



# User Manual

Version 1 · macOS and Windows · October 2026

---

Autonomous dialogue editing and levelling for multitrack AAF sessions.  
Voice detection, bleed rejection, clip-by-clip loudness and fades, delivered as a new AAF.  
Your original file is never modified.



Kraxton Labs LLC · [kraxton.com](https://kraxton.com)

All product and company names mentioned are trademarks of their respective owners and are used only to describe compatibility. See Appendix D.

# Contents

|                                                                   |           |
|-------------------------------------------------------------------|-----------|
| <b>1. Introducing CompDX</b>                                      | <b>6</b>  |
| 1.1 What the software does                                        | 6         |
| 1.2 What the software never does                                  | 7         |
| 1.3 How a job works                                               | 7         |
| <b>2. Requirements and installation</b>                           | <b>9</b>  |
| 2.1 System requirements                                           | 9         |
| 2.2 Installing on macOS                                           | 9         |
| 2.3 Installing on Windows                                         | 9         |
| 2.4 Preparing an AAF with embedded audio                          | 10        |
| 2.5 Licence, trial and activation                                 | 10        |
| 2.5.1 Periodic verification and offline use                       | 11        |
| 2.6 Virtual machines and Windows 11 virtualisation-based security | 11        |
| 2.6.1 Why a physical Windows 11 PC can be reported as virtual     | 11        |
| 2.6.2 How to activate on a physical Windows 11 PC                 | 11        |
| 2.7 Where the application stores its data                         | 12        |
| <b>3. A tour of the interface</b>                                 | <b>13</b> |
| 3.1 Header: session controls and presets                          | 13        |
| 3.2 Tracks panel                                                  | 14        |
| 3.3 Phase panels                                                  | 14        |
| 3.3.1 Working with the controls                                   | 15        |
| 3.4 Run area and footer                                           | 15        |
| <b>4. Quick start: your first job</b>                             | <b>16</b> |
| <b>5. Loading an AAF</b>                                          | <b>17</b> |
| 5.1 Drag and drop or Import AAF                                   | 17        |
| 5.2 AAF Info                                                      | 17        |
| 5.3 External media and skipped tracks                             | 18        |
| 5.4 Choosing tracks                                               | 18        |
| <b>6. Phase 1: Dialogue Detection</b>                             | <b>19</b> |
| 6.1 How detection works                                           | 19        |
| 6.2 VAD Threshold                                                 | 19        |
| 6.3 Pre-Padding and Post-Padding                                  | 20        |
| 6.4 Minimum Pause                                                 | 20        |
| <b>7. Phase 2: Bleed Control (Vertical Judge)</b>                 | <b>21</b> |

|                                                         |           |
|---------------------------------------------------------|-----------|
| 7.1 The problem it solves . . . . .                     | 21        |
| 7.2 How the Judge decides . . . . .                     | 21        |
| 7.3 Enable Bleed Rejection (Vertical Judge) . . . . .   | 22        |
| 7.4 Max Lag . . . . .                                   | 22        |
| 7.5 Kill Threshold . . . . .                            | 22        |
| 7.6 Preserve removed clips (New muted tracks) . . . . . | 23        |
| 7.7 What you will see in the session . . . . .          | 23        |
| <b>8. Phase 3: Leveler Clip by Clip . . . . .</b>       | <b>24</b> |
| 8.1 What levelling means in CompDX . . . . .            | 24        |
| 8.2 Cleaning what the detector let through . . . . .    | 24        |
| 8.3 Enable LUFS Automation . . . . .                    | 25        |
| 8.4 Target LUFS . . . . .                               | 25        |
| 8.5 Leveling Interval . . . . .                         | 25        |
| 8.6 Dynamic Range . . . . .                             | 25        |
| 8.7 Gain staircase . . . . .                            | 26        |
| <b>9. Phase 4: Batch fades . . . . .</b>                | <b>27</b> |
| <b>10. Running a job and using the result . . . . .</b> | <b>28</b> |
| 10.1 Process . . . . .                                  | 28        |
| 10.2 Cancelling . . . . .                               | 29        |
| 10.3 Completion and failure . . . . .                   | 29        |
| 10.4 Importing into Pro Tools . . . . .                 | 30        |
| 10.5 Importing into other applications . . . . .        | 30        |
| 10.6 Reading the output file . . . . .                  | 30        |
| <b>11. Presets . . . . .</b>                            | <b>31</b> |
| 11.1 Factory presets . . . . .                          | 31        |
| 11.2 Managing presets . . . . .                         | 32        |
| <b>12. Support, diagnostics and privacy . . . . .</b>   | <b>33</b> |
| 12.1 The Support button . . . . .                       | 33        |
| 12.2 What the report contains . . . . .                 | 33        |
| 12.3 Logs . . . . .                                     | 34        |
| 12.4 Licence agreement and privacy policy . . . . .     | 34        |
| <b>13. Troubleshooting and tuning . . . . .</b>         | <b>35</b> |
| <b>Appendix A — Parameter reference . . . . .</b>       | <b>36</b> |
| <b>Appendix B — Files and folders . . . . .</b>         | <b>37</b> |

**Appendix C — Glossary . . . . . 38**

**Appendix D — Legal notices . . . . . 39**

## About this manual

This manual describes CompDX Version 1 for macOS and Windows. It explains what the application does, how to prepare an AAF with embedded audio, what every control in the four processing phases means and how it affects the result, how to run a job and how to bring the processed AAF back into your digital audio workstation. It also documents the presets, the licence system, the support report and the most common questions.

**Conventions.** Names of buttons, panels and options appear in **bold**. Values are written as they appear in the interface: milliseconds (ms), loudness units (LUFS), decibels (dB) and picture frames (f). The word **clip** means an audio region on a track exactly as your editing application shows it; **tile** is the shorter internal unit CompDX levels independently (see chapter 8). **Bleed** (also spill or crosstalk) is the sound of one actor picked up by another actor's microphone.

**Trademarks.** Pro Tools, Media Composer, Premiere Pro, DaVinci Resolve, Nuendo and the other product names in this manual belong to their respective owners; they are named only to describe how CompDX works with them. Full notices are in Appendix D.

**Screenshots.** The images in this manual were taken from the macOS version. The Windows version has the same layout and the same controls; only the window frame and the file dialogs differ.

**NOTE** — CompDX never modifies the AAF you load and never renders, re-encodes or normalises audio. Every decision — cuts, clip gain, fades — is written as ordinary edit metadata into a **new** AAF, so the result opens in Pro Tools, Media Composer, Nuendo or Resolve as a normal, fully editable session.

# 1. Introducing CompDX

## 1.1 What the software does

Dialogue editing for film and television starts with hours of mechanical work: finding where each actor speaks, muting or cutting the microphones that only hear the other actor, bringing every line to a consistent loudness and putting a fade on every cut so that nothing clicks. CompDX automates that first pass.

You load an AAF exported from your picture or sound workstation with the audio **embedded**. CompDX reads the timeline, listens to every dialogue track with an on-device neural voice detector, compares the microphones against each other to decide who is really speaking, levels each surviving clip towards a loudness target with clip gain, and writes fades and crossfades at every boundary. The result is a new AAF that you import into your session: cut, cleaned of bleed, levelled and faded, ready for the creative work — restoration, mixing, performance choices — that only a human can do.




Figure 1.1 — The CompDX main window after a completed job: session controls and presets at the top, the track list on the left, the four processing phases in the centre, progress and the Process button at the bottom.


## 1.2 What the software never does

- It never modifies the AAF you load. The input file is opened read-only and a new file is written next to it (or wherever you choose).
- It never renders, bounces, re-encodes or normalises audio. Loudness is applied as **clip gain** metadata, which your workstation applies non-destructively and which you can still adjust or reset clip by clip.
- It never removes tracks you did not select. Unchecked tracks travel to the output exactly as they were.
- It never sends your audio, your session or your project names anywhere. Analysis runs entirely on your computer. The only network traffic is the licence verification described in section 2.5.
- It never uses "AI enhancement" on the sound itself: no denoising, no de-reverb, no voice synthesis. The neural model is used only to *detect* where speech is.

## 1.3 How a job works

A job runs as a fixed pipeline. Every stage receives the decisions of the previous one and refines them; nothing is written until the last stage has validated the whole timeline.

**TIP** — A first run with a factory preset (chapter 11) is the fastest way to understand the result. Import the output into a fresh session next to the original tracks and compare them; then adjust one control at a time.



Every stage works on analysis data only. No audio is rendered or re-encoded: the result is written as edit decisions, clip gain and fades in a new AAF.

Figure 1.2 — The CompDX pipeline. Phases 1 to 4 correspond to the four panels of the interface.

| Stage                               | What happens                                                                                                                                                                                                                                                                                                                                              |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Read AAF</b>                     | The composition is parsed: tracks, clips, source positions, embedded audio, frame rate, sample rate. Tracks with clips are numbered A1, A2, ... in timeline order.                                                                                                                                                                                        |
| <b>Phase 1 — Dialogue Detection</b> | A neural voice-activity model scans each selected track and marks where speech is. Padding and the minimum pause shape the raw detection into editorial clips. The original clip boundaries drawn by the editor are respected and used to rescue short lines the detector is unsure about.                                                                |
| <b>Phase 2 — Bleed Control</b>      | The <b>Vertical Judge</b> compares every pair of tracks over time, measures the acoustic time difference of arrival between microphones and the similarity of their waveforms, and cuts the copies (bleed) while protecting the microphone that is physically nearest to the speaker. Optionally, everything it removes is preserved on new muted tracks. |

| Stage                                 | What happens                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Phase 3 — Leveler clip by clip</b> | Each surviving clip is divided into tiles of the <b>Leveling Interval</b> ; the loudness of every tile is measured and a clip-gain value is written to bring it to the <b>Target LUFS</b> , within a peak guard and a maximum boost. The <b>Dynamic Range</b> control decides how much of an actor's momentary shouts and whispers is preserved. Room-tone tiles that the detector let through are removed; large gain steps are smoothed into a staircase. |
| <b>Phase 4 — Batch fades</b>          | Fade-ins, fade-outs and crossfades of the requested length in frames are written at every clip boundary.                                                                                                                                                                                                                                                                                                                                                    |
| <b>Write AAF</b>                      | A new AAF is assembled with the edited timeline, the embedded audio, the clip gain and the fades, and validated before the job is reported as complete.                                                                                                                                                                                                                                                                                                     |

## 2. Requirements and installation

### 2.1 System requirements

| Component            | Requirement                                                                                                                                                                                                                                 |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| macOS                | macOS 13.3 Ventura or later. Universal application (Apple silicon and Intel), signed and notarised.                                                                                                                                         |
| Windows              | Windows 10 or Windows 11, 64-bit, on a physical PC with hypervisor / virtualisation-based security features disabled. Virtual machines and hypervisor-backed systems are refused by the licensing system (see 2.6).                         |
| Processor and memory | Any current 64-bit CPU. Voice detection runs on the CPU; a 45-minute, 10-track session is analysed in a few minutes on a recent laptop. 8 GB of RAM is a comfortable minimum; large feature-film reels with many tracks benefit from 16 GB. |
| Disk                 | The output AAF contains a copy of the embedded audio, so allow roughly the size of the input AAF as free space in the output folder.                                                                                                        |
| AAF file             | An AAF exported <b>with embedded audio</b> (see 2.4). Mono tracks are processed; stereo and multichannel tracks pass through unchanged. Any sample rate and bit depth supported by the source application.                                  |
| Workstation          | Any application that imports AAF: Pro Tools (recommended), Media Composer, Nuendo, DaVinci Resolve, Premiere Pro.                                                                                                                           |
| Network              | Internet access is needed to start the trial or activate a licence and, afterwards, for periodic verification (see 2.5). No audio or session data is ever transmitted.                                                                      |
| Display              | The window is 1122 × 810 points at 100 % and can be scaled to 75 % or 125 % from the chip in the bottom-right corner.                                                                                                                       |

### 2.2 Installing on macOS

- 1 Open the disk image and drag **Comp DX** to your Applications folder.
- 2 Launch the application. The first launch shows the End User Licence Agreement, then the activation screen.
- 3 If macOS reports that the application was downloaded from the internet, confirm that you want to open it. The application is signed and notarised by Kraxton Labs.

### 2.3 Installing on Windows

- 1 Run the installer and follow the steps.
- 2 Launch **Comp DX** from the Start menu. The first launch shows the End User Licence Agreement, then the activation screen.
- 3 Windows SmartScreen may ask you to confirm the first launch of a newly downloaded application; choose **More info > Run anyway**. The installer is signed with a Kraxton Labs certificate.

## 2.4 Preparing an AAF with embedded audio

CompDX reads the audio **inside** the AAF. An AAF that only *links* to WAV, AIFF or MXF files on disk carries no sound for the detector to analyse: CompDX will warn you when you load it and skip every track. Export with the following settings.

| Application     | Export settings                                                                                                                                                                                                                                       |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pro Tools       | <b>File &gt; Export &gt; Selected Tracks as New AAF/OMF</b> . Set <b>Audio Format</b> to <b>Embedded</b> (not WAV, AIFF or MXF). Enforce Media Composer compatibility can stay on. Export mono dialogue tracks; stereo tracks pass through unchanged. |
| Media Composer  | <b>File &gt; Output &gt; Export to File</b> , format AAF. In the export settings choose <b>Embed Media</b> for audio (not Link to media). Consolidate media if the sequence points to long source files.                                              |
| Premiere Pro    | <b>File &gt; Export &gt; AAF</b> . Tick <b>Embed Audio</b> ; choose a 48 kHz, 24-bit render. Untick "Breakout to mono" only if your tracks are already mono.                                                                                          |
| DaVinci Resolve | <b>File &gt; Export &gt; Timeline &gt; AAF</b> from the Edit page, or Deliver page with an AAF preset. Select <b>Embedded</b> audio (not linked) and render audio as individual mono clips.                                                           |
| Nuendo / Cubase | <b>File &gt; Export &gt; AAF with Embed audio</b> (or "Copy media into AAF") enabled.                                                                                                                                                                 |

**IMPORTANT** — Handles matter. CompDX only ever *shortens* clips (it cuts silence and bleed) and applies gain; it never needs material beyond the clip boundaries. However, the crossfades of Phase 4 are drawn inside the existing clips, so exporting with a little handle on each clip gives your workstation room to extend fades later.

## 2.5 Licence, trial and activation

CompDX uses a licence key issued by Kraxton Labs. On first launch you must accept the licence agreement; the **Activate CompDX** screen then offers two ways in:

- **Activate** — type or paste the licence key you received after purchase and press Activate. Activation requires an internet connection and binds one seat of your licence to this computer.
- **Start Free Trial** — begins the free 10-day evaluation on this computer. The trial is verified with the licence server, so it also needs an internet connection. The same button later reads **Continue Trial (N days left)** and, when the period is over, **Trial Expired**.

The licence status is always visible in the bottom-left corner of the main window: **Seat: Active** for a full licence, or **Trial: N days left** during the evaluation. Clicking the trial chip re-opens the activation screen so you can enter a purchased key. If a seat has been released but the key is still stored on this computer, the application shows a **Seat Released** screen at start-up with **Claim Seat** and **Change License Key**.

**CAUTION** — With a full licence, clicking **Seat: Active** opens a **Release License Seat** confirmation. Choosing **Release Seat** frees the licence for another computer and CompDX on this machine returns to the Claim Seat screen; **Cancel** leaves everything as it was. Use it deliberately when moving to another workstation; claiming the seat back requires an internet connection.

### 2.5.1 Periodic verification and offline use

An activated copy re-verifies its licence with Kraxton Labs in the background. This never interrupts your work and needs only occasional internet access; days of offline work on location or during travel are expected and supported. If the computer stays offline for an extended period the application first shows a discreet reminder with the time remaining, then pauses until the licence can be verified again; verification resumes automatically the moment a connection is available. Your presets and settings are never affected.

## 2.6 Virtual machines and Windows 11 virtualisation-based security

CompDX licences are bound to a cryptographic fingerprint of the physical computer. For security reasons the licensing system refuses to activate or run a trial inside virtual machines, emulators or hypervisor-backed environments, and Kraxton Labs does not issue exceptions to this policy. When this happens the activation screen shows **Virtual Environment Detected**.

### 2.6.1 Why a physical Windows 11 PC can be reported as virtual

Windows 11 can run the operating system itself under the Hyper-V hypervisor to provide virtualisation-based security (VBS). When any of these features is active, the licensing engine cannot reach the real hardware and classifies the system as a virtual machine, even though the PC is physical: **Core Isolation / Memory Integrity**, **Credential Guard**, **Windows Hello Enhanced Sign-in Security**, **Hyper-V**, **Virtual Machine Platform** (installed by WSL 2 and Docker Desktop), **Windows Hypervisor Platform**, **Windows Sandbox**, **Application Guard**, or virtualisation options in the BIOS/UEFI. Many new PCs ship with several of these enabled.

You can confirm the cause with **System Information** (msinfo32): on the Summary page, the line *Virtualization-based security* reads *Running*, and the last line says *A hypervisor has been detected*. Once the steps below are complete, both should disappear.

### 2.6.2 How to activate on a physical Windows 11 PC

- 1 **Turn off Core Isolation.** Windows Security > Device security > Core isolation details > set **Memory integrity** to Off. Restart. (On many systems this alone is not enough; continue with the next steps.)
- 2 **Remove the hypervisor features.** Start > type *Turn Windows features on or off* > untick, where present, **Hyper-V**, **Virtual Machine Platform**, **Windows Hypervisor Platform**, **Windows Sandbox** and **Microsoft Defender Application Guard**. Click OK and restart. Note that WSL 2 and Docker Desktop stop working while Virtual Machine Platform is off.
- 3 **Stop the hypervisor from loading at boot.** Start > type *cmd* > right-click Command Prompt > **Run as administrator**, then run: `bcdedit /set hypervisorlaunchtype off` and restart.
- 4 **Check the firmware.** If msinfo32 still reports a hypervisor, enter the BIOS/UEFI setup and disable *Intel VT-x / VT-d* or *AMD-V / SVM* and, on Intel systems, *Trusted Execution (TXT)*. Some firmware updates re-enable these options.
- 5 Launch CompDX again and activate or start the trial.

**NOTE** — Disabling these features is also what most professional audio vendors recommend to reduce DPC latency on Windows workstations. It does lower the security posture of the PC: weigh that against your needs, and re-enable the features later if you stop using CompDX on that machine.

**IMPORTANT** — On a managed corporate PC these features are usually enforced by group policy and cannot be changed by the user. In that case CompDX cannot be activated on that specific machine; use a non-managed physical PC or a Mac. Kraxton Labs cannot override the policy with a special key.

## 2.7 Where the application stores its data

| Item                                                  | macOS                                                   | Windows                            |
|-------------------------------------------------------|---------------------------------------------------------|------------------------------------|
| <b>Settings (window scale, last folders, sliders)</b> | ~/Library/Application Support/Comp DX/Comp DX.settings  | %APPDATA%\Comp DX\Comp DX.settings |
| <b>Presets</b>                                        | .../Comp DX/CompDX_Presets.json                         | ...\Comp DX\CompDX_Presets.json    |
| <b>Diagnostic logs (one per job)</b>                  | .../Comp DX/Logs/CompDX_Pipeline_<date>.log             | ...\Comp DX\Logs\                  |
| <b>Licence agreement flag</b>                         | .../Comp DX/CompDX_License.json                         | ...\Comp DX\CompDX_License.json    |
| <b>Output AAF (default)</b>                           | A folder named <b>Output AAF</b> next to the input file | Same                               |

## 3. A tour of the interface

The window has four areas: the header with the session controls and the presets, the track list on the left, the four phase panels in the centre, and the run area with the footer at the bottom. Before an AAF is loaded, the centre of the window is a drop zone.




Figure 3.1 — CompDX on launch. Drop an AAF anywhere in the window or click **or Open AAF File**.

### 3.1 Header: session controls and presets




Figure 3.2 — The header. Left: logo and version. Centre: Import AAF, Output Folder and AAF Info with the current file and folder next to them. Right: the Presets bar.

- **Import AAF** — opens a file dialog to load an .aaf. Dropping a file anywhere in the window does the same. The file name appears to the right of the button.
- **Output Folder** — chooses where the processed AAF is written. By default CompDX creates a folder called **Output AAF** next to the input file. The folder name is shown to the right; the final file name is confirmed in a save dialog when you press Process.

- **AAF Info** — opens a read-only summary of the loaded file (section 5.2). Click again, or anywhere else, to close it.
- **Presets** — a drop-down with the factory and user presets and five buttons: new, save, import, export, delete (chapter 11).

## 3.2 Tracks panel




Figure 3.3 — The Tracks panel lists every audio track of the composition with its number, name, AAF track label and clip count.

Each row shows the track as your editing application named it, its position (**A1**, **A2**, ...) and how many clips it holds. Tick the tracks that carry dialogue microphones; untick anything CompDX should not touch — music, effects, temp tracks, reference mixes you want to keep intact. Unchecked tracks are copied to the output unchanged. The list scrolls when a session has more tracks than fit.

**NOTE** — Track numbering follows the order of the tracks in the AAF. Stereo or multichannel tracks are listed but always pass through untouched, because CompDX edits mono dialogue microphones only.

## 3.3 Phase panels

The four panels in the centre are the heart of the interface. Each corresponds to one stage of the pipeline and is described in its own chapter: **Phase 1: Dialogue Detection** (chapter 6), **Phase 2: Bleed Control** (chapter 7), **Phase 3: Leveler Clip by Clip** (chapter 8) and **Phase 4: Batch fades** (chapter 9). Phases 2 and 3 have an enable checkbox; Phases 1 and 4 always run (set the fade lengths to 0 to write no fades).

### 3.3.1 Working with the controls

Every parameter is a horizontal slider with its value on the right. Drag the slider for coarse changes, or click the value to type it: a single click selects the whole value so you can type straight over it; **Return** applies the value (out-of-range entries are corrected on the spot) and releases the field; **Escape** discards what you typed. Clicking anywhere else also applies the value. Values snap to the step shown in Appendix A (for example 5 ms for padding, 0.5 LUFS for the target).

Controls that depend on a disabled option are dimmed and cannot be changed: when **Enable Bleed Rejection** is off, the Max Lag and Kill Threshold sliders and the **Preserve removed clips** checkbox are dimmed, and the same applies to the Phase 3 sliders when **Enable LUFS Automation** is off. A dimmed checkbox still shows its stored state, which takes effect again when its phase is re-enabled.

**NOTE** — While a job is running, every control that feeds it is locked: the four phase panels, the Tracks panel, the Presets bar, **Import AAF**, **Output Folder** and drag and drop. The settings in force when you pressed Process are the ones used for the whole job. Everything unlocks as soon as the job completes, fails or is cancelled.

## 3.4 Run area and footer




Figure 3.4 — The run area: status line, progress bar with percentage and elapsed time, the Process button, and the footer with the licence chip, the Kraxton Labs logo, Support and the window-scale chip.

- **Status line** — "Ready." when a file is loaded, "Scanning AAF structure..." while a file is being read, "Processing." during a job, then "Process Completed." (green) or "FAILED — ..." (red) with the reason on a second line.
- **Progress bar** — percentage and elapsed time of the current job.
- **Process** — starts a job (chapter 10). During a job the button turns red and reads **STOP**; pressing it cancels the job cleanly and removes any partial output.
- **Seat: Active / Trial** — the licence chip (section 2.5).
- **Support** — exports a diagnostic report for Kraxton Labs support (chapter 12).
- **100 %** — window scale: 75 %, 100 % or 125 %. The choice is remembered.

## 4. Quick start: your first job

This chapter runs a complete job with the default factory preset, **TV Show –24dB LUFS**. Each step is explained in depth in the following chapters.

- 1 **Export an AAF with embedded audio** from your workstation (section 2.4). Include the dialogue microphones as mono tracks; a reference mix or the boom can travel along as well.
- 2 **Load it.** Drop the file anywhere in the CompDX window, or click **Import AAF**. The status line reads "Scanning AAF structure..." for a few seconds; then the Tracks panel fills and the status reads "Ready."
- 3 **Check the file.** Click **AAF Info** and confirm the source software, sample rate, video rate and that the media is listed as **embedded**.
- 4 **Choose the tracks.** Leave every dialogue track ticked; untick music, effects and anything you want passed through untouched.
- 5 **Pick a preset.** **TV Show –24dB LUFS** for broadcast-style delivery, **Cinema –27dB LUFS** for a quieter target with more dynamics (chapter 11). Adjust nothing else on the first run.
- 6 **Press Process.** A save dialog proposes <name>\_CompDX . aaf inside the **Output AAF** folder next to your file. Confirm. The progress bar shows the percentage and elapsed time; a 45-minute episode with ten tracks typically takes one to three minutes.
- 7 **Import the result.** In Pro Tools, **File > Import > Session Data** on the new AAF, as new tracks, and compare against the originals. Every cut, clip gain value and fade is an ordinary, editable session element.
- 8 **Refine.** If a line was cut that should have stayed, or bleed survived, chapter 13 tells you which control to move and in which direction. Save your adjusted settings as a preset.

**IMPORTANT** — Always import the processed AAF into a copy of your session or as new tracks. CompDX is a first pass: it is designed to be reviewed, not to replace your judgement on a difficult line.


## 5. Loading an AAF

### 5.1 Drag and drop or Import AAF

CompDX accepts a single .aaf at a time. Drop it anywhere in the window, use **Import AAF**, or click **or Open AAF File** in the empty state. Loading a new file replaces the previous one; your slider settings are kept.

While the file is read the status line shows "Scanning AAF structure..." and the drop zone explains what is being read (mobs, descriptors, components). Large feature-film reels with several gigabytes of embedded audio can take some seconds. Dropping a file of another type shows "Unsupported file type".

### 5.2 AAF Info



| AAF Info           |                      |
|--------------------|----------------------|
| Session name       | AAF TEST COMPDX      |
| Source software    | ProTools 25.12.0.127 |
| Created            | 2026-08-19 01:08:41  |
| Start TC           | 00:00:00:00          |
| End TC             | 24:00:00:00          |
| Composition length | 24:00:00:00          |
| Audio tracks       | 12 (571 clips)       |
| Video tracks       | 0 (0 clips)          |
| Media files        | 571 embedded         |
| AAF version        | AAF 1.1              |
| Container format   | AAF                  |
| Video rate         | 23.976 fps           |
| Sample rate        | 48000 Hz             |
| Bit depth          | 24-bit               |

Figure 5.1 — AAF Info for a Pro Tools export: session name, source software, creation date, timecodes, track and clip counts, media, AAF version, container, video rate, sample rate and bit depth.

| Row                                    | Meaning                                                                                                                                                                            |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Session name / Source software         | The composition name and the application and version that wrote the file. CompDX adapts its reading to the dialect of Pro Tools, Media Composer, Premiere Pro, Resolve and Nuendo. |
| Created                                | Creation date recorded in the file.                                                                                                                                                |
| Start TC / End TC / Composition length | Timeline extent as declared by the exporting application. Pro Tools declares the session length (often 24 hours); this does not affect processing.                                 |
| Audio tracks / Video tracks            | Number of tracks and clips of each kind.                                                                                                                                           |

| Row                                   | Meaning                                                                                                                                                                                       |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Media files</b>                    | Number of audio sources and whether they are <b>embedded</b> . If this row says the media is external (linked), CompDX cannot process the file — re-export with embedded audio (section 2.4). |
| <b>AAF version / Container format</b> | Structural version and container of the file.                                                                                                                                                 |
| <b>Video rate</b>                     | Picture frame rate. Phase 4 fade lengths and the gain staircase of Phase 3 are expressed in these frames.                                                                                     |
| <b>Sample rate / Bit depth</b>        | Of the embedded audio. All tracks must share the sample rate declared here.                                                                                                                   |

## 5.3 External media and skipped tracks

If the AAF only links to media on disk, an orange warning appears as soon as the file is loaded and, if you press Process anyway, every track is skipped and the job fails with an explanation of the export setting to change. CompDX cannot follow links to external files by design: the audio must be inside the AAF.

## 5.4 Choosing tracks

Tick every mono track that carries a dialogue microphone: lavaliers, boom, plant mics. Mix or reference tracks may be included; CompDX recognises that they contain the same signal as the microphones (a digital copy rather than an acoustic one) and does not let them condemn the microphones. Untick music, effects, temp tracks and anything you want to keep exactly as it is — unchecked tracks are copied to the output verbatim.

**TIP** — Stereo and multichannel tracks are always passed through, even when ticked. If you need a stereo pair processed, export it as two mono tracks.

## 6. Phase 1: Dialogue Detection




Figure 6.1 — Phase 1 panel.

### 6.1 How detection works

Phase 1 listens to each selected track with a neural voice-activity model that runs on your computer. For every few milliseconds of audio it produces a probability that a human voice is present. This probability is turned into speech regions with a two-level trigger: a region starts when the probability rises above the **VAD Threshold** and ends when it falls clearly below it, so a momentary dip inside a word does not split it. The regions are then shaped by **Pre-Padding**, **Post-Padding** and **Minimum Pause** into the clips you would draw by hand.

CompDX also knows where the editor placed the original clip boundaries and uses them as evidence. A short line at the edge of a clip — an "Okay", a "Yes", a laugh — often scores lower than the threshold because the model has little context. When such a region lies inside an editor's clip, is close to the threshold and carries the same energy as the rest of that actor's speech on the track, it is **rescued** instead of deleted. Regions that are quiet relative to the actor's level, however confident the model, are left to Phases 2 and 3 to judge.

### 6.2 VAD Threshold

Range 0.05 – 0.95, default 0.40; factory presets use 0.60. The confidence the model must reach for audio to count as speech.

- **Lower values** (0.20 – 0.40) catch whispers, off-mic lines and breathy deliveries, at the price of more false detections on noisy tracks: doors, footsteps, rustle and bleed are more likely to be kept as "speech".
- **Higher values** (0.60 – 0.80) demand clear, present voice. Cleaner results on well-recorded lavaliers; very quiet lines may be missed unless the clip-boundary rescue described above brings them back.
- Start at 0.60. Lower it only if real dialogue is being lost; raise it if silence and room tone are being kept as clips.

## 6.3 Pre-Padding and Post-Padding

Pre-Padding (0 – 500 ms, default 200 ms) is added before the start of every detected region; Post-Padding (0 – 1000 ms, default 400 ms) after the end. They protect the soft consonants at the start of a word (F, S, TH) and the natural decay at the end (final consonants, breath, room reverberation). Both are clamped at the editor's original clip boundaries, so padding never reaches into another clip.

Typical values are 250 – 350 ms before and 300 – 400 ms after. A very reverberant room or a slow speaker benefits from longer post-padding; tight comedy timing with many short lines can use shorter values so that the clips do not run into each other.

## 6.4 Minimum Pause

Range 0 – 2000 ms, default 250 ms. The shortest silence between two detected regions that CompDX is allowed to cut. Pauses shorter than this stay inside one clip; longer ones become a gap. A larger value gives fewer, longer clips and a more natural editorial rhythm (a phrase with its breaths stays whole); a smaller value produces finer cuts and removes more room tone between words.

**NOTE** — Pre-Padding + Post-Padding + Minimum Pause together define the **bridge**: the longest gap that CompDX treats as part of a phrase rather than as a cut. Later phases use the same figure so that pauses which survived Phase 1 are not re-opened by Phase 2 and their room tone is not boosted by Phase 3.

## 7. Phase 2: Bleed Control (Vertical Judge)




Figure 7.1 — Phase 2 panel, with *Preserve removed clips (New muted tracks)* enabled.

### 7.1 The problem it solves

When two actors wear lavaliers, each microphone hears both of them. The voice detector of Phase 1 correctly reports "speech" on both tracks, and a leveler would then bring the bleed up to full loudness. Someone has to decide, moment by moment, which microphone owns the voice. That is the job of the **Vertical Judge**: it looks at the session vertically, across tracks, instead of along one track.

### 7.2 How the Judge decides

Sound travels at roughly 34 cm per millisecond. The voice of actor A reaches A's own lavalier first and B's lavalier a little later; how much later depends on the distance between the two actors. CompDX measures that time difference of arrival between every pair of tracks that are speaking at the same time, in short overlapping windows, using a cross-correlation technique that is robust to level differences and room reverberation.

The Judge works in two passes. It first **measures** every window of every pair — arrival order, stability of that order over neighbouring windows, waveform similarity, and the level of each microphone relative to its own speaking level. Only then does it issue **verdicts**. The rules are physical, not statistical guesses:

- The track whose copy arrives **later**, with a stable delay inside **Max Lag** and a similar waveform, is a copy: it is **cut** for that window. The earlier track is the owner and is protected.
- A copy can never be **louder than its source** once each microphone's own speaking level is taken into account. If the "copy" is the louder, closer signal, the Judge abstains rather than cut the wrong actor — unless the geometry of that pair proves over the scene that both directions occur (two actors facing each other), in which case arrival order is sovereign.
- Two tracks that are **digital duplicates** (a mix track and its microphones, stems, a copied region) have zero delay and near-identical waveforms. They are recognised and neither is condemned by the other in that neighbourhood.

- Delays too short to be acoustic can come from console latency or from a distant source that reaches both microphones at once; they never cut a speaking microphone. When both microphones are far below their own speaking level in such a window, the window is room tone shared by both and both are cut — except inside a phrase, where a breath between two words is handed back to its owner.
- Verdicts are applied at a fine time resolution, then smoothed so that no clip is chopped into confetti. A clip is only removed **whole** when the condemned material is essentially all of it; a clip with an uncondemned run of its own voice — even a very short reply — is **split**, keeping the voice and cutting the bleed.

**NOTE** — The Judge protects the microphone nearest to the speaker, not the loudest track. A boom that is louder than a lavalier but further from the actor will lose to the lavalier where the acoustic evidence says so, and both will be kept where the evidence is inconclusive.

### 7.3 Enable Bleed Rejection (Vertical Judge)

Turns Phase 2 on or off. Turn it off for single-microphone material (one actor, one track, a voice-over) or when you want only detection and levelling. With a single selected track the Judge has nothing to compare and is skipped automatically.

### 7.4 Max Lag

Range 1 – 30 ms, default 10 ms; factory presets use 1 ms. The largest delay between two microphones that CompDX will still accept as acoustic bleed. It corresponds to the distance between actors: 1 ms  $\approx$  34 cm, 3 ms  $\approx$  1 m, 10 ms  $\approx$  3.4 m.

- **1 ms** is the right value for actors in conversation at normal distance and for any scene where bleed is close and clear: it lets the Judge concentrate on the strongest, unambiguous evidence.
- Raise it (3 – 10 ms) for wide blocking — actors across a room, a table of six — where bleed arrives later. Very high values admit more coincidental matches and should be paired with a lower Kill Threshold.

### 7.5 Kill Threshold

Range 10 – 100 %, default 50 %; factory presets use 85 – 90 %. Despite its name this is not a percentage of clips to remove: it sets **how much evidence the Judge needs before it cuts**. At 10 % the Judge is strict — it cuts only with a very clear, stable delay and a well-correlated waveform. At 100 % it is permissive — weaker and less stable evidence is enough. In between, the requirements relax progressively.

- **High values (80 – 95 %)** remove more bleed and are the right choice for typical multi-lavalier dialogue. The physical protections described above (arrival order, level, duplicates, own-voice runs) remain in force at every setting, so a high value does not by itself delete real dialogue.
- **Lower values (30 – 60 %)** are useful when the recording is noisy or reverberant and the Judge is cutting lines you want to keep; they trade some surviving bleed for safety.

## 7.6 Preserve removed clips (New muted tracks)

Off by default, and available only while **Enable Bleed Rejection** is on. Normally the material the Judge removes is simply cut out of the timeline. With this option on, nothing the Judge removes is lost: CompDX moves it to **new muted tracks** added at the end of the output AAF, one for each processed track that lost material, at exactly the same timeline position it had on the original track.

- Each new track is named after its source track followed by a hyphen and the source track number, so the material removed from **A3** appears on **A3-3** and the material removed from a track called **Mix L** in position 1 appears on **Mix L-1**.
- The new tracks are **muted** when the AAF is imported, so they never play by accident or change your mix. Unmute one to audition what the Judge removed.
- The preserved clips keep their **original clip gain**: they are not levelled by Phase 3. They do receive the Phase 4 fades, so a clip you move back to the main track is ready to play without clicks.
- Only what the Judge removed is preserved. Silence removed by Phase 1 and room-tone tiles removed by Phase 3 are not, and padding that the main track kept is not duplicated on the new track.
- The main tracks are identical whether this option is on or off: it only adds tracks, it never changes the edit.

To restore a line the Judge removed, find it on the muted track directly below the processed tracks, at the same position, and drag it (or copy and paste it) back onto the main track. This is much faster than locating it on the original, unprocessed track. When you are satisfied with the edit, delete the muted tracks or keep them hidden in the session as a safety net.

**TIP** — Turn this option on for the first jobs on a new production, or whenever you want to review the Judge's decisions quickly. The extra tracks add no audio to the AAF: they reference the same embedded media as the main tracks.

## 7.7 What you will see in the session

Bleed does not disappear silently: the affected region is cut out of the clip, so the result is visible as a shorter clip or a hole in the track, with Phase 4 fades at the new edges. The main tracks are never muted or gain-reduced to hide bleed. If you disagree with a cut, restore the clip from the muted **Preserve removed clips** track when that option is on, or from the original track next to it.

## 8. Phase 3: Leveler Clip by Clip




Figure 8.1 — Phase 3 panel.

### 8.1 What levelling means in CompDX

Phase 3 brings every surviving piece of dialogue to a consistent loudness using **clip gain** — the same non-destructive gain you would draw on each clip in Pro Tools. The audio is never processed: the measured loudness of each piece is compared with the target and the difference is written as a gain value that your workstation applies on playback and shows on the clip.

Each clip is divided into **tiles** of the **Leveling Interval**. The integrated loudness of each tile is measured (ITU-R BS.1770, in LUFS) and a gain is computed to bring it to the **Target LUFS**, subject to a **peak guard** that never lets the boosted peak exceed  $-3$  dBFS and to a bounded range of boost and cut. A tile that contains a single sharp transient (a plosive, a click, a slammed door) is handled with a **transient splitter**: the few milliseconds around the peak are isolated and protected so that the rest of the tile still receives the gain the dialogue needs, instead of the whole line being punished for one spike.

### 8.2 Cleaning what the detector let through

Not every tile is speech. Room tone, faint bleed the Judge could not attribute and rustle can arrive at Phase 3 as very quiet "speech" tiles that a naive leveler would boost enormously. CompDX applies a **speech floor**: tiles far below any plausible dialogue level, with no speaking-level peaks, are deleted rather than boosted, unless the Judge already certified that microphone as the owner of that moment. A short quiet tile trapped between two measured tiles of the same phrase — the tail of a word, a breath — **inherits** the gain of its neighbours instead of being measured on its own. Near-silent tiles isolated between non-speech are removed.

Genuine **whispers** are protected by the opposite rule: a quiet tile that the engine certifies as speech is treated as dialogue and brought to target (within the peak guard), so an intimate scene is not thrown away as noise. And a pause inside a phrase that is much quieter than the words around it (a **valley**) takes the lower gain of its neighbours so that room tone is not pumped up between words.

## 8.3 Enable LUFS Automation

Turns Phase 3 on or off. Off, clips keep their original gain and only the cuts of Phases 1–2 and the fades of Phase 4 are written.

## 8.4 Target LUFS

Range –36 to –8 LUFS in 0.5 steps, default –24. The loudness every tile is levelled towards. –24 LUFS is the anchor of the ATSC A/85 broadcast standard and the value of the **TV Show** preset; –23 LUFS matches EBU R 128; –27 LUFS, the **Cinema** preset, leaves headroom for a theatrical mix. CompDX levels dialogue tiles, not the programme: the final programme loudness is set in your mix.

## 8.5 Leveling Interval

Range 250 – 3000 ms in 10 ms steps, default 1000 ms; factory presets use 350 – 500 ms. The length of the tiles.

- **Short intervals** (250 – 500 ms) follow the performance closely and even out word-to-word variation: the most consistent intelligibility, at the cost of more gain changes per clip.
- **Long intervals** (1000 – 3000 ms) average whole phrases and preserve the natural macro-dynamics of the delivery with fewer, smoother gain changes.

Tiles that are shorter than the interval because a cut fell inside them (the last word before a Judge cut, for example) are not measured on their own — they inherit the gain of the adjacent tile of the same phrase, so a clipped word is never boosted to a different level than the sentence it belongs to.

## 8.6 Dynamic Range

Range 0 – 100 %, default 0; factory presets use 30 – 45 %. Levelling every tile to the same target flattens a performance: a shout and a murmur end up at the same loudness. **Dynamic Range** controls how much of the actor's **momentary excursions** is preserved.

CompDX computes a running loudness reference for each actor from the surrounding scene and looks at how far each tile departs from it. Tiles close to the reference are levelled to target as usual, so the overall law of the leveler never changes. Tiles that depart clearly for a short time (a shout, a sudden whisper) are compressed rather than flattened: at 0 % the excess is removed entirely and the tile lands on target; at 100 % the whole excess is kept and the shout stays as far above the target as it was above the actor's own average; at 30 % roughly a third of the excursion survives. Sustained changes of level — the actor moving off-mic for a whole scene — are still fully corrected.

- **0 – 20 %:** broadcast promo, documentary interviews, anything that must sit at a fixed level.
- **30 – 45 %:** drama and comedy; shouts stand out, drops in level are less over-compensated. This is what the factory presets use.
- **60 – 100 %:** theatrical dialogue where you want the performance almost untouched and only the average corrected.

## 8.7 Gain staircase

When two adjacent tiles of the same clip differ by more than a few decibels, a single clip-gain step would be audible as a jump. CompDX inserts a **staircase**: part of the louder tile is used to write a series of small intermediate steps, each a whole number of picture frames long and aligned to the frame grid, between the two values. Every step is peak-guarded against the audio it covers. The result in the session is a short sequence of small clip-gain steps that the ear reads as a natural, continuous change.



Figure 8.2 — A completed job. The status line turns green; the button returns to Process.

**TIP** — Clip gain written by CompDX is ordinary Pro Tools clip gain: select the clips and use **Clip > Clip Gain > Clear** to reset a passage, or nudge it with the clip-gain fader, exactly as if you had drawn it yourself.

## 9. Phase 4: Batch fades




Figure 9.1 — Phase 4 panel. Values are picture frames of the AAF.

Every boundary created by Phases 1 and 2 is a hard cut until Phase 4 writes a transition on it. Fade lengths are expressed in **frames** of the composition's video rate (shown in AAF Info) so that they match the way editors think and stay proportional across 23.976, 25 and 29.97 fps material. Fades are written as standard AAF transitions and appear in the session as ordinary fade handles you can drag.

| Control          | Range    | Effect                                                                                                                                                                                                             |
|------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Fade In</b>   | 0 – 15 f | Length of the ramp at the head of every clip whose left edge faces a gap. 1 – 2 frames removes the click of a cut without softening the first consonant.                                                           |
| <b>Fade Out</b>  | 0 – 15 f | Length of the ramp at the tail of every clip whose right edge faces a gap. 2 – 3 frames lets the room decay settle instead of stopping abruptly.                                                                   |
| <b>Crossfade</b> | 0 – 15 f | Length of the crossfade written where two clips are contiguous or overlap after processing — including the steps of the gain staircase and the cuts made by the Judge inside a phrase. 3 frames is a good default. |

A fade is always kept inside the clip it belongs to: it is shortened automatically when a clip is shorter than the requested fade, so a one-frame word is never faded into silence. Set a control to 0 to write no transition of that kind.

## 10. Running a job and using the result

### 10.1 Process

With an AAF loaded and at least one track ticked, press **Process**. CompDX first asks where to save the result.




Figure 10.1 — The save dialog proposes <original name>\_CompDX in the Output AAF folder.

Confirm the name and location. The job then runs in the background: the status line reads "Processing.", the bar fills, and the elapsed time counts up. The button turns red and reads **STOP**. While the job runs, the phase panels, the Tracks panel, the Presets bar and the file controls are locked (section 3.3), so the settings cannot drift away from the ones the job is using.




Figure 10.2 — A job in progress: percentage, elapsed time and the STOP button. The tracks, phase panels, presets and file controls are dimmed and locked until the job ends.

## 10.2 Cancelling

Press **STOP** at any time. The current stage finishes its step, the partial output file is deleted and the status reads "Cancelled.". Nothing on disk is left behind.

## 10.3 Completion and failure

On success the status turns green: "Process Completed.". The output AAF is in the folder you chose. On failure the status turns red and a second line states the cause in plain language — for example that the AAF has no embedded audio, or that the output folder is not writable. Every job also writes a detailed log (section 2.7) that the **Support** button can package for Kraxton Labs.

## 10.4 Importing into Pro Tools

- 1 Open the session that contains the original dialogue tracks, or a copy of it.
- 2 **File > Import > Session Data**, select the `_CompDX.aaf`.
- 3 Import the audio tracks as **New Track**. Keep the source sample rate; choose **Copy** or **Link** to media as your workflow requires.
- 4 The processed tracks appear beneath the originals with the same names. Cuts, clip gain and fades are visible immediately; Phase 2 cuts show as shortened clips or holes; clip gain shows as the gain line on each clip.
- 5 If **Preserve removed clips** was on, the muted tracks with the material removed by the Judge (named **A1-1, A2-2, ...**) follow the processed tracks. Leave them muted while you review.
- 6 Review scene by scene against the originals. Restore or trim anything you disagree with from the preserved muted track or from the original track; all of them are sample-aligned.

## 10.5 Importing into other applications

Media Composer, Nuendo, DaVinci Resolve and Premiere Pro import the processed AAF through their normal AAF import. Clip gain is carried as the AAF clip-gain effect that these applications support; fades are standard transitions. Applications that ignore AAF clip gain will show the cuts and fades but play the clips at their original level.

## 10.6 Reading the output file

The output AAF contains a full copy of the embedded audio of the processed and unprocessed tracks alike, so it is self-contained and roughly the same size as the input. It carries the original track names, the original source timecodes of every clip, and the same composition start time, so it can be laid against the original session without offsets. When **Preserve removed clips** is on, the output also contains the muted tracks described in section 7.6, after all the original tracks; they reference the same embedded audio, so the file is not noticeably larger.

# 11. Presets




Figure 11.1 — The Presets bar with the factory presets in the drop-down and the five action buttons.

A preset stores the settings of the four phases: the two enable checkboxes, the **Preserve removed clips** checkbox and every slider. It does not store file paths, the output folder or the track selection. The last selected preset is restored at the next launch.

## 11.1 Factory presets

| Setting                        | TV Show –24dB LUFS (default) | Cinema –27dB LUFS |
|--------------------------------|------------------------------|-------------------|
| VAD Threshold                  | 0.60                         | 0.60              |
| Pre-Padding                    | 300 ms                       | 300 ms            |
| Post-Padding                   | 300 ms                       | 350 ms            |
| Minimum Pause                  | 350 ms                       | 600 ms            |
| Bleed Rejection                | On                           | On                |
| Max Lag                        | 1 ms                         | 1 ms              |
| Kill Threshold                 | 90 %                         | 85 %              |
| Preserve removed clips         | Off                          | Off               |
| LUFS Automation                | On                           | On                |
| Target LUFS                    | –24.0 LUFS                   | –27.0 LUFS        |
| Leveling Interval              | 350 ms                       | 500 ms            |
| Dynamic Range                  | 30 %                         | 45 %              |
| Fade In / Fade Out / Crossfade | 1 f / 3 f / 3 f              | 2 f / 2 f / 3 f   |

**TV Show** is tuned for episodic broadcast delivery: tight cuts, a –24 LUFS anchor, fast levelling and a moderate dynamic range. **Cinema** keeps phrases longer (600 ms minimum pause), levels over longer tiles, targets –27 LUFS and preserves more of the performance. Factory presets cannot be overwritten or deleted; save your variations under a new name.

## 11.2 Managing presets

- **+ (new)** — saves the current settings as a new preset and asks for a name.
- **Save** — overwrites the selected user preset with the current settings.
- **Import** — loads a .cdxpreset file, for example one received from a colleague or another workstation.
- **Export** — writes the current settings to a .cdxpreset file.
- **Delete** — removes the selected user preset.

Presets are stored together in `CompDX_Presets.json` in the application data folder (section 2.7); copying that file moves your whole library to another computer.

## 12. Support, diagnostics and privacy

### 12.1 The Support button




Figure 12.1 — The Export Support Report dialog explains what the report contains and how to send it.

If a job fails, or a result looks wrong, press **Support**. CompDX packages a diagnostic report as a single .zip and lets you choose where to save it. Nothing is sent automatically: open the Support page at **kraxton.com**, submit a ticket describing what happened and attach the file. The report is the only way the team can see what the engine decided and why.

### 12.2 What the report contains

- The diagnostic logs of recent jobs: every decision of the four phases with its timecode — never any audio or media content.
- System information: operating system, CPU, memory.
- The licence status (activated, trial, days remaining) — never your product key.
- Recent crash reports for Comp DX, if any.

Privacy protection is applied before packaging: project and media file names are replaced with neutral pseudonyms (Media01, Media02, ...) and user account names are removed from every path. Some technical metadata — track labels, timecodes, clip counts — remains because it is required for diagnosis. You can open the zip and read every file before sending it.

## 12.3 Logs

Each job writes one log, `CompDX_Pipeline_<date>_<time>.log`, to the Logs folder (section 2.7). Logs are rotated automatically. They are plain text and can be read with any editor; the support report includes the most recent ones.

## 12.4 Licence agreement and privacy policy

The End User Licence Agreement is shown once, at first launch, and must be accepted before the application can be used. It and the privacy policy are published on the Kraxton Labs website.

# 13. Troubleshooting and tuning

| Symptom                                                    | Cause and remedy                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| "External-media AAF" warning; every track skipped          | The AAF links to audio on disk instead of embedding it. Re-export with embedded audio (section 2.4).                                                                                                                                                                                                            |
| A track is listed but not processed                        | Stereo or multichannel tracks pass through untouched. Export the pair as two mono tracks if it must be processed.                                                                                                                                                                                               |
| Real dialogue was cut (whole clip missing)                 | First check the original clip: was it very quiet or off-mic? Lower <b>VAD Threshold</b> by 0.10. If the line coincides with another actor speaking on another track, lower <b>Kill Threshold</b> (e.g. 90 → 70 %) so the Judge demands stronger evidence. Very wide blocking may need a higher <b>Max Lag</b> . |
| I want to review or undo what the Judge removed            | Turn on <b>Preserve removed clips (New muted tracks)</b> in Phase 2 and process again. Everything the Judge removed lands on muted tracks after the processed tracks, at its original position, ready to be moved back (section 7.6).                                                                           |
| Ends of phrases are trimmed                                | Increase <b>Post-Padding</b> (e.g. 300 → 400 ms) and <b>Minimum Pause</b> so breaths stay inside the phrase.                                                                                                                                                                                                    |
| Bleed survived next to the speaking actor                  | Raise <b>Kill Threshold</b> (85 → 95 %). Check that both microphones are ticked — the Judge can only compare tracks it processes. If the actors were far apart, raise <b>Max Lag</b> to 3 – 5 ms.                                                                                                               |
| Room tone or noise kept as clips                           | Raise <b>VAD Threshold</b> (0.60 → 0.70). Phase 3's speech floor removes most of it when LUFS Automation is on.                                                                                                                                                                                                 |
| Levelling sounds "pumped" or flat                          | Increase <b>Leveling Interval</b> (350 → 800 ms) for smoother macro-dynamics, and raise <b>Dynamic Range</b> to keep more of the performance.                                                                                                                                                                   |
| A shout is still contained                                 | Raise <b>Dynamic Range</b> . Remember that the peak guard never lets a boosted peak exceed -3 dBFS.                                                                                                                                                                                                             |
| Clicks at cuts                                             | Set <b>Fade In/Fade Out</b> to at least 1 frame and <b>Crossfade</b> to 2 – 3 frames.                                                                                                                                                                                                                           |
| Job fails at once with an unreadable AAF                   | Some exports omit part of their internal index. CompDX repairs most of them; if not, re-export, or send the support report so the dialect can be added.                                                                                                                                                         |
| Pro Tools refuses or crashes on the output                 | Send the support report and, if possible, the input AAF through the support page. The output is validated before completion; a workstation-specific dialect problem can be corrected quickly with the file at hand.                                                                                             |
| "Cannot open file" on Windows                              | Fixed in Version 1.1.0 for paths with non-Latin characters. Update if you run an older version.                                                                                                                                                                                                                 |
| "Virtual Environment Detected" on a physical Windows 11 PC | Windows is running under the Hyper-V hypervisor (Core Isolation / VBS, Hyper-V, Virtual Machine Platform, Sandbox, BIOS virtualisation). Follow the steps in section 2.6. On a managed corporate PC the features cannot be disabled and CompDX cannot be activated there.                                       |
| Licence screen appears although the licence is valid       | The seat may have been released from another computer, or the trial has ended. Use <b>Claim Seat</b> or enter the key. If the computer has been offline for a long time, connect once so that the licence can be re-verified.                                                                                   |

# Appendix A — Parameter reference

| Panel   | Control                                   | Range         | Step  | Default |
|---------|-------------------------------------------|---------------|-------|---------|
| Phase 1 | VAD Threshold                             | 0.05 – 0.95   | 0.01  | 0.40    |
| Phase 1 | Pre-Padding                               | 0 – 500 ms    | 5 ms  | 200 ms  |
| Phase 1 | Post-Padding                              | 0 – 1000 ms   | 5 ms  | 400 ms  |
| Phase 1 | Minimum Pause                             | 0 – 2000 ms   | 5 ms  | 250 ms  |
| Phase 2 | Enable Bleed Rejection                    | On / Off      | —     | On      |
| Phase 2 | Max Lag                                   | 1 – 30 ms     | 1 ms  | 10 ms   |
| Phase 2 | Kill Threshold                            | 10 – 100 %    | 1 %   | 50 %    |
| Phase 2 | Preserve removed clips (New muted tracks) | On / Off      | —     | Off     |
| Phase 3 | Enable LUFS Automation                    | On / Off      | —     | On      |
| Phase 3 | Target LUFS                               | –36 – –8 LUFS | 0.5   | –24.0   |
| Phase 3 | Leveling Interval                         | 250 – 3000 ms | 10 ms | 1000 ms |
| Phase 3 | Dynamic Range                             | 0 – 100 %     | 1 %   | 0 %     |
| Phase 4 | Fade In                                   | 0 – 15 f      | 1 f   | 0 f     |
| Phase 4 | Fade Out                                  | 0 – 15 f      | 1 f   | 0 f     |
| Phase 4 | Crossfade                                 | 0 – 15 f      | 1 f   | 0 f     |

"Default" is the value of a control before any preset is applied. On a fresh installation the **TV Show –24dB LUFS** preset is loaded automatically, so the values you see on first launch are those listed in chapter 11.

## Appendix B — Files and folders

| Extension / item               | Description                                                                                                    |
|--------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>.aaf (input)</b>            | Advanced Authoring Format composition with embedded audio, exported by your workstation. Read-only for CompDX. |
| <b>&lt;name&gt;_CompDX.aaf</b> | Processed composition: edited timeline, clip gain, fades and a copy of the embedded audio.                     |
| <b>.cdxpreset</b>              | Preset file (JSON) exported from the Presets bar.                                                              |
| <b>CompDX_Presets.json</b>     | Your preset library in the application data folder.                                                            |
| <b>Comp DX.settings</b>        | Window scale, last folders and slider positions.                                                               |
| <b>CompDX_Pipeline_*.log</b>   | One diagnostic log per job, in the Logs folder; rotated automatically.                                         |
| <b>.zip</b>                    | Support report generated by the Support button.                                                                |

# Appendix C — Glossary

| Term                            | Definition                                                                                                                                                                           |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AAF</b>                      | Advanced Authoring Format: the interchange file that carries a timeline (tracks, clips, fades, gain) and, optionally, the audio itself.                                              |
| <b>Embedded / linked media</b>  | Audio stored inside the AAF versus audio referenced as external files. CompDX requires embedded media.                                                                               |
| <b>VAD</b>                      | Voice Activity Detection: the model that estimates, moment by moment, whether human speech is present.                                                                               |
| <b>Bleed (spill, crosstalk)</b> | The voice of one actor captured by another actor's microphone.                                                                                                                       |
| <b>Vertical Judge</b>           | The Phase 2 engine that compares tracks against each other over time to attribute each moment of speech to one microphone.                                                           |
| <b>TDOA</b>                     | Time difference of arrival: how much later a sound reaches one microphone than another; about 1 ms per 34 cm.                                                                        |
| <b>Preserved track</b>          | A muted track, named after its source track and number (for example A3-3), that holds the clips the Vertical Judge removed from that track when <b>Preserve removed clips</b> is on. |
| <b>Duplicate</b>                | Two tracks carrying the same digital signal (a mix and its sources, stems, copies), as opposed to two acoustic recordings of the same voice.                                         |
| <b>Tile</b>                     | The piece of a clip, of the Leveling Interval in length, that Phase 3 measures and levels independently.                                                                             |
| <b>LUFS</b>                     | Loudness Units relative to Full Scale (ITU-R BS.1770). –24 LUFS is the ATSC A/85 anchor; –23 LUFS the EBU R 128 target.                                                              |
| <b>Clip gain</b>                | Non-destructive gain attached to a clip and applied on playback by the workstation. The only way CompDX changes level.                                                               |
| <b>Peak guard</b>               | The –3 dBFS ceiling that limits any boost.                                                                                                                                           |
| <b>Speech floor</b>             | The loudness below which a detected tile is treated as room tone and removed.                                                                                                        |
| <b>Gain staircase</b>           | A sequence of small clip-gain steps written between two tiles whose gains differ markedly.                                                                                           |
| <b>Dynamic Range</b>            | The Phase 3 control that preserves a chosen share of an actor's short, obvious departures from their running loudness.                                                               |
| <b>Bridge</b>                   | Pre-Padding + Post-Padding + Minimum Pause: the longest gap treated as part of a phrase.                                                                                             |
| <b>Handle</b>                   | Short padding region written before or after a phrase (attack head, release tail) at a reduced gain.                                                                                 |
| <b>Seat</b>                     | One activation of a licence on one computer. Released and claimed from the licence chip.                                                                                             |

## Appendix D — Legal notices

---

CompDX is a product of Kraxton Labs LLC. Use of the software is governed by the End User Licence Agreement shown at first launch and published on the Kraxton Labs website.

Pro Tools, Avid and Media Composer are trademarks or registered trademarks of Avid Technology, Inc. in the United States and other countries. Adobe and Premiere Pro are trademarks of Adobe Inc. DaVinci Resolve is a trademark of Blackmagic Design Pty Ltd. Nuendo and Cubase are trademarks of Steinberg Media Technologies GmbH. macOS is a trademark of Apple Inc. Windows is a trademark of Microsoft Corporation. These names are used solely to describe compatibility; CompDX is not affiliated with, endorsed by or sponsored by any of these companies.

CompDX includes a speech-segmentation neural network derived from the pyannote.audio project (pyannote/segmentation-3.0, MIT licence, © CNRS), executed locally with ONNX Runtime (MIT licence, © Microsoft Corporation), and is built with the JUCE framework (© Raw Material Software Limited) under a commercial JUCE licence. Libraries embedded in JUCE are also used: FLAC, Ogg Vorbis, HarfBuzz, ShreenBidi, CHOC, QuickJS, libpng, zlib and software based in part on the work of the Independent JPEG Group. The full licence texts are reproduced in the file THIRD\_PARTY\_NOTICES.txt shipped with the application: inside the application bundle at `Comp DX.app/Contents/Resources/` on macOS (right-click the application, **Show Package Contents**), and in the installation folder next to `Comp DX.exe` on Windows. No audio leaves your computer.

The licensing system relies on TurboActivate / LimeLM by wyDay, LLC. Licence verification transmits an anonymised machine identifier and the licence status only; it never transmits audio, session, project or media data.

© 2026 Kraxton Labs LLC. All rights reserved. This manual may be copied and distributed together with the software; it may not be modified.