SpiritFone User Guide

Version 1.0.2

Contents

1. What SpiritFone does
2. Quick start
3. The main window
4. Global Settings (left sidebar)
5. Presets
6. Recording
7. Output Analysis
8. The plugin cards (Plugins 1 to 8)
9. Keyboard shortcuts
10. Where SpiritFone keeps its files
11. Troubleshooting
12. System requirements
13. Credits
14. Disclaimer
15. License

1. What SpiritFone does

SpiritFone is a real-time audio processor for EVP (electronic voice phenomena) listening sessions. It creates its own white noise and then passes that noise through a chain of up to eight processing stages ("plugins"). The stages shape, filter and clean the sound so that any voice-like content in it is easier to hear. You listen to the result live through headphones or speakers, and you can record it to WAV files.

The plugins run in a fixed order, top to bottom:

#	Plugin	What it does
1	Source & EQ	Creates the white-noise source and shapes its tone with a 6-point equalizer.

2	Noise filter	A spectral noise-reduction stage that removes steady background noise while trying to keep voice-like "body".
3	Time mod.	The Time Modulator. Slows down short slices of one frequency band and blends them back in.
4	Dialogue	A neural dialogue extractor (DeepFilterNet3) that pulls speech-like sound out of the noise, followed by optional tone curves and enhancers.
5	ReaEQ	Cockos ReaEQ, a hosted equalizer for extra tone shaping.
6	Vocal Lift	Helps voices stand out and turns down the quiet noise in between.
7	ReaFIR	Cockos ReaFIR, a hosted spectral gain curve for precise per-frequency cuts.
8	SPAN	Voxengo SPAN, a hosted live spectrum display. It shows the sound; it does not change it.

After Plugin 8 the sound passes through the **Master Gain**, the **Output Channels** setting (Stereo or Mono), the optional **Output Compressor**, and finally a **limiter** that stops the output from getting too loud.

Required: Cockos ReaPlugs (64-bit)

SpiritFone requires Cockos ReaPlugs (64-bit) for Plugin 7 to work properly. ReaPlugs is a free plugin package from Cockos. It provides **Plugin 7 (ReaFIR)**, which is required, and also **Plugin 5 (ReaEQ)**, which is optional. ReaPlugs is not included with SpiritFone, because Cockos' license does not allow other programs to redistribute it. **Without ReaPlugs, Plugin 7 (ReaFIR) does not work, and SpiritFone does not process audio as intended.**

How the SpiritFone installer handles it. If 64-bit ReaPlugs is not already installed, the installer shows a **ReaPlugs is required** page with the box **Download and install ReaPlugs (64-bit) from Cockos now**, ticked by default. If you leave it ticked:

1. When you click **Install**, the installer downloads Cockos' own 64-bit ReaPlugs installer (`reaplugs236_x64-install.exe`, about 1.2 MB) directly from `www.reaper.fm`. It checks that the downloaded file is exactly the expected Cockos file before using it.

2. It then opens Cockos' installer. Windows asks for permission to let it make changes, because ReaPlugs installs into `C:\Program Files`.
3. In Cockos' installer, read and accept Cockos' license, and keep its default install folder (`C:\Program Files\VSTPlugins\ReaPlugs`).
4. When Cockos' installer closes, the SpiritFone installer checks that ReaPlugs (64-bit) is now installed, and continues.

If you untick the box, cancel the download, cancel Cockos' installer, or refuse Windows' permission prompt, the installer shows a warning and SpiritFone is still installed, but Plugin 7 (ReaFIR), and the optional Plugin 5 (ReaEQ), will not work until you install ReaPlugs. To install it later, use the **Download ReaPlugs** button on the Plugin 5 or Plugin 7 card, or go to `https://www.reaper.fm/reaplugs/`, and download **ReaPlugs VST v2.36 64-bit**. **The 32-bit version does not work with SpiritFone.**

What SpiritFone does NOT do

- **It does not connect to the internet.** SpiritFone itself sends and receives nothing online. The only internet connections are these, and each happens only when you choose it:
    - **The installer** downloads Cockos' ReaPlugs installer from `www.reaper.fm`, only if ReaPlugs (64-bit) is not already installed and you leave the **Download and install ReaPlugs** box ticked (see **Required: Cockos ReaPlugs (64-bit)** above). Nothing else is downloaded or sent, and a silent (unattended) install never downloads it.
    - **The Download buttons** on the Plugin 5, 7 and 8 cards open a web page in your own web browser, so you can download those plugins yourself if needed. SpiritFone does not download anything itself.
- **It does not install or use any pre-recorded audio files.** Everything you hear is created live from SpiritFone's own white noise.

2. Quick start

1. Make sure **Cockos ReaPlugs (64-bit)** is installed. It is required for Plugin 7 (ReaFIR) (see **Required: Cockos ReaPlugs (64-bit)** in section 1). The Plugin 7 card shows green when it was found.
2. Put on headphones and turn your volume down before starting.
3. Launch SpiritFone. The installer sets it up to load the **SpiritFoneDefault** preset at startup.
4. Make sure the round power button at the top left is **green** (Plugin Chain on). Press it or **`Ctrl+Alt+P`** to turn it on and off.
5. Listen. Raise your volume slowly to a comfortable level.

6. To record, press **Start Recording** in the Audio Recorder card, then **Stop Recording** when you are done. Press **Open Internal Recordings Folder** to find the file.
7. If you change settings and like the result, type a name in **Preset Name** and press **Save Preset**.

> **Tip:** Hover the mouse over any **(?)** mark in SpiritFone for an explanation of the control next to it.

3. The main window

The window is laid out in these areas:

- **Left sidebar:** global settings, the chain order, bypass and reset buttons, shortcuts and an output status summary.
- **SpiritFone Presets card** (top): save, load and delete complete presets.
- **Audio Recorder card:** manual recording on the left, auto-batch recording on the right.
- **External Mic Recorder card:** optional microphone recording and monitoring.
- **Plugin tab row:** eight buttons, "Plugin 1" to "Plugin 8". Clicking one shows that plugin's card first, followed by the next three.
- **Output Analysis card** (left of the plugin cards): live level meters and spectrum graph.
- **Plugin cards:** four cards at a time. Each card's header shows the plugin name with an **on/off checkbox** at the top right. The header stays in place while the card's contents scroll.
- **Status line** (bottom): audio engine status, sample rate and buffer size.

Using the controls

- **Number controls** have **-** and **+** buttons for small steps. You can also click in the box and type a value.
- **Collapsing sections** (headers with a small arrow) open and close when clicked. Most are closed by default to keep the cards short.
- **Plugin on/off checkboxes** take effect immediately.

4. Global Settings (left sidebar)

Plugin Chain power button

The round button at the top of the sidebar. **Green** means the plugin chain is on; **dark red** means it is off.

Click it (or press **`Ctrl+Alt+P`**) to switch. Hovering shows the current state.

Output Channels

- **Stereo:** each ear hears a different channel. Each channel is created from its own separate noise, so a voice can appear in one ear, the other, or move between them.
- **Mono:** both channels are blended and both ears hear the same sound.

Recordings save exactly what you hear. Files are always 2-channel WAVs, but in Mono both channels are identical. Switching during a recording changes the recording from that moment on. Saved in presets.

Output Level & Compressor (closed by default)

Control	What it does
Master Gain	Final output level, applied after the last plugin and before the Output Channels setting, compressor and limiter. Default +3.5 dB. Lowering it makes the loudest syllables hit the limiter less. Saved in presets.
Plugin 4 Output Trim	The same control as Output Trim on the Plugin 4 card: a plain volume adjustment at the very end of Plugin 4. Saved in presets.
Output Compressor	Off by default. Gently pulls the loudest syllables down just before the limiter, so the limiter rarely has to make sharp level dips (which can distort strong syllables). Quiet sound below the threshold is untouched. Saved in presets.
Comp Threshold	Level (dBFS) above which the compressor starts working. Lower = more of the sound is compressed. Default -14 dB.
Comp Ratio	How firmly levels above the threshold are pulled down. 2:1 turns every 2 dB above the threshold into 1 dB. 1:1 does nothing. Default 2.5:1.
Comp gain reduction	Read-only. How many dB the compressor is currently taking off.

Exit SpiritFone

Closes SpiritFone safely, the same as closing the window. Your current settings are saved to the "last session" snapshot on the way out (see *Presets*).

Window Startup

- **Maximize on startup:** open SpiritFone maximized.
- **Remember window state:** open SpiritFone the way it was (maximized or not) when you last closed it. This takes priority over "Maximize on startup".

Current Chain Order

Lists the eight plugins in processing order. The dot on the right of each row shows whether that plugin is on.

Bypass All Plugins

Turns Plugins 1 to 8 off in one click. Their settings are kept, so you can switch each one back on from its card.

> Plugin 1 creates the noise, so with Plugin 1 off SpiritFone is silent.

Shortcuts (closed by default)

Lists the keyboard shortcuts. See section 9.

Output

A status summary: whether the plugin chain is on, whether Plugin 1 (the noise source) is on, whether Output Analysis is on, and the limiter ceiling.

5. Presets

A **complete preset** (`.sfpreset` file) stores every SpiritFone setting: all plugin settings, the Master Gain, the Output Channels setting, the compressor and the plugin on/off states. The full settings of ReaEQ and ReaFIR are stored in it too. For SPAN, which only provides a live spectral display of the sound, just its on/off state is stored.

Control	What it does
Preset Name	Type the name to save under. Choosing a preset from Saved Presets also fills this box in.
Save Preset	Saves the current settings under the name in Preset Name. If a preset with that name already exists, you are asked whether to **Overwrite** it or **Cancel**. Shortcut: **`Ctrl+Alt+S`**.

Saved Presets	Drop-down list of your presets. Start typing when it is open to filter the list.
Load Selected Preset	Loads the preset chosen in Saved Presets. Shortcut: **`Ctrl+Alt+L`**.
Open Presets Folder	Opens SpiritFone's presets folder in File Explorer.
Delete Selected Preset	Permanently deletes the chosen preset file after you confirm. This cannot be undone.
Refresh Presets	Re-reads the presets folder, for example after copying preset files in by hand.
Import Preset...	Adds a `.sfpreset` file from anywhere on your computer, such as one another SpiritFone user sent you, to Saved Presets. See *Sharing presets* below.
Export Preset...	Saves the settings SpiritFone is running right now as a `.sfpreset` file in a folder you choose, ready to send to someone. The file name starts out as the name of the loaded preset. Any changes you made after loading it are included. Your presets folder is not changed.

Control	What it does
Back Up Presets...	Copies all your presets, or just the ones you tick, to a folder you choose. See *Backing up and restoring presets* below.
Restore Presets...	Copies presets back from a backup folder. See *Backing up and restoring presets* below.
Restore Default Preset...	Puts the original **SpiritFoneDefault.sfpreset** that came with SpiritFone back into your presets folder, replacing the current one. Nothing is loaded and your other presets are not changed. See *Restoring the default preset* below.
Reset All to Defaults	**Careful:** this loads the SpiritFoneDefault preset straight away, with no confirmation, replacing **all** your current settings. Shortcut: **`Ctrl+Alt+R`**. Save a preset first if you might want your current settings back.
Load selected preset on startup	When ticked, the preset chosen in Saved Presets loads automatically each time SpiritFone starts. The name is shown next to it as "Startup Preset".
Load last settings on startup	When ticked, SpiritFone starts with exactly the settings it had when you last closed it. If both startup options are ticked, this one wins.

The "SpiritFone Last Session" preset

Every time SpiritFone closes normally, it saves all current settings to **SpiritFone Last Session.sfpreset**. This is a safety net in case you changed settings and forgot to save them.

It also works as an "undo" for the current session: if "Load last settings on startup" was ticked, choose **SpiritFone Last Session.sfpreset** in Saved Presets and press **Load Selected Preset** to go back to the settings you started with. Don't save a preset over it during the session, or that restore point is lost.

Sharing presets

A complete preset is a single small file, so you can share your settings with anyone else who has SpiritFone installed.

To send a preset:

1. Get SpiritFone sounding the way you want, for example by loading one of your presets.

2. Press **Export Preset...**, choose where to save it (the Desktop or Documents is easy to find) and, if you like, change the file name. Press **Save**.
3. Send that one file by email, a message, a USB stick or a forum post.

To use a preset someone sent you:

1. Save the file anywhere on your computer, for example in Downloads.
2. Press **Import Preset...** and choose the file. SpiritFone copies it into your presets folder and selects it in Saved Presets.
3. Press **Load Selected Preset** to use it.

> **Good to know:** importing never replaces one of your own presets. If you already have a preset with the same name, the imported copy is renamed, for example `My Session (2).sfpreset`. Importing does not load the preset or change your startup preset. A file made by a different SpiritFone version, or one that isn't a SpiritFone preset, is refused and nothing is copied.

The settings of ReaEQ and ReaFIR travel inside the preset file, so there is nothing else to send. SPAN only provides a live spectral display of the sound, so just whether it is on or off is shared, not how its display is set up.

Backing up and restoring presets

Back Up Presets... copies your presets to a folder you choose, such as a USB stick or a cloud drive folder. Choose **All presets** (every SpiritFone preset plus the ReaEQ, ReaFIR and SPAN preset folders) or **Selected presets only**. Each backup goes into its own new dated folder, so earlier backups are never overwritten. Every backup also saves your SpiritFone settings: the startup preset, the window size and position, the recording folders and the Auto-Batch values. SpiritFone never backs up on its own, so make a backup whenever you've saved presets you want to keep.

Restore Presets... copies presets back from a backup folder. You can restore all of them or just the ones you tick. A preset with the same name is replaced by the backup's copy. Presets that aren't in the backup are left alone.

Restore Presets never restores the four `.cfg` files that Back Up Presets saves with every backup. They stay exactly as they are:

- `startup_preset.cfg`: whether a preset loads when SpiritFone starts, and which one; and whether "Load last settings on startup" is ticked.
- `window_startup.cfg`: the SpiritFone window's size and position, and the "Maximize on startup" and "Remember window state" options.

- `recorder_folders.cfg`: the folders your internal and microphone recordings are saved to, and the "Create a merged file as well" option.
- `recorder_auto_batch.cfg`: the Auto-Batch Files, Minutes and Seconds values.

To restore one of these `.cfg` files from a backup, copy it by hand. See *Restoring a settings file* below.

Restoring the default preset

If you saved over SpiritFoneDefault.sfpreset, or deleted it, press **Restore Default Preset...** and then **Restore**. SpiritFone copies the original back from the **Default Preset Backup** folder, which Setup puts in the SpiritFone folder. This only puts the file back. To use it, press **Reset All to Defaults**, or choose it in Saved Presets and press **Load Selected Preset**. If the Default Preset Backup folder is missing, run Setup again.

Restoring a settings file

Each setting is kept in its own small file in the presets folder. If one gets damaged, or you want an older setting back, copy that one file back from a backup by hand:

Settings file	What it holds
`startup_preset.cfg`	"Load selected preset on startup", the startup preset, and "Load last settings on startup"
`window_startup.cfg`	The window size and position, "Maximize on startup" and "Remember window state"
`recorder_folders.cfg`	The recording folders and "Create a merged file as well"
`recorder_auto_batch.cfg`	The Auto-Batch Files, Minutes and Seconds values

1. Press **Open Presets Folder**. Leave that File Explorer window open.
2. **Close SpiritFone.** This matters: SpiritFone saves its settings when it closes, so a file copied in while it is running can be overwritten.
3. In another File Explorer window, open the backup folder that has a good copy of the file.
4. Copy that one file into the presets folder and choose **Replace the file in the destination**.
5. Start SpiritFone.

Good to know: if a settings file is deleted rather than damaged, SpiritFone simply goes back to its default for that setting and makes a new file the next time you change it.

6. Recording

Audio Recorder

Records exactly what SpiritFone plays, after all processing, to a 16-bit WAV file. If Stereo is selected as the Output Channel, SpiritFone saves the file in stereo. Likewise if Mono is selected as the Output Channel, then the saved output will be mono.

- **Start Recording / Stop Recording:** starts and stops a recording. "Recording" is shown in red while it runs.

- **Recording time / Total length:** the length of the current or last recording.
- **Open Internal Recordings Folder:** opens the folder where recordings are saved. The folder path is shown above the button. By default this is `Recorded Files\Internal Recordings` in the SpiritFone folder.
- **Change Internal Recordings Folder...:** choose a different folder. **Reset to Default** appears once you have chosen one. Folder changes are only allowed while not recording.

Each recording is named `SpiritFone_InternalOutput_<preset>_<date_time>.wav`. A matching `..._settings.sfpreset` file is saved next to it with the exact settings used, so you can reload them later: bring it in with **Import Preset...** (see *Presets*), then load it from Saved Presets.

Auto-Batch Recording

Records several files one after another automatically, for long unattended sessions.

- **Files:** how many files to record.
- **Minutes Each / Seconds Each:** the length of each file.
- Press the **Auto Record ...** button to start. Progress is shown as "file X of Y".
- **Stop Auto Batch + Save Current:** stops early and keeps the file in progress.

Auto-batch is only available while the normal recorder is stopped. Your last Files/Minutes/Seconds values are remembered.

External Mic Recorder

Records your computer's microphone at the same time as the Audio Recorder, so you can capture your own questions during a session. It starts and stops automatically with the Audio Recorder.

- **Microphone Access:** turns microphone recording on. Mic recordings are saved separately as `SpiritFone_Microphone_<date_time>.wav`, by default in `Recorded Files\Mic Recordings` in the SpiritFone folder.
- **Create a merged file as well:** when stopping, also saves a third WAV that blends the SpiritFone recording and the mic recording, named `SpiritFone_Merged_<preset>_<date_time>.wav`, in the Audio Recorder's folder. The two separate files are still saved. Requires Microphone Access, and must be ticked *before* recording starts.
- **Monitor the Mic (headphones only!):** plays the microphone live to your output device, whether or not you are recording.
- **Open Mic Recorder Folder / Change Mic Recorder Folder...:** same as for the Audio Recorder.

Use headphones for mic monitoring. On speakers the microphone hears SpiritFone and feeds it back, causing loud feedback. On speakers, a merged file will also contain SpiritFone's sound twice (the second copy slightly later, like a small room echo).

The Microphone Access, merged-file and folder settings can't be changed while a mic recording is running.

7. Output Analysis

Shows what SpiritFone is actually sending out. It only reads the sound and never changes it. Use the checkbox at the top right of the card to turn it on or off.

- **RMS / Peak:** average and peak output level in dBFS (0 dBFS is the maximum).
- **Spectrum graph:** the live frequency content of the output, low frequencies on the left.

Analyzer

- **Freeze Display:** holds the graph still so you can study it.
- **Clear Analysis:** resets the graph and meters.
- **Display Smoothing:** higher values make the graph move more slowly and smoothly.
- **Visual Gain:** moves the graph up or down on screen only. It does not change the sound.

Reference Overlay

Draws a saved reference curve over the live graph so you can compare against a known good sound.

- **Show Reference Overlay:** turns the reference curve on.
- **Reference Curve:** which reference to show.
- **Auto-align Reference / Align Band / Align Reference Now:** line the reference curve up with the live graph at the chosen frequency band, automatically or once on demand.
- **Show Band Balance Table:** shows a table comparing the live sound and the reference, band by band.
- **Reference Offset:** moves the reference curve up or down by hand.

Spectrogram

The **Show Spectrogram** button, next to the Output Analysis label, opens a separate window that draws the sound as a picture: time runs across the picture, pitch runs up the picture, and the colour shows how loud each pitch is. It only reads the sound and never changes it. **View** chooses **Fast only**, **Fine only** or **Dual** (both,

stacked). A Fast picture uses a short analysis window, so it follows quick changes in time. A Fine picture uses a long window, so it shows pitch more exactly but blurs the timing. **Source** chooses what is shown: **Final output** (what leaves SpiritFone), **Before voice stage** (the sound going into Vocal Lift), or **Both** side by side with the same settings (left: before the voice stage, right: final output), so you can see what Vocal Lift did. Other controls set the window sizes, **Overlap** (50%, 75% or 87.5%), **Window shape**, **Time span** (5, 10 or 20 seconds), a linear or logarithmic frequency axis with **Low Hz** and **High Hz**, the **Floor dB** and **Ceiling dB** of the colours, and the **Colours** (Heat, Viridis-like, Grey or Rainbow). Move the mouse over a picture to read the time, frequency and level at that point. **Pause** holds the pictures still. **Snapshot** (or the S key) saves the pictures as a PNG image in the **SpectrogramSnapshots** folder next to SpiritFone.exe, with a caption that gives the date and time, the voice stage, the last preset loaded or saved, and the settings used; **Open folder** then shows that folder. The window can be resized, and SpiritFone remembers its settings, but they are not saved in presets.

8. The plugin cards

Every card has an on/off checkbox in its header and (except the display-only ones) a **Reset Plugin N Defaults** button at the bottom that resets that one plugin. Controls with a (?) mark have more detail when you hover over them.

Plugin 1: Source & EQ

Creates the white-noise source that everything else works on, then shapes its tone with an equalizer. **If Plugin 1 is off, SpiritFone is silent.**

Control	What it does
Main WET	Level of the generated noise, in dB.
EQ Dry/Wet	How much of the equalized noise is used (100% = fully equalized, 0% = plain noise).
EQ Shift	Shifts the whole EQ curve lower or higher in frequency (-100 to +100).
Active EQ Points 1-6	Click a number to select that EQ point for editing (green = selected). Selected points show their settings below.
Type	Peak, Low Shelf, High Shelf, Notch, Bandpass or Allpass.
Frequency / Gain / Q	The point's centre frequency, boost or cut in dB, and width (Q). Notch, Bandpass and Allpass have no gain; Q sets their width.

Plugin 2: Noise filter

A spectral noise-reduction stage that removes steady background noise while keeping the "body" of voice-like sound. The status table at the top shows the Mode (Stereo or Mid/Side), whether Emphasis is on, and whether Solo Noise is on.

Section / control	What it does
Main Reduction	
Reduction Amount	How strongly noise is reduced, in dB.
Maximum Attenuation	The most any frequency can be turned down.
Soft Knee	How gradually reduction starts. Higher = smoother, gentler transitions.
Reaction Time	How fast the filter reacts to changes, in milliseconds.
Adaptation Time	How fast the filter's picture of the background noise updates, in seconds.
Dry Remainder	A small amount of the unfiltered sound mixed back in (0-5%), which can make the result sound more natural.
Emphasis EQ	
Emphasis Enabled	Turns the four-band emphasis EQ on.
Low Shelf / Low Mid / High Mid / High Shelf	Frequency, gain and slope or bandwidth for each band. Emphasis tells the filter which frequencies to favour.
Routing	
Mid/Side Mode	Processes the centre and side parts of the stereo sound separately instead of left and right.
Solo Noise	Plays only what is being removed. Useful for checking the filter isn't taking away too much. Turn it off for normal listening.

Advanced sections (Broad-Boost Counter, Gate Bias Curve, Spectral Floor Search, Emphasis Boost Headroom) control the filter's inner workings directly. Their defaults are carefully calibrated, so change them

with care. Each control has a (?) explanation.

- **Restore Calibrated Gate Tuning** resets only those advanced sections.
- **Reset Plugin 2 Main Controls** resets the main controls only (and turns Noise Reduction and Emphasis off), leaving the advanced sections alone.

Plugin 3: Time mod.

The Time Modulator takes short slices of the sound in one frequency band, plays them back slowed down, and blends them with the live sound. The status table shows the band, the Mix, the loop settings and whether Solo Band is on.

Control	What it does
Tempo lock (checkbox)	Locks the Tempo so it can't be changed by accident.
Tempo	The tempo in BPM that slice lengths are based on.
Loop Length	Length of each slice, as a note value at the chosen tempo (1/16 up to 8 Bars).
Mode	How much each slice is slowed down: 1.5x, 2x or 4x.
Fade In / Fade Out	How each slice fades in and out, from Instant to 16 Bars (automatically kept shorter than the slice).
Smooth	Smooths the level changes between slices, in milliseconds.
Band Low / Band High	The frequency range that is processed.
Mix	How much of the processed band is blended in (0-100%).
Solo Band	Plays only the processed band. Turn it off for normal listening.

Plugin 4: Dialogue

A neural dialogue extractor. It uses the **DeepFilterNet3** model to pull speech-like sound out of the noise. The **Status** line shows whether the engine is ready. The status table shows the pitch range and whether the Pitch and Rhythm enhancers are on.

Section / control	What it does
Live Focus	
Harmonic Boost Strength	How strongly voice harmonics are reinforced.
Pitch Range Low / High	The voice pitch range (in Hz) that the enhancers focus on.
Smoothing	How smoothly the enhancers change over time, in milliseconds.
Output Trim	Volume adjustment at the very end of Plugin 4 (same control as in the sidebar).
Rhythm Boost	Strength of the Rhythm Enhancer.
Reference Curve	
Reference Curve Enabled	Master switch for this whole section. When off, Plugin 4 outputs the engine's raw result. Ticking it also turns on the three EQ bands.
Sensitivity	Blend of the engine's output with the original input. 100% = all engine output (default). Lower values bring back more of the original: less noise removal, more noise.
Maximum Attenuation	The most any frequency can be reduced. 0 dB = no noise reduction; -60 dB is the default; -100 dB or lower = no limit.
Low Shelf / Peak / High Shelf	An output EQ after the engine, each band with its own Enabled box, frequency, Emphasis (boost or cut) and slope or bandwidth.
Routing	
Pitch Enhancer Enabled	Reinforces sound in the voice pitch range.
Rhythm Enhancer Enabled	Reinforces speech-like rhythm.

Solo Enhanced Pitch Only	Plays only the pitch-enhanced part. Turn it off for normal listening.
Solo Noise	Plays only what Plugin 4 removes. Turn it off for normal listening.

Plugins 5, 7 and 8: hosted plugins

Plugins 5, 7 and 8 are third-party plugins that run inside SpiritFone. Each card has the same basic layout:

- **Plugin:** whether the plugin was found and loaded (green = ready).
- **Editor:** whether the plugin's own window is open.
- **Show Editor / Hide Editor:** opens or closes the plugin's own window, where its real controls are.
- **Last loaded preset:** the name of the last preset file loaded into that plugin this session. For ReaEQ and ReaFIR this is exact; for SPAN it is SpiritFone's best guess, and can be out of date if you then change settings by hand.

Plugin 5: ReaEQ

Cockos ReaEQ, an equalizer. Shape the tone in ReaEQ's own editor (Show Editor).

> **Optional:** ReaEQ is not required for SpiritFone to work properly. It comes from **Cockos ReaPlugs (64-bit)**, the same package as the required Plugin 7 (ReaFIR), so once ReaPlugs is installed ReaEQ is available too. ReaPlugs is not included with SpiritFone. See **Required: Cockos ReaPlugs (64-bit)** in section 1.

- **Load Preset (.fxp)... / Save Preset (.fxp)...:** load or save ReaEQ settings as `.fxp` files. The dialog opens in SpiritFone's own ReaEQ presets folder.
- **Download ReaPlugs:** opens the Cockos ReaPlugs download page (reaper.fm/reaplugs) in your web browser. Choose **ReaPlugs VST v2.36 64-bit**; the 32-bit version does not work with SpiritFone.
- **Plugin File Location / Browse for reaeq-standalone.dll...:** shows where ReaEQ was found (**Found at:**), or where SpiritFone expected it if it was not found. SpiritFone finds ReaPlugs in the folder Cockos' installer used, even if that is not the default folder. **Browse** lets you point SpiritFone at a different copy; **Reset to Default Location** undoes that.

Plugin 6: Vocal Lift

Vocal Lift helps you hear voices more clearly. It is built into SpiritFone. It splits the sound into separate pitch bands and treats each band on its own. It lifts faint sound toward a target level (**Boost**), and it turns a band down while that band is quiet, so that the static between words is reduced (**Static removal**). A band's gate

opens when a word starts, stays open briefly after the word so the ends of words are not chopped, and then closes gently. **Hysteresis** stops a gate from opening and shutting again and again around its opening point. **Tone and output** has a low cut, a gentle bass and treble tilt, one adjustable peak, an output level, a limiter ceiling and a Mix control (1 is all processed sound, 0 is the untouched sound). Your settings are saved in presets.

Show Vocal Lift Editor opens Vocal Lift in its own window. At the top is a table with one row for each pitch band, with live readouts, and below it are dials for the other settings. **Number of bands** sets how many bands the sound is split into: more bands give a finer gate. **Solo band** lets you listen to just one band to hear what it contains. **Reset all settings** puts every Vocal Lift setting back to its starting value. A listening sweep lets you hear only the frequencies inside a green bracket, either after Vocal Lift's processing or before it, so you can hear what Vocal Lift does to that slice. Hover over any dial to see what it does. **Hide Vocal Lift Editor** (or the window's close button) closes the window.

Two more processors from the LSP plug-in collection can follow Vocal Lift inside this card. They work differently from Vocal Lift, which learns how loud the static is in each band and opens or closes by itself. The **LSP Multiband Expander** is a plain expander: for each band you set a threshold yourself and choose a mode. In **Down** mode the band is turned down whenever it is quieter than that threshold. In **Up** mode the band is turned up whenever it is louder than that threshold. It does not learn anything. Use Down mode for a second, fixed pass at quiet sound that Vocal Lift left in. The sections after this one explain how to set up its bands. The **LSP Multiband Compressor** does the opposite job: for each band you set a threshold, and the band is turned down whenever it is louder than that threshold. Vocal Lift has nothing like this, apart from its output limiter, which only acts on peaks. The thresholds, amounts and speeds are set in each processor's own window, which the **Show LSP Expander** and **Show LSP Compressor** buttons open. **LSP Expander On** and **LSP Compressor On** switch each one on or off. Each has a button for its own window: **Show LSP Expander** / **Hide LSP Expander** and **Show LSP Compressor** / **Hide LSP Compressor**. **Compressor Before Expander** chooses which of the two comes first. These settings are saved in presets as well.

The two processors are in the file **lsp-plugins.dll**, which the installer puts in the same folder as SpiritFone.exe. They come from LSP Plugins (Linux Studio Plugins Project), version 1.2.35, which is free software under the GNU Lesser General Public License version 3. The license texts are in the **redistributables\LSP-Plugins** folder inside the SpiritFone folder. The complete source code is at https://github.com/pkshadygrove/SpiritFone in the **LSP-Plugins** folder. If lsp-plugins.dll is missing, the two LSP buttons cannot start and the card shows a problem message under them.

LSP Multiband Expander: setting up the bands

The LSP Multiband Expander splits the sound into up to 8 frequency bands. Each band has its own expander settings, so you can treat the low part of a voice, the middle and the highs differently. Press **Show LSP Expander** to open its window.

How the bands are laid out. Each band has one frequency setting, **Range**, which is where that band **starts**. A band ends where the next band above it starts. Band 1 has no Range knob: it always starts at 0 Hz. The two numbers under each Range knob are that band's start and end. The end of the top band is shown as 22050, which is half the 44.1 kHz sample rate. A band whose **Band button** is off (the buttons marked #2 to #8 at the top of each column) is skipped, and the band below it simply stretches up to the next band that is on.

The one rule when you set the ranges. When several bands are on and you change a band's Range, the editor moves the other bands to keep them in order. If the bands start out at the same frequency, they can all jump to the value you just set. Follow the steps below so that does not happen.

Example: 5 bands that start at 150, 500, 1500 and 4000 Hz

1. Click the Band buttons **#3 to #8** so that they are no longer lit. Leave band 2 on.
2. Set band 2's Range to **150**. The numbers for band 1 should now read 0 and 150.
3. Band 3: while its Band button is **off**, set its Range to **500**. Then click its Band button to turn it on.
4. Band 4: while its Band button is off, set its Range to **1500**, then turn it on.
5. Band 5: while its Band button is off, set its Range to **4000**, then turn it on.
6. Leave bands 6, 7 and 8 off.
7. Check the numbers under the Range knobs. From band 1 to band 5 they should read: 0 and 150; 150 and 500; 500 and 1500; 1500 and 4000; 4000 and 22050. If a number is slightly off, such as 149.9 or 500.5, or two bands show the same start, a band was pushed. Turn the bands off again and repeat the step where it went wrong.

Changing a range later. You can turn a band's Range knob freely anywhere between its neighbours' values. For example, band 3 can go anywhere above band 2's start and below band 4's start. When you move band 3's start, the end number of band 2 moves with it. That is the same boundary seen from both sides, and it is correct. Do not drag a band **past** a neighbour while the bands are on. To move a band past a neighbour, turn its Band button off, set the Range, and turn it on again.

Each band has two on/off buttons, and they do different jobs. The **Band button** (#1 to #8, at the top of the column) decides whether the band exists. The **ON button** (in the Controls row) decides whether the band's expander works. A band's expander only works when both are on.

LSP Multiband Expander: what each band's controls do

These controls are in the strip along the bottom of the window.

Control	What it does
Band button (#2 to #8)	Decides whether the band exists. With it off, the band is skipped and the band below it stretches up to the next band that is on. Band 1's button cannot be turned off.
ON	Decides whether the band's expander works. With ON off, the band keeps its frequency range, but its sound passes through unchanged. If the band's Band button is off, its ON button makes no difference.
SOLO / MUTE	For listening only. SOLO turns every other band down by 36 dB, so you hear just that band. MUTE turns that band down by 36 dB. Turn them off when you are finished listening.
Makeup	Gain added to the band's output after the expander has acted, in dB.
Ratio	How strong the expansion is. A ratio of 1:1 does nothing. Higher values act more strongly.
Knee	How gradual the change is around the Attack Level. A larger value gives a softer transition.
Attack Level (dB)	The level at which the band starts to act.
Attack time (ms)	How quickly the band reacts when the sound crosses the Attack Level.
Release Level (dB)	A second level, set relative to the Attack Level, that controls when the release starts. The starting value shows -inf.
Release time (ms)	How slowly the band lets go.
Hold (ms)	How long the band keeps its action before the release begins.
Expander Mode (Up or Down)	**Down**: sound quieter than the Attack Level is turned down, and the further below the level it is, the more it is turned down. Sound louder than the level is left alone. **Up**: sound louder than the Attack Level is turned up, and sound quieter than the level is left alone.

LSP Multiband Expander: what each band measures

The sidechain is the copy of the sound that each band measures to decide when to act. These controls change what the band measures. They do not change the sound you hear.

Control	What it does
Sidechain Preamp	Boosts the copy of the sound that the band measures.
Sidechain Mode	How the band measures the level of the sound. **Peak** follows instant peaks. **RMS** uses an average, which is smoother. **LPF** and **SMA** are two other kinds of smoothing.
Sidechain Reactivity (ms)	How fast the measurement follows the sound.
Sidechain Lo-cut / Hi-cut	An extra filter on the sound the band measures. Each has an on/off button and a frequency knob.

Mode. The **Mode** box at the bottom left of the window chooses how the bands are cut apart: **Classic**, **Modern** or **Linear Phase**. Linear Phase adds a small delay.

> **Tip:** The LSP window has its own preset menu at the top, separate from SpiritFone's presets. SpiritFone also saves the expander's current settings in its own presets.

Plugin 7: ReaFIR

> **REQUIRED: SpiritFone requires ReaFIR to function properly.** ReaFIR is part of **Cockos ReaPlugs (64-bit)**. ReaPlugs is not included with SpiritFone and must be installed separately (the SpiritFone installer offers to do this for you). Without it, Plugin 7 does not work and SpiritFone does not process audio as intended. See **Required: Cockos ReaPlugs (64-bit)** in section 1.

Cockos ReaFIR, a spectral gain curve for precise per-frequency cuts. It works the same way as the Plugin 5 card: Show Editor, Load/Save Preset (.fxp), Download ReaPlugs (choose the **64-bit** version), Plugin File Location, and Browse for reafir_standalone.dll.

> **How SpiritFone uses ReaFIR:** ReaFIR runs in **Subtract** mode. In this mode the curve you draw is treated as a noise profile and subtracted from the sound, so **lowering** a point **keeps more** of that frequency band, and raising it removes more.

Plugin 8: SPAN

Voxengo SPAN, a detailed live spectrum display. It only shows the sound and does not change it. Open it with Show Editor. It is included with SpiritFone; **Download SPAN** and **Browse for SPAN.dll...** are there if you want to use your own copy.

9. Keyboard shortcuts

All shortcuts use **Ctrl+Alt**. They don't work while you are typing in a text or number box.

Keys	Action
`Ctrl+Alt+P`	Plugin Chain on/off
`Ctrl+Alt+1 ... 8`	Plugin 1 ... Plugin 8 on/off
`Ctrl+Alt+S`	Save Preset (using the name in Preset Name)
`Ctrl+Alt+L`	Load the preset selected in Saved Presets
`Ctrl+Alt+R`	Reset **all** settings to defaults (no confirmation!)

10. Where SpiritFone keeps its files

SpiritFone is installed for your Windows user only, by default in:

`C:\Users\<you>\AppData\Local\Programs\SpiritFone`

Installing again. Setup always asks where to install. If the folder you choose already contains SpiritFone, Setup tells you and asks what to do. Click **Yes** to install over it: Setup replaces SpiritFone's program files in that folder, keeps the presets and recordings you made, and replaces the starter preset **SpiritFoneDefault.sfpreset**, and the original copy in **Default Preset Backup**, with the ones that come with this version. Click **No** to go back and choose a different folder (Browse, then Make New Folder, or type a new folder name). Each folder is a separate copy of SpiritFone with its own presets. To remove a copy, run `unins000.exe` in its folder.

What	Where (inside the SpiritFone folder)
Complete presets (`.sfpreset`), including SpiritFone Last Session	`presets\`
Hosted-plugin presets (ReaEQ, ReaFIR, SPAN)	Their own folders inside `presets\`
Original copy of SpiritFoneDefault.sfpreset (used by Restore Default Preset...)	`Default Preset Backup\`
Hosted-plugin saved state	`plugin_state\`
Audio recordings (and merged recordings)	`Recorded Files\Internal Recordings\` (you can change this in the Audio Recorder card)
Microphone recordings	`Recorded Files\Mic Recordings\` (you can change this in the External Mic Recorder card)

A fresh install includes a starter preset for each of these plugins, made from SpiritFone's default settings: **ReaEQDefault.fxp** and **ReaFIRDefault.fxp**. Load one to see or return to the settings SpiritFone starts with.

The easiest way to find a recording is the **Open ... Folder** buttons in the recorder cards.

11. Troubleshooting

Problem	Things to check
No sound at all	Is the power button green? Is Plugin 1 on (sidebar "Source" line)? Is Windows sending sound to the device you are listening on? Is the volume up?
Sound is harsh or distorted on loud syllables	Lower the **Master Gain**, or turn on the **Output Compressor** (sidebar, Output Level & Compressor).
Much louder or noisier than expected	Check that no **Solo** option is on (Plugin 2 Solo Noise, Plugin 3 Solo Band, Plugin 4 Solo options).
Plugin 5 (ReaEQ) or Plugin 7 (ReaFIR) shows red / not loaded	Cockos ReaPlugs (64-bit) is probably not installed. It is required for Plugin 7 (ReaFIR); Plugin 5 (ReaEQ) is optional. Press **Download ReaPlugs** on the card and install **ReaPlugs VST v2.36 64-bit** (not the 32-bit version). SpiritFone keeps checking for it, so the card turns green once it is installed. If it is installed in an unusual place, use **Browse for ...dll**.
Plugin 8 shows red / not loaded	Use the card's **Browse for ...dll** button to point SpiritFone at the plugin file, or **Download** to install it.
Plugin 4 status is red	The DeepFilterNet engine files are missing. Reinstall SpiritFone.
Loud whistling or feedback	Turn off **Monitor the Mic**, or switch to headphones.
Changed too many settings and want to go back	Load your last saved preset, or **SpiritFone Last Session.sfpreset** (see section 5).
Still need help?	Send your request to the author, Paul Kelly: webmaster@spiritit c.org

12. System requirements

	Minimum	Recommended
Operating system	Windows 10 or Windows 11, 64-bit	Windows 11, 64-bit
Memory (RAM)	4 GB	8 GB or more
Processor	2 or more cores	A modern processor running at about 2.3 GHz or faster
Disk space	About 100 MB free	Extra space for recordings (see below)
Audio	A sound output device (headphones or speakers)	Headphones
Cockos ReaPlugs	**Required** for Plugin 7 (ReaFIR): ReaPlugs VST v2.36 **64-bit** (free; the installer offers to download it, see section 1)	
Internet	Only needed if the installer downloads ReaPlugs	
Microphone	Not required	Needed only for the External Mic Recorder
PDF viewer	Needed to open this Help file from the Help button	

Recordings and disk space. Each recording uses about 10 MB per minute (about 635 MB per hour). A merged recording and a mic recording each use about that much again.

Recordings and memory. SpiritFone saves a recording to disk while it records, so recording needs only about 8 MB of extra memory however long the recording is. A recording in progress is a file with "InProgress" in its name, for example `SpiritFone_InternalOutput_InProgress_<date_time>.wav`. If SpiritFone or the computer stops unexpectedly, that file keeps everything up to the last few seconds and still plays. When you press Stop, the file is renamed to its normal name.

Longest single recording. A WAV file can't be larger than 4 GB, which is about 6.7 hours of recording. If a recording reaches that size, SpiritFone saves what fits and shows a warning. Use Auto-Batch Recording to

split longer sessions into separate files.

The installer checks the memory and processor-core minimums and won't install if they aren't met. It doesn't check processor speed, because Windows doesn't report it reliably.

13. Credits

This work is based upon the foundational work of the following pioneers in the field of ITC research: Konstantin Raudive, Friedrich Jürgenson, Anabela Cardoso and others. A limitless amount of credit must also be given to those who followed those pioneers; Ron & Lourdes Yacovetti and Tony & Cherie Rathman, for without them, this project would not exist. Thank you.

- **DeepFilterNet3**, the dialogue extraction model used by Plugin 4.
- **ReaEQ and ReaFIR** (ReaPlugs) by Cockos. ReaPlugs is not included; it is installed separately from Cockos. ReaFIR is required; ReaEQ is optional.
- **SPAN** by Voxengo, redistributed with permission.
- **LSP Plugins** (Multiband Expander and Multiband Compressor, used by the Vocal Lift card) by the Linux Studio Plugins Project, version 1.2.35, free software under the GNU Lesser General Public License version 3.

The full credits are in `CREDITS.txt` in the SpiritFone folder.

14. Disclaimer

For entertainment purposes only. SpiritFone is provided for entertainment and personal experimentation. It is not a scientific instrument, and nothing it produces should be treated as evidence or proof of anything.

No guarantee of spirit voices. SpiritFone generates its own white noise and processes it with filters, equalizers and a speech-enhancement model. These tools are designed to bring out voice-like sounds, and the human brain naturally tends to hear words and voices in random sound (an effect known as pareidolia). Any voices, words or messages you hear or record may come from the noise itself, the processing, or your own interpretation. We make no claim or guarantee that SpiritFone detects, receives or reproduces the voices of spirits, the deceased, or any other paranormal source.

No warranty. SpiritFone is provided "as is", without warranty of any kind. You use it at your own risk, and the author is not responsible for any loss, damage or decision arising from its use or from how you interpret anything you hear.

Protect your hearing. Keep your volume at a comfortable level, especially with headphones. Processed noise can produce sudden loud sounds.

15. License

SpiritFone is free, for personal, non-commercial use only. Nobody may sell it, or charge money for it or for access to it. It is available only from its official source, `https://github.com/pkshadygrove/SpiritFone`, and may not be copied, shared, posted or uploaded for other people. Its program files may not be modified or taken apart, except as the law allows and as the licenses of the third-party software allow for that software.

The software made by other people that SpiritFone includes or works with (LSP Plugins, Voxengo SPAN, DeepFilterNet, and Cockos ReaEQ and ReaFIR) is not covered by the SpiritFone license. It stays under its own existing terms and conditions. See section 13, Credits.

The complete license is in `SpiritFone License.txt` in the SpiritFone folder. Setup also shows it before installing.