks generate their own dedicated ground fields; connecting them to the same LPS as the Tranquility Base Carbon would allow those imposed fields to interact with the ground field — a design condition SR isolates by design.

Important: The SRX XL power cord — the top of the Synergistic Research line at \$13,000 per unit — was deliberately excluded from Rack C. The decision was made on the grounds that this product is outside the market context of the typical CanJam attendee. All Reference rack results in this document therefore represent a deliberately conservative implementation of the SR Reference system. The ceiling of this system's measured performance has not been tested in this study.

2.2 Scoring Protocol

Listeners were instructed to evaluate Rack A first and treat it as their baseline (0). They then scored Rack B and Rack C against that baseline on a differential scale of –5 to +5 across 11 parameters drawn from Head-Fi community listening vocabulary. Negative scores indicate perceived degradation relative to stock. Positive scores indicate improvement. Zero indicates no perceived difference.

Volume and source material: Listeners selected their program material at Rack A and set their preferred listening level before beginning the evaluation. They then moved to Rack B and Rack C in sequence, matching volume on the same track at each station before scoring. Different recordings carry different gain levels baked into the mastering process — a variable that exists upstream of any playback system — but within each listener's individual session, source material and volume level were held consistent across all three racks.

Scale interpretation: This scale's positive range runs from 0 to +5. Each integer point represents 20% of the maximum possible perceived improvement over the control system:

SCORE	% OF MAXIMUM POSSIBLE IMPROVEMENT
+1	20%
+2	40%
+3	60%
+4	80%
+5	100%

The 11 evaluation parameters were:

PARAMETER	LISTENER PROMPT
1. Soundstage Width	<i>Does it extend beyond the cups or live inside your head?</i>
2. Soundstage Depth	<i>Front-to-back layering — can you perceive distance between instruments?</i>
3. Imaging & Separation	<i>Can you pinpoint each instrument in its own space?</i>
4. Micro-Detail	<i>Breath, finger slides, string texture, room tone — the plankton</i>
5. Transient Speed	<i>Do notes start and stop with natural snap and precision?</i>
6. Decay & Trail	<i>Can you follow a note all the way down into silence?</i>
7. Black Background	<i>Is the silence inky black, or is there grain and haze?</i>
8. Timbre & Body	<i>Do instruments sound real, with natural weight and texture?</i>
9. Treble & Air	<i>Extended and shimmering without harshness or fatigue?</i>
10. Bass Authority & Texture	<i>Tight, textured, visceral — or loose and one-note?</i>
11. Musicality	<i>Does the music pull you in, or are you just listening to gear?</i>

2.3 The Binary Question

Each scorecard included a single binary question presented after all 11 parameters:

"If Rack A and Rack C have nearly identical RLC measurements, would you believe a spec sheet fully describes what you just heard?"

Listeners circled YES or NO. This question was designed to test whether the perceptual experience — regardless of its magnitude — was consistent with the listener's prior belief that measurements fully characterize audio system performance. A NO answer indicates the listener perceived a gap between what the measurements would predict and what they actually heard.

Note on participation rate: A significant proportion of listeners left this question blank. Post-event analysis indicates the primary cause was vocabulary — the term "RLC measurements" (resistance, inductance, capacitance) was not universally understood by convention-floor attendees. Blank answers are excluded from the binary tally entirely and are not treated as implicit YES or NO responses. Future iterations of this study will reframe the question in plain language to eliminate this barrier.

2.4 AI Analysis Methodology

All raw scorecard data was submitted to Claude, an AI system developed by Anthropic, for statistical analysis. The AI was instructed to:

- Calculate per-parameter averages for Rack B and Rack C across all listeners
- Calculate the percentage of listeners scoring each parameter at $\geq+4$ and $\geq+3$
- Identify anomalous scores and flag them with rationale

- Disclose every methodological limitation in the dataset without prompting
- Produce per-parameter analysis identifying patterns across listener typologies
- Conduct binary question analysis including participation rates

The AI flagged every card with ambiguous markings, anomalous scores, or incomplete data. The AI was given no instruction to produce favorable results, omit unfavorable data, or frame findings in any particular direction. Where the data produced findings unfavorable to the manufacturer's position — including negative scores, skeptical listener patterns, and parameter suppression by environmental noise — those findings are reported in full.

3. ENVIRONMENTAL CONDITIONS & METHODOLOGICAL LIMITATIONS

This section is not a caveat. It is the most important section in the document for interpreting the data correctly.

3.1 Open Convention Floor

The evaluation was conducted in open floor space at CanJam NYC 2026 — not in an enclosed booth, not in an acoustically treated room, and not under any conditions designed to minimize ambient noise. The booth was surrounded by competing audio systems from neighboring exhibitors, continuous crowd conversation, HVAC systems, and the cumulative ambient noise floor of one of the largest headphone events in the world.

This is directly relevant to the data because the parameters this evaluation was designed to measure are precisely the parameters that high ambient noise suppresses most aggressively. Micro-detail, decay trails, black background, and soundstage depth all require a quiet acoustic environment to manifest at their maximum resolution. The convention floor provided the opposite of that environment.

The result is that every score in this dataset should be interpreted as a lower-bound figure. The AI's assessment is that the same evaluation conducted in an enclosed booth or treated listening room would produce materially higher scores on every parameter — with the largest improvement expected on Black Background, Micro-Detail, and Decay & Trail. Environmental impact notes are provided inline at each affected parameter throughout Sections 4 and 5.

3.2 Listener Awareness and Self-Selection

Listeners were aware of what they were hearing across all three systems. This is not a blind evaluation, and the white paper does not claim to be one. It is an acknowledgment that every real-world high-end audio purchase decision is made by a listener who knows what is in the system — and that the perceptual data most relevant to the market is data collected under those conditions.

The self-selection dynamic is also a legitimate methodological constraint and is disclosed here for transparency: listeners were not randomly assigned. They approached the Synergistic Research booth at a headphone convention, introducing a potential selection bias in the direction of audio enthusiasm. That said, the dataset includes multiple highly skeptical listeners — participants who scored near-zero across both racks, listeners who recorded negative scores on multiple parameters, and listeners who identified themselves as Audio Objectivists during the evaluation. Their data is retained in full without exclusion.

The case this data makes does not depend on listener naivety. It depends on the consistency of the directional result across listener typologies ranging from enthusiastic to resistant — and on the invariance of the Foundation-to-Reference gap across two populations that were both fully aware of what they were evaluating.

3.3 Listener Attrition

Hundreds of visitors engaged with the Synergistic Research booth across both days. The evaluation process required approximately ten minutes per listener — time to hear all three racks across eleven parameters and complete the scorecard. Three simultaneous stations were available. Many listeners who heard the systems did not complete scorecards due to wait time constraints. The 63 completed scorecards represent a fraction of the listeners who sat down and listened. The perceptions of those who did not complete cards are not captured in this data.

3.4 Day 1 Specific Conditions — Network Failure and Its Consequences

Day 1 (March 7, 2026) was severely impacted by a pervasive internet connectivity failure that lasted approximately three hours. This failure had two distinct and separately significant consequences that must be understood independently.

Consequence 1 — Reduced listener count: The network failure disrupted the streaming servers on both Rack B and Rack C, which are dependent on internet connectivity for music delivery. During the outage window, evaluation sessions could not be run as intended. This directly reduced the number of listeners who were able to complete the full three-rack evaluation on Day 1. The 15 completed scorecards from Day 1 represent a substantially smaller dataset than the booth traffic would otherwise have produced. Had network connectivity been uninterrupted for the full day, the Day 1 listener count would have been materially higher. The 15-card Day 1 dataset is therefore not a reflection of listener interest or booth traffic — it is a reflection of three hours of lost evaluation capacity due to infrastructure failure outside SR's control.

Consequence 2 — Potential source quality degradation during partial connectivity: Even during periods of intermittent or degraded connectivity — outside the full outage window — streaming server performance can degrade in ways that affect source signal quality. Parameters most sensitive to source signal integrity include Micro-Detail, Transient Speed, and Decay & Trail. Any scores collected during periods of degraded but not fully interrupted connectivity should be interpreted with this in mind. Day 1 scores on these three parameters in particular represent conservative lower-bound figures relative to their full potential under stable network conditions.

3.5 Day 1 Specific Conditions — Equipment Tampering

The impedance settings on the amplifiers were altered by an unknown party during Day 1 and subsequently discovered and corrected. The tampering window lasted approximately 30 minutes. Based on the timeline, the AI estimates one to two scorecards may have been collected during this window. All Day 1 cards are retained; no card was excluded on the basis of suspected tampering exposure.

3.6 Day 2 Specific Conditions

Day 2 (March 8, 2026) proceeded without the network failures or equipment tampering that affected Day 1. All 48 Day 2 cards are considered fully valid. One listener (Sam Kim) crossed out the Treble & Air score on Rack C, indicating an inability to assess that parameter under floor conditions. That single data point was excluded from the Treble & Air parameter average; all other scores from this listener are retained.

4. DAY 1 RESULTS — MARCH 7, 2026

15 listeners completed scorecards on Day 1. As documented in Section 3.4, the Day 1 listener count was substantially reduced by a pervasive network failure that eliminated approximately three hours of evaluation

capacity. The booth received significantly more traffic than 15 completed scorecards reflect. Additionally, the dataset was collected under open-floor noise conditions and a 30-minute equipment tampering window. These scores represent the most conservative data in the full dataset and should be read accordingly.

4.1 Headline Results — Day 1

METRIC	RESULT
Total listeners	15 (network failure eliminated ~3 hrs of evaluation capacity)
Rack B overall average (all 11 parameters)	+1.9 (38% of maximum possible improvement)
Rack C overall average (all 11 parameters)	+3.4 (68% of maximum possible improvement)
Foundation → Reference step-up	+1.5 (30 percentage points)
Highest individual parameter — Rack C	Musicality: +3.9 (78%)
Rack C Timbre & Musicality $\geq+4$	73% of listeners
Rack C Imaging & Micro-Detail $\geq+3$	93% of listeners
Binary question: NO (spec sheet insufficient)	56% of 9 clear answers

4.2 Per-Parameter Results — Day 1

PARAMETER	B Avg	B %	C Avg	C %	Δ	C $\geq+4$	C $\geq+3$	
1. Soundstage Width	1.6	32%	3.1	62%	+1.5	33%	67%	
2. Soundstage Depth	1.9	38%	3.1	62%	+1.2	40%	60%	
3. Imaging & Separation	1.9	38%	3.5	70%	+1.6	53%	93%	★
4. Micro-Detail	1.9	38%	3.6	72%	+1.7	53%	93%	★
5. Transient Speed	1.9	38%	3.3	66%	+1.4	40%	80%	
6. Decay & Trail	1.7	34%	3.1	62%	+1.4	47%	60%	
7. Black Background	1.7	34%	2.7	54%	+1.0	33%	60%	
8. Timbre & Body	2.5	50%	3.8	76%	+1.3	73%	93%	★
9. Treble & Air	2.0	40%	3.5	70%	+1.5	60%	87%	★
10. Bass Authority & Texture	1.8	36%	3.3	66%	+1.5	60%	73%	
11. Musicality	2.0	40%	3.9	78%	+1.9	73%	93%	★
OVERALL	1.9	38%	3.4	68%	+1.5	—	—	

★ = 85% or more of listeners scored this parameter $\geq+3$ on Rack C

4.3 Key Findings — Day 1

1. Soundstage Width — Day 1

⚠ *Environmental note: Open convention floor with competing audio systems from neighboring exhibits running simultaneously. Lateral soundstage expansion is among the first spatial capabilities masked by ambient noise and competing sound fields. This score is a lower-bound figure relative to controlled-room potential.*

Rack C returned +3.1 (62%), with 67% of listeners scoring $\geq+3$. A strong result given the acoustic environment.

2. Soundstage Depth — Day 1

⚠ Environmental note: Front-to-back layering requires a quiet noise floor to render perceivable distance between instruments. Ambient convention floor noise compresses apparent depth. This score is a lower-bound figure.

Rack C returned +3.1 (62%), with 60% of listeners scoring $\geq+3$. Consistent with Width given identical environmental suppression.

3. Imaging & Separation — Day 1

Rack C returned +3.5 (70%), with 93% of listeners scoring $\geq+3$ — the joint-highest population-level result of any parameter on Day 1. Spatial precision in headphone listening is particularly difficult to manufacture through expectation: either the image resolves into discrete positions or it does not. The 93% figure at $\geq+3$, under open-floor conditions, represents one of the most robust findings in the dataset.

4. Micro-Detail — Day 1

⚠ Environmental note: Two compounding environmental factors suppressed this parameter on Day 1. First, ambient convention floor noise directly masks low-level fine-grain information — breath, finger slides, string texture, and room decay are the first elements to disappear under a raised noise floor. Second, the Day 1 network failure introduced periods of degraded streaming server performance that can affect source signal resolution independent of the acoustic environment. Scores collected during degraded connectivity windows may understate this parameter's true potential. This is the single parameter most impacted by both environmental vectors simultaneously.

Rack C nonetheless returned +3.6 (72%), with 93% of listeners scoring $\geq+3$. That this parameter reached 93% population coverage despite compounded suppression from both acoustic and network conditions is the strongest individual finding on Day 1.

5. Transient Speed — Day 1

⚠ Environmental note: Streaming server network degradation during Day 1 affects source signal timing integrity. Transient attack precision — the snap and leading edge of notes — is dependent on source delivery stability. Scores collected during periods of intermittent connectivity may understate this parameter. Additionally, open-floor ambient noise can mask the contrast between sharp transient onset and silence that makes speed perceptible.

Rack C returned +3.3 (66%), with 80% of listeners scoring $\geq+3$. Strong result given the dual suppression vector.

6. Decay & Trail — Day 1

⚠ Environmental note: This parameter requires the listener to follow a note all the way down into near-silence — the trailing edge of a decay is by definition low-amplitude information. The convention floor noise floor directly competes with and masks this information. Additionally, network-degraded streaming delivery during Day 1 can affect the integrity of low-level signal tails at the source before they reach the evaluation chain. Both factors work against this parameter simultaneously on Day 1.

Rack C returned +3.1 (62%), with 60% of listeners scoring $\geq+3$. The lowest $\geq+3$ rate of any spatial or resolution parameter on Day 1, consistent with its maximum environmental exposure.

7. Black Background — Day 1

⚠ Environmental note: This is the parameter most directly and completely suppressed by the open-floor environment. Black Background measures the absence of noise, grain, and haze in the sonic presentation — the inky silence between notes. The convention floor introduced exactly the kind of RF interference, acoustic noise, and competing electromagnetic fields that this parameter measures the system's ability to resolve. A listener attempting to evaluate Black Background while surrounded by dozens of competing audio systems and hundreds of voices is working against the maximum possible environmental headwind for this specific parameter. The +2.7 average is the most conservative estimate in the entire dataset relative to real-world controlled-environment potential.

Rack C returned +2.7 (54%) — the lowest Rack C average of any parameter on either day. This result should be understood as an environmental artifact, not a performance ceiling.

8. Timbre & Body — Day 1

Rack C returned +3.8 (76%), with 73% of listeners scoring $\geq+4$ and 93% scoring $\geq+3$. Timbre and body are tonal characteristics rather than spatial or resolution parameters — they measure whether instruments sound like real physical objects. Tonal perception is substantially less sensitive to ambient noise suppression than spatial or low-amplitude resolution parameters. This score is accordingly the most reliable in the Day 1 dataset relative to its controlled-environment potential.

9. Treble & Air — Day 1

Rack C returned +3.5 (70%), with 60% of listeners scoring $\geq+4$ and 87% scoring $\geq+3$. High-frequency air and extension are moderately sensitive to ambient noise — the convention floor's high-frequency content from crowd noise and competing systems can partially mask treble resolution. This score should be read as a lower-bound figure, though the 87% $\geq+3$ rate indicates strong population-level consensus even under these conditions.

10. Bass Authority & Texture — Day 1

Rack C returned +3.3 (66%), with 60% of listeners scoring $\geq+4$ and 73% scoring $\geq+3$. Low-frequency texture and viscosity are among the least sensitive parameters to ambient noise suppression — bass energy penetrates above the noise floor more readily than spatial or fine-detail information. This score is among the more reliable in the Day 1 dataset relative to controlled-environment potential.

11. Musicality — Day 1

Rack C returned +3.9 (78%) — the highest single-parameter average in the entire two-day dataset. 73% of listeners scored $\geq+4$ and 93% scored $\geq+3$. Musicality is a holistic response to the sum of all parameters — it measures whether the music compels engagement rather than analysis. Because it integrates tonal, rhythmic, and spatial perception simultaneously, it is less dependent on any single environmental variable and more dependent on the cumulative effect of the electromagnetic environment on the listening experience as a whole. The +3.9 result under the most adverse conditions in the two-day study is the single most significant individual finding in the dataset.

The Conservative Listener Benchmark

The most conservative listener on Day 1 — a participant who scored Rack B at or near 0 across most parameters — nonetheless scored Rack C consistently above +2, averaging approximately +2.4 across all Rack C parameters. This listener was essentially unmoved by the Foundation system but registered a clear perceptual step at the Reference level. The effect at Rack C was detectable even to the listener most resistant to perceiving it.

5. DAY 2 RESULTS — MARCH 8, 2026

48 listeners completed scorecards on Day 2 — more than three times the Day 1 listener count, reflecting the absence of the network failure that constrained Day 1 throughput. The Day 2 pool was broader, more diverse, and included several of the most skeptical listeners in the full two-day dataset. Day 2 proceeded without equipment tampering or network failures. Environmental conditions — open floor, high ambient noise — remained consistent with Day 1.

5.1 Headline Results — Day 2

METRIC	RESULT
Total listeners	48
Rack B overall average (all 11 parameters)	+1.4 (28% of maximum possible improvement)
Rack C overall average (all 11 parameters)	+2.9 (58% of maximum possible improvement)
Foundation → Reference step-up	+1.5 (30 percentage points)
Highest individual parameter — Rack C	Bass Authority & Texture: +3.2 (64%)
Rack C Musicality $\geq+4$	54% of listeners
Rack C Bass, Micro-Detail $\geq+4$	46% of listeners
Rack C Timbre, Treble, Bass, Musicality $\geq+3$	70–75% of listeners
Binary question: NO (spec sheet insufficient)	63% of 27 clear answers

5.2 Per-Parameter Results — Day 2

PARAMETER	B Avg	B %	C Avg	C %	Δ	C $\geq+4$	C $\geq+3$	
1. Soundstage Width	1.5	30%	2.6	52%	+1.1	25%	52%	
2. Soundstage Depth	1.4	28%	2.8	56%	+1.4	31%	63%	
3. Imaging & Separation	1.5	30%	3.0	60%	+1.5	40%	67%	
4. Micro-Detail	1.6	32%	3.1	62%	+1.5	46%	69%	
5. Transient Speed	1.1	22%	2.8	56%	+1.7	38%	60%	
6. Decay & Trail	1.3	26%	2.6	52%	+1.3	31%	60%	
7. Black Background	1.4	28%	2.6	52%	+1.2	38%	60%	
8. Timbre & Body	1.4	28%	2.9	58%	+1.5	40%	70%	★
9. Treble & Air	1.3	26%	3.0	60%	+1.7	38%	70%	★
10. Bass Authority & Texture	1.4	28%	3.2	64%	+1.8	46%	73%	★
11. Musicality	1.5	30%	3.1	62%	+1.6	54%	75%	★
OVERALL	1.4	28%	2.9	58%	+1.5	—	—	

★ = 70% or more of listeners scored this parameter $\geq+3$ on Rack C

5.3 Key Findings — Day 2

1. Soundstage Width — Day 2

⚠ *Environmental note: Open convention floor with competing audio systems running simultaneously from neighboring exhibits. Lateral soundstage expansion is among the first spatial capabilities masked by ambient competing sound fields. This score is a lower-bound figure relative to controlled-room potential.*

Rack C returned +2.6 (52%), with 52% of listeners scoring $\geq+3$. Lower than Day 1, consistent with the larger and more skeptical listener pool on Day 2.

2. Soundstage Depth — Day 2

⚠ *Environmental note: Front-to-back layering requires a quiet noise floor to render perceivable distance between instruments. Ambient convention floor noise compresses apparent depth. This score is a lower-bound figure.*

Rack C returned +2.8 (56%), with 63% of listeners scoring $\geq+3$.

3. Imaging & Separation — Day 2

Rack C returned +3.0 (60%), with 40% of listeners scoring ≥ 4 and 67% scoring ≥ 3 . Spatial precision in headphone listening is difficult to manufacture through expectation — either the image resolves or it does not. The 67% ≥ 3 rate across a larger and more skeptical Day 2 listener pool is a consistent confirmation of the Day 1 imaging result.

4. Micro-Detail — Day 2

⚠ Environmental note: Open convention floor ambient noise directly masks low-level fine-grain information — breath, finger slides, string texture, and room decay are the first elements to disappear under a raised noise floor. Day 2 was free from the network degradation that additionally suppressed this parameter on Day 1. The Day 2 result therefore reflects acoustic suppression only, and is a more reliable baseline for this parameter than the Day 1 figure.

Rack C returned +3.1 (62%), with 46% of listeners scoring ≥ 4 and 69% scoring ≥ 3 . Slightly lower than Day 1 in absolute terms, but produced by a much larger and more skeptical listener pool under a single suppression vector rather than two.

5. Transient Speed — Day 2

⚠ Environmental note: Open convention floor ambient noise can mask the contrast between sharp transient onset and silence that makes speed perceptible. Day 2 was free from the network degradation that additionally suppressed this parameter on Day 1. The Day 2 result is accordingly a more reliable estimate of this parameter's performance under real-world convention conditions.

Rack C returned +2.8 (56%), with 38% of listeners scoring ≥ 4 and 60% scoring ≥ 3 .

6. Decay & Trail — Day 2

⚠ Environmental note: This parameter requires the listener to follow a note all the way into near-silence. Convention floor ambient noise directly competes with and masks this low-amplitude trailing information. Day 2 was free from the network degradation that additionally affected this parameter on Day 1. The Day 2 figure is a more isolated measure of acoustic-only suppression.

Rack C returned +2.6 (52%), with 31% of listeners scoring ≥ 4 and 60% scoring ≥ 3 .

7. Black Background — Day 2

⚠ Environmental note: This is the parameter most directly suppressed by the open-floor environment on both days. Black Background measures the absence of noise, grain, and haze — the silence between notes. The convention floor introduced RF interference, acoustic noise, and competing electromagnetic fields that directly oppose what this parameter measures. One Day 2 listener wrote directly on his scorecard beside this parameter: "hard to tell in show" — and then scored it +4 anyway. That annotation is the most candid single data point in the two-day dataset. The +2.6 average is an environmental artifact, not a performance ceiling.

Rack C returned +2.6 (54%), with 38% of listeners scoring ≥ 4 and 60% scoring ≥ 3 .

8. Timbre & Body — Day 2

Rack C returned +2.9 (58%), with 40% of listeners scoring ≥ 4 and 70% scoring ≥ 3 . Tonal characteristics are substantially less sensitive to ambient noise suppression than spatial or low-amplitude resolution parameters. This score is among the more reliable in the Day 2 dataset relative to controlled-environment potential.

9. Treble & Air — Day 2

⚠ Environmental note: High-frequency air and extension have moderate sensitivity to ambient noise — convention floor crowd noise and competing systems can partially mask treble resolution. This score should be read as a lower-bound

figure. Note: one listener (Sam Kim) crossed out this score on Rack C, indicating an inability to assess it under floor conditions. That data point was excluded from the average; all other Treble & Air scores are retained.

Rack C returned +3.0 (60%), with 38% of listeners scoring $\geq+4$ and 70% scoring $\geq+3$.

10. Bass Authority & Texture — Day 2

Rack C returned +3.2 (64%) — the highest individual parameter average in the Day 2 dataset and across the full two-day study. 46% of listeners scored $\geq+4$ and 73% scored $\geq+3$. Low-frequency texture and viscosity are among the least sensitive parameters to ambient noise suppression — bass energy penetrates above the noise floor more readily than spatial or fine-detail information. This is accordingly the most reliable parameter result in the Day 2 dataset and the strongest single finding of Day 2.

11. Musicality — Day 2

Rack C returned +3.1 (62%), with 54% of listeners scoring $\geq+4$ — the highest $\geq+4$ rate of any parameter across both days of the study. 75% of listeners scored $\geq+3$. Musicality integrates tonal, rhythmic, and spatial perception simultaneously and is less dependent on any single environmental variable. Its strength on Day 2 — produced by the largest and most skeptical listener pool in the study — mirrors its strength on Day 1 and confirms it as the most environmentally robust parameter in the evaluation.

A Harder Listener Pool — Same Directional Result

Day 2 included several of the most skeptical listeners in the full two-day dataset. One listener scored near-zero across all 22 data points on his card and answered NO to the binary question — meaning he perceived no meaningful difference between any of the three racks, yet concluded that a spec sheet could not describe what he just heard. A second listener recorded negative Rack C scores on three spatial parameters before stepping up strongly to +3 and +4 on the tonal and musical parameters. A third recorded negative scores across both Rack B and Rack C on multiple parameters. All cards are retained without exclusion. The directional pattern — Rack C higher than Rack B, Rack B higher than Rack A — held even within the most skeptical segment of the dataset.

The Press Benchmark

Gary Barnes — editor at Headphone.Guru, a credentialed audio press professional — sat down, listened without coaching, and wrote unprompted across the top of his scorecard:

| "MOST DRAMATIC CHANGE WAS DYNAMIC RANGE"

Barnes scored the Reference system at an average of +2.7 across all 11 parameters — a 54% improvement over stock — and answered NO to the binary question. His spontaneous annotation identified dynamic range as the primary perceptual effect, a finding consistent with the dataset's strongest results appearing on Bass Authority, Musicality, and Timbre — the parameters most directly connected to dynamic resolution. Barnes had no financial relationship with Synergistic Research and no prior contact with the evaluation team on the day of his assessment.

The Binary Question Strengthened

63% of Day 2 listeners who answered the binary question concluded that identical RLC measurements do not describe the perceptual difference they heard — up from 56% on Day 1, with nearly three times as many clear answers. This result was produced by a larger, harder, more skeptical listener population on a day without equipment tampering or network failures.

6. TWO-DAY SYNTHESIS & COMBINED ANALYSIS

6.1 Combined Parameter Results

The following table presents weighted combined averages across both days (15 Day 1 listeners, 48 Day 2 listeners, 63 total) alongside the individual day results for direct comparison.

PARAMETER	DAY 1 B	DAY 1 C	D1 Δ	DAY 2 B	DAY 2 C	D2 Δ	COM B C	COM B %
1. Soundstage Width	1.6	3.1	+1.5	1.5	2.6	+1.1	2.72	54%
2. Soundstage Depth	1.9	3.1	+1.2	1.4	2.8	+1.4	2.87	57%
3. Imaging & Separation	1.9	3.5	+1.6	1.5	3.0	+1.5	3.12	62%
4. Micro-Detail	1.9	3.6	+1.7	1.6	3.1	+1.5	3.22	64%
5. Transient Speed	1.9	3.3	+1.4	1.1	2.8	+1.7	2.92	58%
6. Decay & Trail	1.7	3.1	+1.4	1.3	2.6	+1.3	2.72	54%
7. Black Background	1.7	2.7	+1.0	1.4	2.6	+1.2	2.62	52%
8. Timbre & Body	2.5	3.8	+1.3	1.4	2.9	+1.5	3.11	62%
9. Treble & Air	2.0	3.5	+1.5	1.3	3.0	+1.7	3.12	62%
10. Bass Authority & Texture	1.8	3.3	+1.5	1.4	3.2	+1.8	3.22	64%
11. Musicality	2.0	3.9	+1.9	1.5	3.1	+1.6	3.29	66%
OVERALL AVERAGE	1.9	3.4	+1.5	1.4	2.9	+1.5	3.02	60%

6.2 The Invariant Gap

The most significant finding in the two-day dataset is not the absolute level of the scores. It is the consistency of the gap between Foundation and Reference across two days with fundamentally different listener populations and conditions.

	DAY 1	DAY 2	DELTA VARIANCE
Foundation (Rack B) avg	38%	28%	10%
Reference (Rack C) avg	68%	58%	10%
Foundation → Reference gap	30%	30%	0%

The absolute averages moved between Day 1 and Day 2 — predictably, as a larger and more skeptical listener pool produced lower overall scores. Both Foundation and Reference averages declined by 10 percentage points from Day 1 to Day 2. The gap did not move.

This is the result that requires explanation. If the scores were driven by listener expectation, brand enthusiasm, or psychological priming, the gap would be expected to vary with listener typology. A more skeptical population should produce a compressed gap if the effect is expectation-dependent. Instead, the gap held to the decimal point. Two different populations, hearing the same systems under the same open-floor conditions on consecutive days, independently arrived at the same 30-point differential between Foundation and Reference.

The parsimonious explanation is that the 30-point gap reflects an actual and consistent perceptual property of the difference between these two electromagnetic environments — not a statistical artifact, not a listener bias pattern, and not a coincidence.

6.3 Binary Question — Combined Results

	D1 YES	D1 NO	D2 YES	D2 NO	COMBINED NO%
Clear answers	4	5	10	17	—
% of clear answers	44%	56%	37%	63%	61%
Blank / ambiguous	6	—	21	—	—

The combined binary result — 61% NO across both days, with 37 clear answers out of 63 total listeners — represents the most direct population-level test of whether the perceptual gap between Rack A and Rack C is consistent with the predictions of standard electrical measurement. It is not.

The NO result strengthened from Day 1 to Day 2 despite a larger, more skeptical, and more technically informed listener pool. If the binary question were being driven by SR brand loyalty or audiophile tribalism, the result would be expected to weaken as the listener pool diversified. It strengthened.

6.4 Cross-Day Pattern Analysis

Parameters That Held Strongest Across Both Days

Three parameters showed the most robust results across both listener populations: Musicality, Bass Authority & Texture, and Micro-Detail. These three parameters also share a functional characteristic: they are the least susceptible to ambient noise masking. Musicality is a holistic response — listeners either feel compelled by the music or they do not. Bass Authority measures low-frequency texture and viscosity, which remain perceptible above a moderate noise floor. Micro-Detail, while sensitive to noise, benefited from listeners who actively focused on fine-grain texture rather than spatial presentation.

Parameters Most Suppressed by Environment

Black Background, Soundstage Width, and Decay & Trail were consistently the lowest-scoring Rack C parameters across both days. All three share a common environmental vulnerability: they measure the system's ability to resolve information in the low-amplitude domain — silence, space, and the trailing edge of notes. The convention floor noise floor worked directly against all three. Their combined two-day averages are the most conservative estimates in the dataset relative to their expected performance in a controlled environment. On Day 1 specifically, Micro-Detail and Transient Speed were additionally suppressed by network degradation affecting source signal quality.

Listener Typology Distribution

The two-day dataset captures a broad spectrum of listener typologies:

- High-enthusiasm listeners (Rack C averages $\geq +4.0$): Luke Spencer (Day 2) scored every Rack C parameter at +4 or +5. Allan Rohde (Day 1) averaged +4.6 on Rack C. These listeners anchor the upper bound of the dataset.
- Moderate positive listeners (Rack C averages +2.5 to +3.9): The largest segment across both days. Gary Barnes (+2.7) and Matthew Griffith (+3.1) represent this category — listeners who heard a clear and consistent improvement without scoring at ceiling.

- Conservative listeners (Rack C averages +1.0 to +2.4): Oleg Stanislavov (Day 1, +2.4), Dustin C. (Day 2, +3.1), and Giordano Zacchini (Day 2, +1.8) fall in this range.
- Skeptical listeners (Rack C averages near 0 or negative): Charles Li (Day 2) averaged effectively 0 across all parameters. Timothy Lam (Day 2) recorded negative scores in both Rack B and Rack C. Both are retained without exclusion.

The directional finding — Rack C higher than Rack B, Rack B higher than Rack A — held across all four typology categories. Even the most skeptical listeners in the dataset did not systematically score Rack C below Rack B.

7. CONCLUSIONS

7.1 Primary Finding

Across two days, 63 independent listeners, and 693 data points collected under conditions systematically adverse to the hypothesis under examination, the data does not support the claim that the electromagnetic environment is an inaudible variable in high-end audio system performance.

The Reference-level electromagnetic environment produced a combined two-day average of 60% of maximum possible improvement over a stock system. The Foundation-level environment produced a combined average of 30% improvement. The gap between them was 30 percentage points on both days, across two different listener populations, under identical open-floor conditions.

7.2 The Measurement Paradox

61% of listeners who answered the binary question concluded that identical RLC measurements do not describe the perceptual difference they heard. This is the study's most direct finding regarding the relationship between standard electrical measurement and perceptual reality in high-end audio.

The standard electrical measurements of the three racks — resistance, inductance, and capacitance — are effectively identical. The systems are, by any conventional engineering specification, the same. The data presented in this document indicates that 63 independent listeners, including a credentialed audio press professional, heard something that the measurements did not predict. The gap between what measurements capture and what listeners perceive is the central unresolved question in high-end audio engineering. This dataset adds 693 data points to that question.

7.3 Environmental Context and Score Interpretation

All scores in this document are lower-bound figures. The evaluation was conducted in one of the most acoustically hostile environments available — an open convention floor at a major audio show. The parameters most directly affected by SR technology are the parameters most suppressed by that environment. A future evaluation conducted in an enclosed or acoustically treated environment would be expected to produce materially higher scores on every parameter, with the largest improvement anticipated on Black Background, Decay & Trail, and Soundstage parameters.

The fact that the 30-point Foundation-to-Reference gap held under these conditions — rather than being compressed by environmental noise — is itself a finding. The gap is robust to acoustic adversity.

7.4 Limitations and Future Directions

This study has several limitations that future iterations should address:

- Enclosed booth evaluation to reduce ambient noise suppression of the parameters most sensitive to electromagnetic field treatment
- Stable dedicated network infrastructure to eliminate source signal degradation from connectivity issues
- Pre-test belief capture to enable segmentation of expectation-bias effects from genuine perceptual response
- Post-test belief capture using matched pre-test language to document belief change events
- Digital scorecard submission to eliminate ambiguous marking interpretation and enable real-time data analysis
- Increased listener throughput via additional simultaneous evaluation stations
- Plain-language reframing of the binary question to eliminate vocabulary-gap blank responses
- Formal statistical testing including Wilcoxon signed-rank tests on the B vs. C parameter distributions

A next-generation evaluation program addressing all of these limitations is in development. The infrastructure established at CanJam NYC 2026 — including the scoring methodology, the AI analysis framework, and the white paper publication format — provides the foundation for that program.

8. DISCLOSURES & TRANSPARENCY STATEMENT

- This study was conducted by Synergistic Research at CanJam NYC 2026. Synergistic Research designed the evaluation, provided the equipment, and collected the scorecards.
- Statistical analysis was performed by Claude, an AI system developed by Anthropic. The AI was given no instruction to produce favorable results or omit unfavorable data.
- All raw scorecard data is published in the Appendix to this document. No scores were excluded except where explicitly documented with rationale in the Flagged Readings section.
- Listeners were not compensated for participation beyond eligibility in a nightly raffle for SR products. The raffle was open to all scorecard completers regardless of the direction of their scores.
- No listener was coached on what to hear or informed of the expected direction of the results before completing their scorecard. Listeners were aware of what they were hearing across all three systems — a condition identical to every real-world audio purchase evaluation. The cost of each system's components was not disclosed during evaluation.
- Gary Barnes of Headphone.Guru had no financial relationship with Synergistic Research and was not aware his scorecard would be referenced in this document at the time of his evaluation.
- The SRX XL power cord (\$13,000) was deliberately excluded from Rack C. Results therefore represent a conservative implementation of the SR Reference system.
- The reduced Day 1 listener count is attributable to a pervasive network failure outside SR's control, not to listener disinterest or evaluation design failure. Booth traffic significantly exceeded the number of completed scorecards on Day 1.

APPENDIX: RAW SCORECARD DATA

The complete raw scorecard data for all 63 listeners across both days — including all flagged readings, anomalous scores, and binary question responses — is available in the published Day 1 and Day 2 statistical reports, which are companion documents to this white paper.

Day 1 Report: 15 listeners | March 7, 2026 | Full per-parameter breakdown and flagged readings

Day 2 Report: 48 listeners | March 8, 2026 | Full per-parameter breakdown and flagged readings

Both companion documents were produced using the same AI analysis methodology described in Section 2.4 of this white paper and are available from Synergistic Research upon request.

Prepared by Synergistic Research | AI Analysis by Claude (Anthropic)
CanJam NYC 2026 | Data collected March 7–8, 2026 | Published March 2026
synergisticresearch.com