
04. BASS MIXING
ADVANCED SPECTRAL
INTERLOCKING
COMPOSITION
BOOK

Advanced Spectral Interlocking: Three-Tier Bass Fusion

A rigorous, three-tier workflow combining frequency splitting, spectral interlocking, and hierarchical sidechaining to produce bass that is massive, clean, and translationally stable across all playback systems. This document is designed for professional producers and audio engineers seeking surgical control over low-end architecture.

Core Philosophy: Why Three Tiers?

Traditional bass mixing relies on a two-part split, sub and everything else. While functional, this approach creates a fundamental problem: the "everything else" band contains both the punch frequencies (80-200 Hz) and the texture frequencies (200 Hz and above), forcing a single layer to serve two conflicting sonic roles. The result is either muddy mids or anemic top-end, rarely both at once. The three-tier system resolves this by assigning each frequency region its own dedicated source track, eliminating role conflict at the architectural level.

By expanding from two bands to three, Lows, Mids, and Mid-Highs, and applying this split across three discrete bass sources, the engineer gains a 9-track matrix with granular, independent control over every region of the bass spectrum. Each track has a singular, non-negotiable purpose. There is no ambiguity about what any given layer should do, and no layer is asked to perform outside its designated role.

The philosophy rests on three interlocking principles. First, separation before fusion: every bass source is split into three bands before any processing occurs, ensuring that EQ, compression, and saturation decisions are band-specific rather than broadband compromises. Second, spectral interlocking: layers are not simply stacked, they are dynamically carved around each other using frequency-selective sidechain routing, so that each band occupies its space without masking the others. Third, **h**ierarchical sidechaining: less-sub bass triggers ducking on more-sub bass, creating a rhythmic priority chain that preserves punch while maintaining weight.

This approach is particularly powerful in dense electronic productions where multiple bass layers compete for the same spectral real estate. When three bass sources, each with its own character, saturation profile, and harmonic content, are simply layered without spectral management, the result is phase cancellation, frequency masking, and a bass that sounds enormous in solo but collapses under the kick in the full mix. The three-tier system prevents this by enforcing discipline at every stage of the signal path.

The end goal is not complexity for its own sake, it is controlled power. A bass that feels unified and enormous while remaining translationally honest: consistent on club systems, earbuds, car stereos, and studio monitors alike. Every design decision in this workflow serves that goal, from linear-phase crossover EQs to mono-enforced lows to frequency-matched sidechain triggers.

Separation

Split each source into three dedicated bands before any processing. No broadband compromises.

Interlocking

Layers dynamically carve around each other via spectral sidechain routing, not static EQ cuts.

Fusion

All nine bands glue back into a single cohesive instrument, sidechained to the kick as one unit.

The 9-Track Layer Matrix

Splitting three distinct bass sources into three frequency bands yields a complete 9-track matrix. This structure is not arbitrary, it is the logical consequence of applying the three-tier philosophy across multiple bass sources. Each cell in the matrix has a precise, non-overlapping role. Understanding the matrix is prerequisite to understanding every processing and routing decision that follows.

The three sources are characterized by their sub content: Bass 1 contains the most sub (the foundation layer), Bass 2 is the mid-range bass (the punch layer), and Bass 3 contains the least sub (the texture layer). This ordering, from most-sub to least-sub, is critical because it determines the direction of all hierarchical sidechain routing in Phase 2. The matrix below maps every track to its spectral role and its source identity.

Source	Mid-Highs (Air / Width)	Mids (Body / Crunch)	Lows (Foundation / Sub)
Bass 1 - Most Sub	h1 - Texture tail	m1 - Harmonic body	l1 - Primary sub foundation
Bass 2 - Mid Bass	h2 - Edge and grit	m2 - Core punch zone	l2 - Sub reinforcement
Bass 3 - Least Sub	h3 - Stereo air and fizz	m3 - Aggressive mid trigger	l3 - Minimal sub presence

Each column in this matrix represents a frequency tier that will be processed and routed as a unit. The Low column (l1, l2, l3) feeds the Fused Bass Bus in strict mono. The Mid column (m1, m2, m3) is width-limited below 300 Hz and serves as the primary sidechain trigger source. The Mid-High column (h1, h2, h3) carries the stereo image and is the primary sidechain target. This column-wise organization is what makes spectral interlocking possible, each tier can be addressed independently in the routing phase.

It is worth emphasizing that the naming convention (h1, m1, l1, etc.) is not merely shorthand, it is a functional map. When the routing documentation specifies "m2 triggers ducking on h2," the engineer knows exactly which two tracks are involved and can set up the sidechain without ambiguity. In a 9-track bass system, clear naming is not optional; it is a technical requirement for maintaining signal flow integrity.

Practically speaking, each of the nine tracks should be color-coded and grouped in your DAW according to its tier. All Low tracks (l1, l2, l3) should share a color and be pre-routed to a "Lows" subgroup. All Mid tracks (m1, m2, m3) should share another color and route to a "Mids" subgroup. All Mid-High tracks (h1, h2, h3) route to a "Mid-Highs" subgroup. These three subgroups then feed the Fused Bass Bus. This organizational discipline is essential for managing the complexity of a 9-track bass system without losing track of signal flow during the mixing process.

One common mistake when first implementing this matrix is treating the three sources as interchangeable. They are not. Bass 1, Bass 2, and Bass 3 should each be designed with their tier role in mind from the synthesis stage. Bass 1 should be a sub-heavy waveform (sine, triangle, or heavily filtered saw). Bass 2 should emphasize the 80–200 Hz punch region. Bass 3 should be a harmonically rich, top-heavy source with minimal sub content. Designing the sources correctly at the synthesis stage dramatically reduces the EQ work required during the splitting phase.

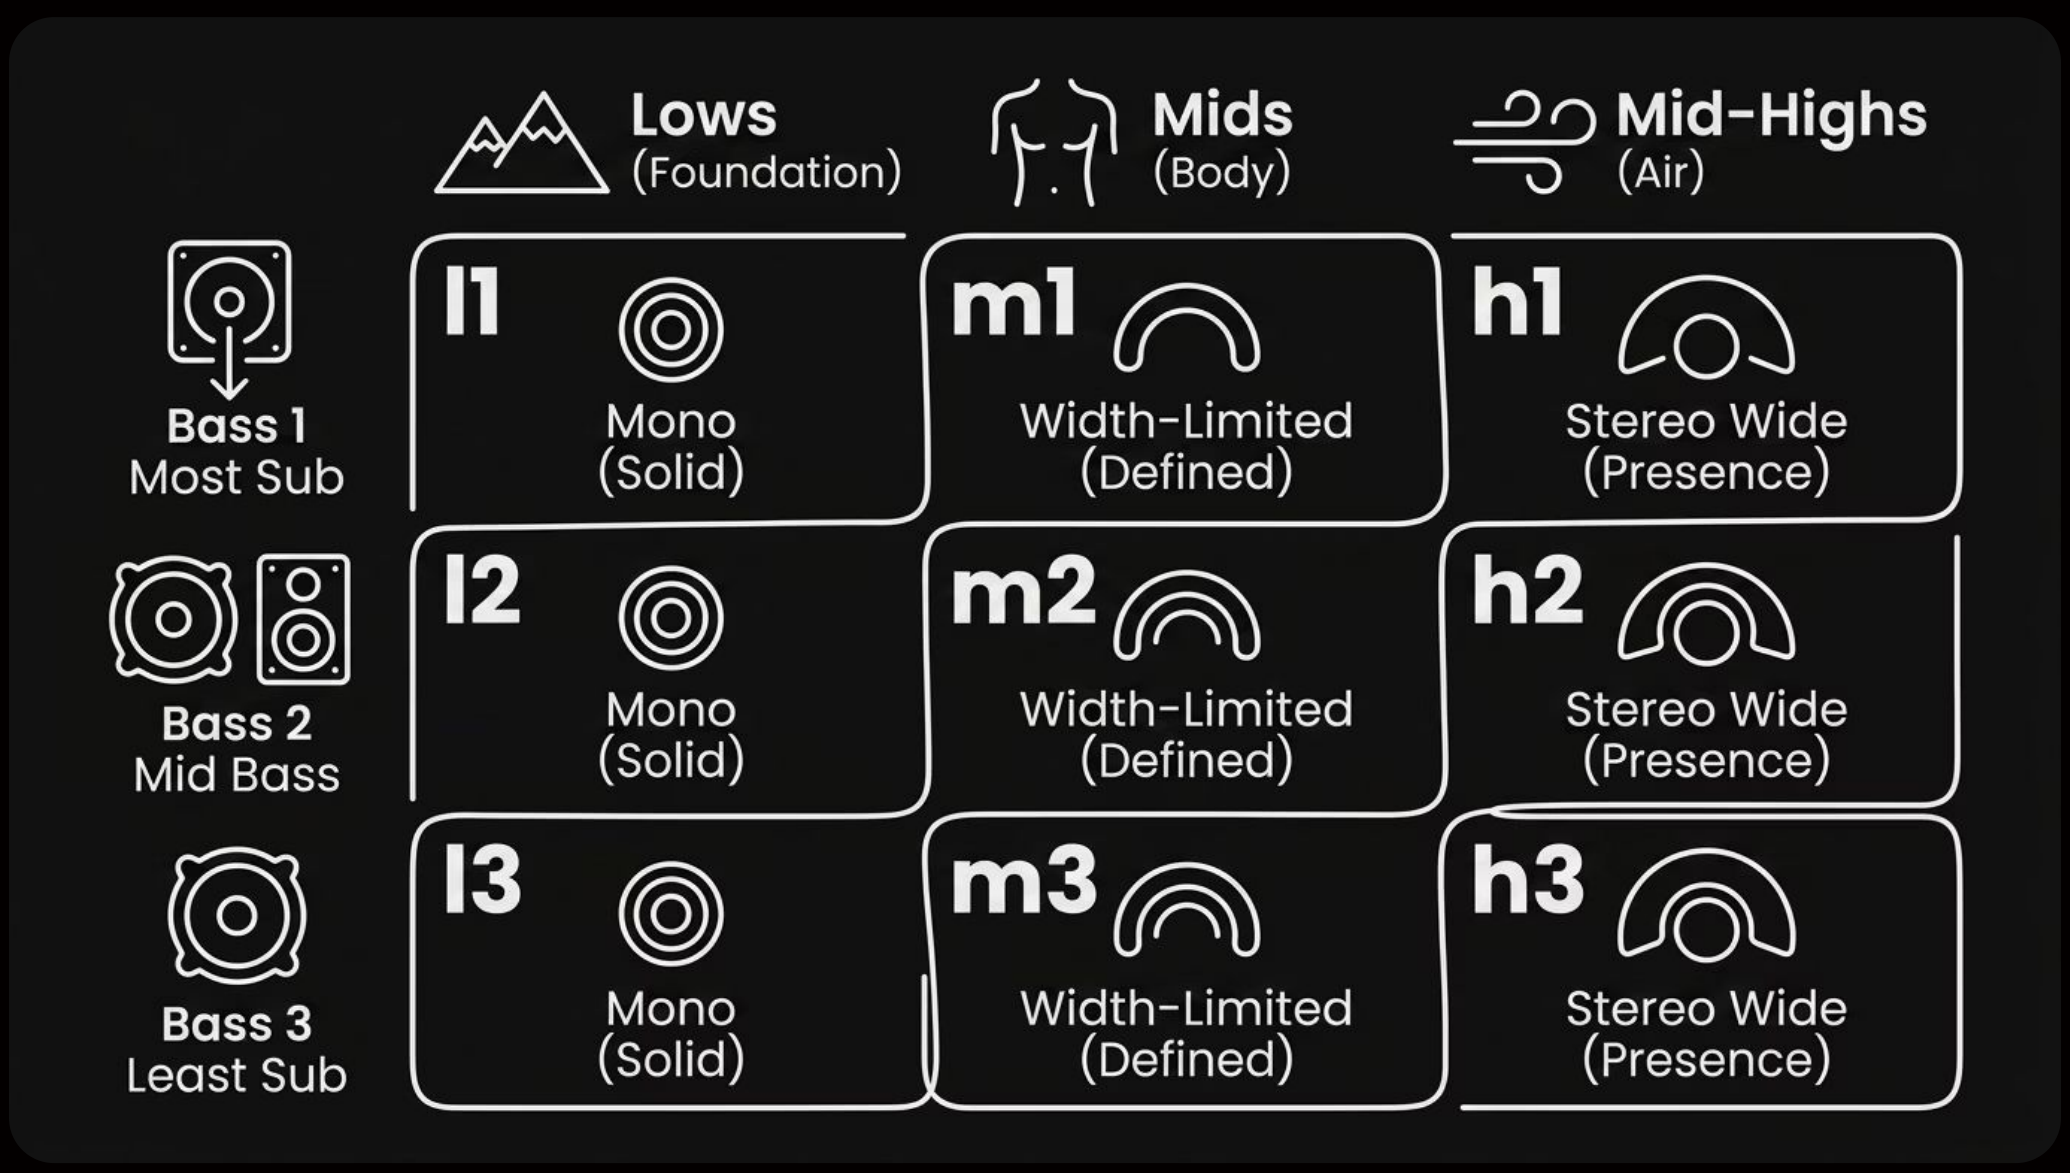

Architecture & Signal Flow

Before any processing decisions are made, the signal routing architecture must be understood in full. Each bass source follows an identical three-path split, but the processing within each path differs according to the tier's role. The diagram below illustrates the routing for a single bass source, this exact architecture is replicated for all three sources, with source-specific EQ frequency adjustments.

The master bass source is the starting point, a fully designed, unprocessed bass sound. From this single source, three parallel signal paths are created using sends or duplicate tracks. Each path is immediately filtered with a linear-phase EQ to isolate its designated frequency band. The Low path applies a low-pass filter at approximately 95 Hz. The Mid path applies a band-pass between approximately 95 Hz and the start of the Mid-High region. The Mid-High path applies a high-pass filter. These three filtered signals are then routed to the Fused Bass Bus, where they are recombined and processed as a unified instrument.

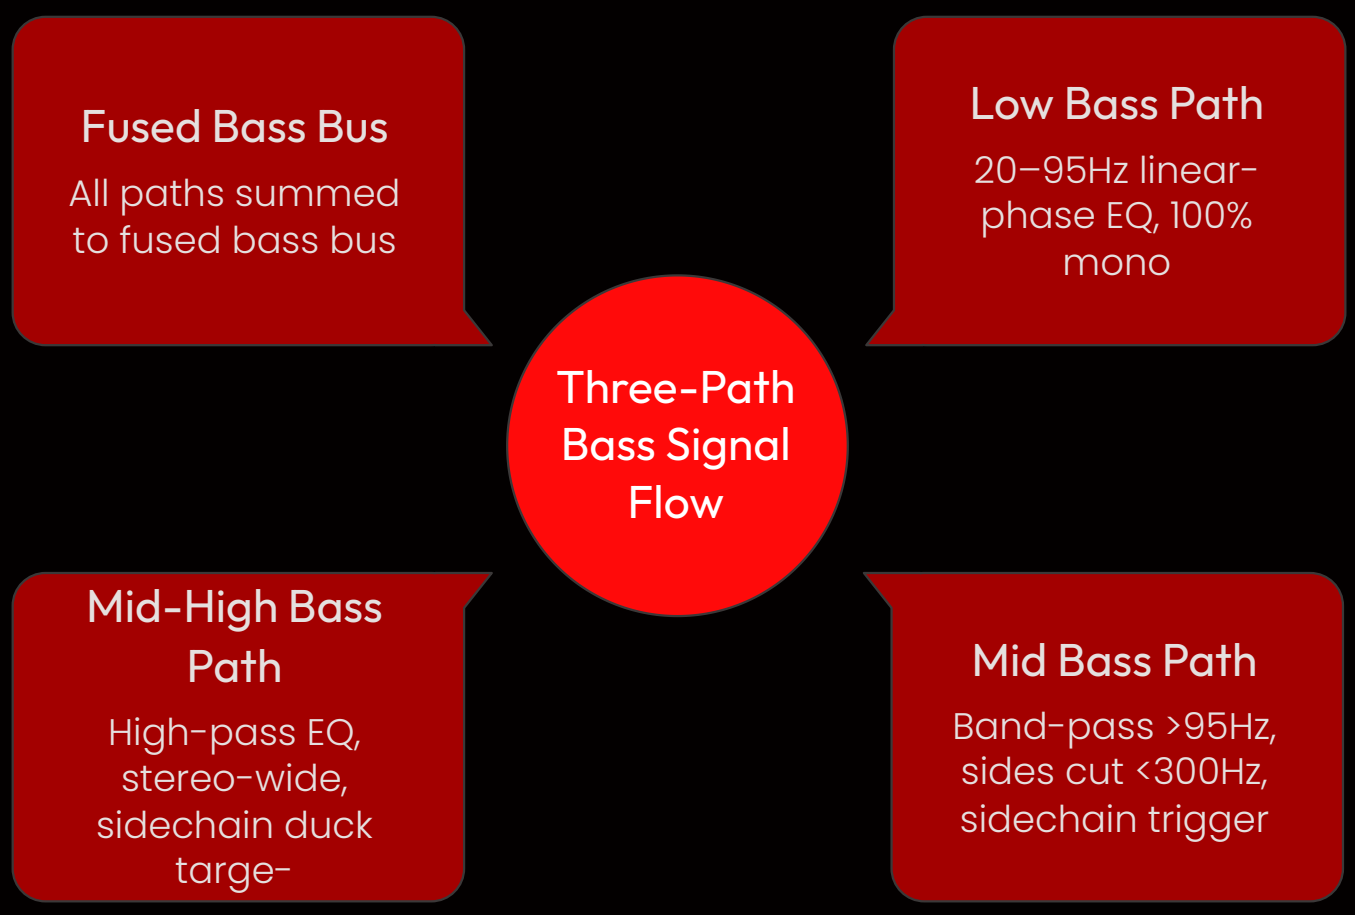

The critical detail in this architecture is that all three paths converge at the Fused Bass Bus, not before it. Each band is processed independently up to the point of fusion. This means EQ, saturation, compression, and width processing on the Low track has no effect on the Mid or Mid-High tracks, and vice versa. The only point at which the bands interact is at the bus, where spectral sidechain routing creates the interlocking behavior described in Phase 2.

The linear-phase EQ requirement at the split stage is non-negotiable for maintaining phase coherence at the crossover points. A standard minimum-phase EQ will introduce phase shift that causes the bands to interfere with each other at the crossover frequency, resulting in comb filtering, a hollow or "sucked out" quality in the transition region, and inconsistent bass response across different playback systems. Linear-phase EQs eliminate this problem by maintaining phase relationships across the entire frequency spectrum, ensuring that the three bands sum coherently at the fusion bus.

Recommended linear-phase EQs for this application include FabFilter Pro-Q 3 (in linear phase mode), Voxengo Curve EQ, or Brainworx bx_digital V3. When setting crossover slopes, matched slopes on adjacent bands are essential, if the Low track uses a 24 dB/octave low-pass at 95 Hz, the Mid track should use a 24 dB/octave high-pass at the same frequency. Mismatched slopes create asymmetric phase behavior and uneven energy distribution at the crossover point.

A common routing mistake is to apply the frequency split after individual track processing rather than before it. The split must happen first, the source is duplicated, then each duplicate is filtered, then each filtered duplicate is processed. Reversing this order means that saturation, compression, and EQ applied to the full-band source will behave differently on each band, undermining the precision of the tier system. Always split first, process second, fuse last.

Phase 1 - Advanced Frequency Splitting

Frequency splitting is the foundational operation of the entire workflow. Done correctly, it creates three clean, phase-coherent bands that can be processed and routed independently. Done incorrectly, it introduces phase cancellation, spectral gaps, and translation problems that no amount of downstream processing can fix. This phase demands precision, discipline, and the right tools.

The three bands are not simply "low, mid, and high" in a generic sense. Each has a precise frequency definition, a width mandate, and a processing philosophy that must be respected throughout the entire mix. The crossover frequencies specified here, approximately 95 Hz between Lows and Mids, and approximately 300 Hz for mid-side width management, are starting points, not absolutes. They should be tuned to the specific content of your bass sources, but the principle of clean, non-overlapping band separation must never be compromised.

Low Track - The Foundation

Strictly mono. Frequencies below approximately 95 Hz. This layer must never have stereo information, phase cancellation between left and right channels at sub frequencies causes the bass to disappear on mono playback systems, which includes most club PA subwoofers. Use a linear-phase low-pass at the crossover point with a steep slope (24 dB/octave minimum). No saturation, no width processing, no exceptions. The only processing on this track beyond the split EQ should be gentle compression for level consistency.

Mid Track - The Body

Band-pass filtered above approximately 95 Hz. The critical processing here is mid-side width management: cut the side channel below 300 Hz using a mid-side EQ or a dedicated mid-side processor. This removes low-mid mud from the stereo field while preserving the mono punch. The result is a bass that feels wide and powerful without sacrificing low-end focus. This track is the primary sidechain trigger source for the spectral interlock in Phase 2.

Mid-High Track - The Texture

High-pass filtered. Full stereo width permitted and encouraged. This is the layer responsible for air, fizz, harmonic complexity, and extreme stereo image. It should feel spacious and detailed without competing with the Mid track for energy in the 100–300 Hz region. This track is the primary sidechain target — it ducks dynamically whenever the Mid track is active, creating the spatial interlock that prevents masking.

⚠ **Phase Watch-Out:** Standard minimum-phase EQs introduce phase shift at crossover points, causing the three bands to interfere with each other where they meet. Always use linear-phase EQs or matched-slope crossover pairs. Verify phase coherence by summing all three bands in solo and checking for comb filtering or a hollow quality at the crossover frequencies. A phase correlation meter should read consistently positive on all three bands.

Phase 2 - Fusion Chain & Routing Logic

With all nine bands split and processed individually, the fusion chain is where spectral interlocking happens. This is the most technically sophisticated phase of the workflow, and the one that most distinguishes this approach from conventional bass layering. The core principle is simple: layers interlock rather than mask. Each band dynamically creates space for the bands around it, using frequency-selective sidechain routing rather than static EQ cuts.

The routing logic operates on two axes simultaneously. The first axis is inter-bass routing, bass source versus bass source within the same frequency tier. The second axis is spectral sidechaining, mid frequencies versus high frequencies within the same bass source. Both axes must be configured correctly for the interlock to function as intended. Misconfiguring either axis will result in either no audible effect or, worse, rhythmic pumping that fights against the groove.

Inter-Bass Routing (Bass vs. Bass)

Less-sub bass triggers ducking on more-sub bass. This creates a hierarchical priority chain where the textural bass (Bass 3) takes rhythmic precedence, the mid bass (Bass 2) follows, and the sub foundation (Bass 1) provides the underlying weight. Critically, low frequencies only sidechain low frequencies, and mid frequencies only sidechain mid frequencies — there is no cross-tier triggering.

- **Lows:** l3 → ducks l2 → ducks l1
- **Mids:** m3 → ducks m2 → ducks m1

Compressor settings on these inter-bass sidechains should be fast-attack, medium-release, with 2–4 dB of gain reduction. The goal is rhythmic breathing, not obvious pumping.

Spectral Sidechaining (Mid vs. Mid-High)

Within each bass source, the mid band triggers ducking on the mid-high band. Whenever the dense mid-range hits, the wide top-end dynamically reduces in level. This is the spatial interlock: the mid provides the punch, and the high provides the space, they never compete for the same moment.

- **Bass 1:** m1 triggers ducking on h1
- **Bass 2:** m2 triggers ducking on h2
- **Bass 3:** m3 triggers ducking on h3

These sidechains should be more aggressive than the inter-bass sidechains, 4–8 dB of gain reduction with a fast attack and a release timed to the track's tempo. The effect should be audible as a dynamic "opening up" of the stereo field between mid hits.

The sidechain compressor selection matters significantly here. For spectral sidechaining, a compressor with a clean, transparent character and precise sidechain filtering is ideal, FabFilter Pro-C 2, Waves SSL G-Master Buss Compressor, or Native Instruments Solid Bus Comp all work well. Enable the sidechain high-pass filter on the compressor to ensure that only the mid frequencies (not the sub) are triggering the ducking on the high bands. Set the sidechain filter to approximately 150 Hz to isolate the punch region from the sub region.

Release time calibration is critical for both types of sidechain routing. The release should be timed to the track's tempo so that the ducking recovers before the next transient, creating a rhythmic "breathing" that reinforces the groove rather than fighting it. A useful starting point is to set the release so that gain reduction returns to zero approximately one-sixteenth note after the triggering transient. Fine-tune by ear while listening to the full mix, the sidechain should feel invisible, contributing to groove rather than drawing attention to itself.

Inter-Bass Routing: Hierarchical Sidechain Chains

The inter-bass routing system is a hierarchical sidechain chain that establishes rhythmic priority across the three bass sources. Understanding why this hierarchy exists, and how to tune it, is essential for achieving the intended result: a bass that feels like a single, unified instrument with internal dynamics, rather than three competing layers.

The hierarchy flows from least-sub to most-sub: Bass 3 (texture) triggers Bass 2 (punch), which triggers Bass 1 (foundation). This ordering is deliberate and musically logical. The textural bass (Bass 3) typically carries the most rhythmic information, it is the layer that follows the melody or the most active part of the bassline. By making it the trigger source, the entire bass system breathes in sync with the most rhythmically active layer. The sub foundation (Bass 1) is the target of the chain, meaning it ducks slightly when the upper layers are active, this prevents the sub from overwhelming the punch and texture layers during busy passages.

1	2	3
Bass 3 - Least Sub	**Bass 2 - Mid Bass**	**Bass 1 - Most Sub**
Texture layer. Primary trigger source. Most rhythmically active. Triggers ducking on Bass 2.	Punch layer. Both trigger and target. Receives from Bass 3, triggers Bass 1.	Foundation layer. Final target in the chain. Provides underlying weight. Ducks subtly under Bass 2.

The key technical detail is frequency-matched sidechain triggering. When l3 triggers ducking on l2, the sidechain input to the compressor on l2 must be filtered to only include low frequencies, specifically, the same frequency range that l2 occupies. If the sidechain input includes mid or high frequencies from l3, the compressor will trigger on content that l2 doesn't even contain, causing erratic and musically incoherent gain reduction. Use a sidechain high-pass filter on the compressor set to approximately 95 Hz to isolate the low-frequency trigger.

Similarly, when m3 triggers ducking on m2, the sidechain input should be band-pass filtered to the mid region (approximately 95 Hz to 300 Hz). This ensures that the compressor on m2 only responds to mid-frequency energy from m3, not to the sub or the high-frequency texture. Most modern compressors with sidechain EQ sections, including FabFilter Pro-C 2 and Waves SSL G-Master, allow you to filter the sidechain input directly within the plugin, eliminating the need for a separate EQ in the sidechain send path.

Gain reduction targets should be conservative on inter-bass sidechains. Two to four decibels is typically sufficient. The goal is subtle dynamic carving, the listener should not perceive individual layers ducking in and out, but rather a cohesive bass that breathes as a single unit. If the sidechain is audible as a distinct pumping effect, the threshold is too low or the release is too slow. Reduce the ratio, raise the threshold, or shorten the release until the effect disappears into the groove.

A useful diagnostic technique is to solo the two tracks involved in a sidechain pair (for example, l2 and l3) and listen to the gain reduction meter on l2's compressor while the bassline plays. You should see gain reduction occurring in sync with l3's transients, and the amount of reduction should feel proportional to the musical content. If l3 is playing a sustained note and l2 is ducking continuously, the release is too short. If l3 is playing fast transients and l2's compressor isn't recovering between hits, the release is too long.

Spectral Sidechaining: Mid vs. Mid-High Interlock

While inter-bass routing handles the vertical hierarchy between bass sources, spectral sidechaining handles the horizontal relationship between frequency tiers within a single source. This is where the "interlocking" in spectral interlocking becomes most audible: the mid and mid-high bands of each bass source dynamically trade space, creating a bass that feels simultaneously punchy and wide without ever becoming muddy.

The mechanism is elegantly simple: the mid band of each bass source is routed as the sidechain trigger for a compressor on the corresponding mid-high band. When the mid hits, the high ducks. When the mid releases, the high returns. The result is a bass where punch and width are never in competition, they occupy the same spectral neighborhood but never the same moment in time.

The Trigger: Mid Band (m1, m2, m3)

The mid band of each bass source serves as the sidechain trigger. This is the dense, punch-focused layer that carries the core harmonic energy of the bass. Because it is width-limited below 300 Hz (from Phase 1), it has a focused, mono-dominant character that makes it an ideal trigger source, clean transients, minimal stereo variance, and consistent energy in the punch region. Route the mid band output to the sidechain input of the compressor on the corresponding mid-high band.

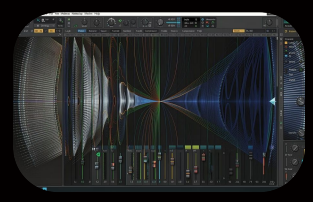

The Target: Mid-High Band (h1, h2, h3)

The mid-high band is the sidechain target, the layer that ducks in response to mid band activity. This is the wide, airy, textural layer that provides stereo image and harmonic complexity. When the compressor engages, the stereo field narrows slightly and the top-end reduces in level. When the compressor releases, the stereo field opens back up. This dynamic movement is what creates the sensation of a bass that "breathes", wide and spacious between hits, focused and punchy on the hits themselves.

Compressor settings for spectral sidechaining should be more aggressive than those used for inter-bass routing. A ratio of 4:1 to 8:1, with a fast attack (1-5 ms) and a release timed to the track tempo, will produce a clear and musical ducking effect. The goal here is not subtlety in the same way as inter-bass routing, the spectral sidechain should be perceptible as a dynamic movement, even if the listener cannot identify the mechanism. A good test is to bypass the spectral sidechain compressor and listen: the bass should feel slightly congested or "stacked." Re-engage the compressor and the bass should feel like it opens up between transients.

One advanced technique is to use different release times for each bass source's spectral sidechain. Bass 3 (the least-sub, most-textural source) can have a faster release, allowing its high-frequency content to return quickly and maintain a sense of air and movement. Bass 1 (the most-sub, most-foundational source) can have a slower release, keeping its high-frequency content more controlled and preventing it from competing with Bass 3's texture. This creates a secondary layer of spectral hierarchy within the interlock itself.

It is also worth noting that the spectral sidechain compressor on the mid-high band should have its sidechain input filtered to exclude frequencies below approximately 150 Hz. Without this filter, the sub content of the mid band will trigger the compressor even when the punch region is not active, causing the high band to duck unnecessarily and reducing the overall sense of air and width in the bass. Most compressors with sidechain EQ sections allow you to set a high-pass filter on the sidechain input, use it here to isolate the trigger to the punch region only.

Phase 3 - Groove Glue: The Fused Bass Bus

Once all nine bands are split, processed, and spectrally interlocked, they must be fused back into a single cohesive instrument. This is the purpose of the Fused Bass Bus, a master bus that receives all nine active bass layers and treats them as a unified source for the final stage of bass processing. Without this bus, the nine tracks remain nine separate instruments. With it, they become one powerful, controlled bass.

The Fused Bass Bus is not simply a summing point, it is an active processing stage with two critical functions. First, it applies a final layer of glue compression to bind the nine bands together into a sonically unified instrument. Second, it receives a sidechain input from the kick drum, ensuring that the entire bass system ducks in perfect sync with the kick's transients. This kick-to-bass sidechain is the final safeguard against low-end conflict, preserving the kick's punch and the bass's weight without either element sacrificing headroom to the other.

All 9 Bands → Fused Bass Bus

Every active bass layer, l1, m1, h1, l2, m2, h2, l3, m3, h3, routes to the Fused Bass Bus. No exceptions. This is the single point of convergence for the entire bass system.

Kick Owns the Bus

The entire Fused Bass Bus is sidechained to the kick drum. All nine bands duck together in response to kick transients, preserving punch and pocket across the full bass spectrum simultaneously.

Optional: Vocal Cleanup

In extremely dense mixes, a light vocal-to-kick sidechain on the Fused Bass Bus, felt, not heard, can enhance rhythmic clarity by creating micro-dips in the bass during vocal phrases. Use sparingly and subtly.

The glue compressor on the Fused Bass Bus should be a clean, transparent bus compressor, the SSL G-Bus, FabFilter Pro-C 2 in bus mode, or the Waves API 2500 are all excellent choices. Settings should be conservative: a ratio of 2:1, a threshold that produces 1-2 dB of gain reduction on peaks, a fast attack (10-30 ms), and an auto or program-dependent release. The goal is cohesion, not compression, the nine bands should feel like they were always one instrument, not like nine instruments being forced together.

The kick sidechain on the Fused Bass Bus requires careful tuning. The attack should be fast enough to catch the kick transient (1-5 ms), but the release is the critical parameter. A release that is too short will cause the bass to pump unnaturally; a release that is too long will cause the bass to remain ducked for too long, reducing its perceived weight and power. A good starting point is a release of 100–200 ms, adjusted by ear while listening to the kick and bass together in the full mix. The bass should feel like it steps aside for the kick and then immediately returns to full weight, a rhythmic "pocket" rather than a sustained dip.

After the glue compressor and kick sidechain, the Fused Bass Bus is the appropriate place for any final tonal shaping, a gentle shelf EQ to balance the overall spectral character, or a subtle saturation stage to add harmonic cohesion across all nine bands. However, these should be finishing touches, not corrective measures. If the bass requires significant EQ or saturation at the bus stage to sound correct, the problem likely originates in one of the earlier phases and should be addressed at the source.

Final Check: Diagnostics & Quality Assurance

The three-tier spectral interlocking system is powerful, but its complexity creates multiple failure points. A systematic diagnostic approach is essential for identifying and resolving problems before they become mix issues. The checks below cover the most common failure modes encountered when implementing this workflow, organized by symptom.

1

Bass Is Loud But Weak

Diagnosis: Phase cancellation. Check phase coherence on all three Low tracks (l1, l2, l3). Use a phase correlation meter, it should read consistently positive. Verify that all Low tracks are in strict mono. Check crossover EQ slopes for matching. Listen in mono to confirm the bass does not collapse.

2

Bass Is Wide But Muddy

Diagnosis: Low-frequency stereo content. Check the Mid tracks (m1, m2, m3) for side content below 300 Hz. Use a mid-side analyzer to verify that the side channel is cut below the crossover point. Verify that no Low track has stereo information. Check that the Mid-High high-pass filter is set correctly and not allowing sub content into the wide layer.

3

Bass Feels Disconnected

Diagnosis: Insufficient glue compression or misconfigured spectral sidechain. Check the Fused Bass Bus glue compressor, it should show 1-2 dB of gain reduction. Verify that all nine bands are routing to the bus. Check spectral sidechain release times, they should be tempo-synced and creating a perceptible "opening up" between mid hits.

4

Kick and Bass Are Fighting

Diagnosis: Kick sidechain misconfigured on the Fused Bass Bus. Verify the sidechain input is receiving the kick signal. Check attack and release settings, attack should be fast (1-5 ms), release should allow full bass recovery before the next kick. Verify that the sidechain high-pass filter on the compressor is not filtering out the kick's fundamental.

> "The resulting bass should feel unified, powerful, and controlled, without sacrificing headroom. If any of these qualities is missing, trace the signal path backward from the Fused Bass Bus until you find the break in the chain."

The ultimate test of a successful three-tier spectral interlock is translation. A properly configured system will sound consistent across club systems, studio monitors, headphones, car stereos, and consumer speakers. If the bass sounds enormous in the studio but disappears on a club system, the problem is almost certainly phase-related in the Low tracks. If the bass sounds wide and detailed on headphones but muddy on speakers, the problem is almost certainly low-frequency stereo content in the Mid or Mid-High tracks. Use these translation checks as your final quality assurance step before committing the mix.

As a final workflow recommendation, save your Fused Bass Bus as a template once it is dialed in. The routing, compressor settings, and sidechain configuration represent significant setup time, and having a template allows you to apply the three-tier system to new productions without rebuilding the architecture from scratch. Document your crossover frequencies, sidechain routing, and compressor settings in the template notes so that you can replicate and adjust the system efficiently on future projects.

9	3
Active Tracks	**Frequency Tiers**
Total bands in the layer matrix across three bass sources	Lows, Mids, and Mid-Highs, each with a dedicated role
2	1
Sidechain Axes	**Fused Output**
Inter-bass routing and spectral mid-to-high interlock	All nine bands unified on a single Fused Bass Bus