

SEG-D MONITOR

Gunlink SEG-D 1.0 Source-Signature Quality Control Suite

USER MANUAL

Windows • Linux • macOS

Revision 2.26.3 — for application version 2.26.3

Built 2026-10-11

Developed by Vladimir Lesnichiy · 2026

Contents

1. Introduction	5
1.1 What the application does	5
1.2 Free features and licensed features	5
1.3 What changed since revision 1.0	5
1.3.1 What changed since revision 2.4.0	6
1.4 About this manual	7
2. Installation	8
2.1 System requirements	8
2.2 Installing on Windows	8
2.3 Installing on Linux	8
2.4 Installing on macOS	8
2.5 Upgrading to a new version	9
3. First Run and Licensing	10
3.1 What happens on first start	10
3.2 Obtaining a license	10
3.3 How the license works	11
3.4 Where the license is stored	11
3.5 Renewals and replacements	12
4. Projects and the Project Database	13
4.1 The project concept	13
4.2 Creating, opening and switching projects	13
4.3 What lives where	15
4.4 Projects on a network share	15
5. The Main Window	16
5.1 Layout	16
5.2 The toolbar groups	16
5.3 Collapsible groups	19
5.4 The SEG-D folder panel	19
5.5 The similarity heatmap	19
5.5.1 Watch your neighbours	20
5.6 The log panel and status labels	22
5.7 Themes	23
5.8 Window behaviour	23
5.9 Layouts and the command palette	23
6. Configuring the Source Array	25
6.1 The model, and how it reaches the app	25
6.2 Building an array in the Designer	26
6.3 The Source Configuration tab	27
6.4 Nav masks	28
6.5 Worked example: 3 sources × 4 sub-arrays × 6 guns	29
6.6 Channel Sets	30
6.7 Shot Points and the spec check	31
6.8 Carrying a configuration to another survey	32
6.9 Far Field	32
7. Monitoring a Line	36

7.1 Preparation checklist	36
7.2 Generating a reference	36
7.3 What happens during monitoring	38
7.4 Reading the log	38
7.5 Inspecting headers	39
7.6 Stopping	40
7.7 Light Mode and the performance settings	40
7.8 Automatic mode	40
8. Live Trend Graphs	43
8.1 What the graphs show	43
8.2 Using the panel	43
8.3 The gun × SP heatmap	44
8.4 The far-field page	44
9. Trace Stacks	46
9.1 The Stacks window	46
9.2 By Sub-arrays	46
9.3 Detaching, zooming and fitting	47
10. Viewers	48
10.1 Data Viewer	48
10.2 Database Viewer	48
10.3 Seismic Viewer	50
10.4 SP Quality	51
10.5 The Refresh badge	52
10.6 Manage Lines	52
11. Bubble Test (BBT)	54
11.1 Selecting files	54
11.2 Parameters	54
11.3 Picking and manual correction	55
11.4 The computation	56
11.5 Saving, comments and saved tests	56
11.6 The BBT report	56
11.7 Choosing k	56
12. Reports	58
12.1 The End-of-Line report	58
12.2 Generating a report for any stored line	62
12.3 Comments	62
12.4 Branding	63
12.5 Missing shots and unreadable files	63
12.6 How reports are rendered	64
12.7 The far-field pages	64
12.8 The hydrophone QC pages	67
13. Session Logs and Diagnostics	68
13.1 The session log file	68
13.2 Reporting a problem	68
14. Backup, Transfer and Multi-Workstation Use	69
15. Troubleshooting and FAQ	70
15.1 Licensing	70

15.2 Monitoring and ingest70

15.3 Reports.....70

15.4 Displays and viewers.....71

Appendix A. Folder and File Reference72

Appendix B. Bubble-Test Formula Summary73

Appendix C. Spec-Check Rules74

Appendix D. Keyboard and Mouse Reference.....75

Appendix E. Glossary76

Appendix F. Far-Field Signature Mathematics77

1. Introduction

1.1 What the application does

SEG-D Monitor is a quality-control application for marine seismic source acquisition. It watches a folder of incoming Gunlink SEG-D format 1.0 near-field hydrophone (NFH) files in real time, decodes every shot as it arrives, and evaluates source-signature similarity by autocorrelation against a reference. Around that core it provides live trend graphs, trace-stack displays, a full project database, bubble-test processing with PDF reporting, shot-point quality analysis, and automatic end-of-line reports.

The application is distributed as a compiled, self-contained package for Windows, Linux and macOS. No Python installation, interpreter, or third-party library is required on the target machine: everything is inside the application folder.

1.2 Free features and licensed features

The core monitor is free forever. Everything needed to watch a line in real time works without any license: folder watching, SEG-D decoding, autocorrelation similarity, the heatmap, live graphs, header inspection, the Source/Gun Map Designer, session logging, and the project database.

MODULE	WHAT IT UNLOCKS
Data / Database Viewers	The  Data and  Database buttons — the filterable table of stored shots, and the gun performance / line replay / sources analysis tabs.
Seismic NPZ Viewer	The  Seismic button — wiggle display of recorded shot NPZ files and trace stacks.
Bubble Test (BBT) Window	The  BBT button — complete bubble-test processing, evaluation, saved tests and its PDF report.
End-of-Line PDF Report	The  Line Report and  SP Quality buttons. Both are covered by the same license feature.
Trace Stacks Viewer	The  Stacks button, including the  By Sub-arrays window.

Locked buttons carry a padlock () and their tooltip names the current license status. Without a license the monitor is limited to **one source**; everything else about it is unrestricted and does not expire.

The Source/Gun Map Designer is free. It was a licensed module in revision 1.0 of this manual and was released as a free feature in application version 1.9.3. Older license files that still list it are accepted — the obsolete entry is simply ignored.

1.3 What changed since revision 1.0

Revision 1.0 of this manual covered application version 1.7.x. If you are returning to the manual after using an older release, these are the additions and behavioural changes that matter most.

AREA	WHAT IS NEW	CHAPTER
Main window	The toolbar is now seven groups, four of which collapse to chips that pop open on demand. The SEG-D folder panel collapses automatically when monitoring starts. An interface theme selector was added.	5
Shot-point quality	 SP Quality — contractual bad-shot acceptance checking, live during a line and as a verdict page in the report.	10, 12

Session logs	Every session now writes a log file into the project, with  Logs and  Save Log buttons.	13
Reports	Redesigned End-of-Line and BBT reports, a Comments box with offline spell-check, and up to three customer logos via  Branding...	12
Missing shots	Missing shot points are detected and reported, and are kept strictly distinct from files the application could not read.	12
Ingest reliability	A reconciliation sweep re-scans the watched folder every 30 s, so a stalled network drive delays shots instead of losing them.	7
Sub-array stacks	Panels can be detached into their own windows without pausing the rest, and each panel keeps its own zoom.	9
Platforms	macOS builds are now provided alongside Windows and Linux.	2
Help	A built-in manual on  Help or F1 .	—

1.3.1 What changed since revision 2.4.0

Revision 2.4.0 of this manual covered application version 2.4.0. These are the additions since.

AREA	WHAT IS NEW	SECTION
Far-field signature	Each shot's far-field signature is reconstructed from the near-field hydrophones and followed along the line, with a check against the source's model. A live graph page, two report pages and its own settings tab. Off by default.	6.9 , 8.4 , 12.7 , F
Gun × SP heatmap	Right-click a depth, pressure or timing graph for a per-gun, per-shot map of deviation from tolerance — a drifting gun reads as a stripe.	8.3
Light Mode	A cross-correlation-only mode for lines arriving faster than the app can keep up with, plus two performance settings. Stop now reports how long each phase took.	7.7
Manage Lines	Every line in the project with its counts, and a permanent delete that also removes the line's NPZ files.	10.6
Source Designer	Irregular layouts — arrays with different numbers of sub-arrays — are fully supported, and the nav mask the panel SHOWS is the one the design applies.	6.2
Spec check	A gun switched off is no longer a bad shot on its own; the percentage rule is now reported live and can recover; a missing shot is charged to the one source whose turn it was; and a shot breaking several rules is counted once.	C
Sources and masks	A sub-array switched off no longer discards the shot, and every window names the source rather than a raw nav mask.	6.4
Automatic mode	Lines start, stop and report themselves: a line is stopped after a quiet time, its End-of-Line report is made without questions, and the next line folder is started when it appears. Comments are added afterwards, also while another line is being monitored, and there is one report per line.	7.8 , 12.2 , 12.3
Watch your neighbours	When a hydrophone stays at or below the colour-scale minimum, the hydrophones next to it are checked: if they changed too, an alarm outlines the area on the similarity heatmap, sounds AIRLEAK in Morse and steams until it is answered; if not, a note marks one suspect hydrophone. Can be switched while monitoring.	5.5.1
Free license online	The SEG-D Monitor website issues a free license with every module, for three months at a time, straight from the License dialog's Get a free licence online button.	3.2

1.4 About this manual

Chapters 2–4 get the application installed, a project created, and a license activated. Chapter 5 describes the main window. Chapter 6 covers configuring the source array — the single most important setup step, and the one that most affects whether everything downstream is correct. Chapters 7–11 follow a survey in the order it happens: monitoring a line, watching the graphs and stacks, then analysing what was recorded. Chapter 12 covers reports, chapter 13 logs and diagnostics, chapter 14 backup and transfer, and chapter 15 troubleshooting. Reference material is collected in the appendices.

Procedures are set out in numbered, titled blocks like this one, and can be followed without reading the surrounding description:

To open the built-in manual

1. Press **F1** anywhere in the main window, or click **? Help** in the Setup group.
2. Use **Ctrl + F** to search within it.

2. Installation

2.1 System requirements

ITEM	REQUIREMENT
Operating system	Windows 10/11 (64-bit) • Linux x86-64 with glibc 2.31 or newer (Ubuntu 20.04+, Debian 11+, RHEL/Rocky 9+ and equivalents) • macOS on Apple Silicon or Intel
Memory	8 GB minimum; 16 GB recommended for long lines
Disk	≥ 2 GB for the application. Project storage grows with the survey — the NPZ shot files, trace stacks, session logs and the database all live in the project folder.
Display	1280×768 or larger. The interface adapts to small screens: the toolbar scales its font and collapses groups rather than forcing a wide window, and every window can be freely resized.
Network	Not required. The application never connects to the internet and licensing is fully offline. A network drive may be used for the watched SEG-D folder; see §4.4 before putting the <i>project</i> on one.

2.2 Installing on Windows

To install on Windows

1. Unpack the distribution ZIP anywhere you like, for example `C:\SEGDMonitor\`. All files must stay together in that folder.
2. Start the application by double-clicking `SEGDMonitor.exe`.
3. Optional: right-click `SEGDMonitor.exe` → **Send to** → **Desktop (create shortcut)**.

Windows SmartScreen may show “*Windows protected your PC*” for a newly downloaded, unsigned application. Click **More info** → **Run anyway**. This warning concerns download reputation only.

2.3 Installing on Linux

To install on Linux

1. Unpack the tarball: `tar -xzf segd-monitor-linux-x86_64.tar.gz`
2. Start the application: `./x_corr_app_v1_5_0.dist/segd-monitor`
3. Optional: create a launcher or symlink, e.g. `ln -s $PWD/x_corr_app_v1_5_0.dist/segd-monitor ~/bin/segd-monitor`

The application runs on both X11 and Wayland.

2.4 Installing on macOS

The macOS build is an application bundle. It is ad-hoc signed and **not notarised**, so the first launch needs one extra step — double-clicking it will be refused, with a message that macOS cannot verify the application. How the refusal is lifted depends on the macOS version.

To install on macOS 15 or later

1. Unpack the ZIP and move **SEG-D Monitor.app** into your **Applications** folder.
2. Double-click the application once. macOS refuses to open it; click **Done** — not *Move to Trash*.
3. Open **System Settings** → **Privacy & Security** and scroll down to **Security**. A line there names SEG-D Monitor as blocked: click **Open Anyway** and confirm with your password or Touch ID.
4. Click **Open Anyway** in the dialog that follows. This is only needed once — subsequent launches behave normally.

To install on macOS 14 or earlier

1. Unpack the ZIP and move **SEG-D Monitor.app** into your **Applications** folder.
2. **Right-click** (or `Control`-click) the application and choose **Open**.
3. In the dialog that appears, click **Open** again. This is only needed once — subsequent launches behave normally.

On any version the same can be done in Terminal, which suits a machine set up remotely:

```
xattr -cr "/Applications/SEG-D Monitor.app"
```

Builds are architecture-specific: an Apple Silicon build will not run on an Intel Mac or vice versa. Check which build you were sent if the application refuses to start.

2.5 Upgrading to a new version

The application folder holds nothing of yours. To upgrade, replace it wholesale with the new one. Your projects, your settings and your license are untouched: settings live in the project database, and the license lives in your user profile ([§3.4](#)).

3. First Run and Licensing

3.1 What happens on first start

On first start the application opens in FREE mode and shows the welcome screen (chapter 4), which on a first start offers **Create your first project...** No activation, registration, or internet connection takes place — the application is usable immediately. Licensing only concerns the professional modules listed in [§1.2](#).

3.2 Obtaining a license

Licenses are issued per machine: a license is generated from your computer's Machine ID and will only validate on the machine it was issued for. There are two ways to get one. A **free license**, with every module, comes from the SEG-D Monitor website, www.segd-monitor.com, and lasts three months at a time. A **full license** — longer, or perpetual — comes from the vendor.

To get a free license online

1. Click  **License** in the Setup group of the toolbar, then **Get a free licence online**. Your browser opens the SEG-D Monitor website with this computer's Machine ID already filled in.
2. Enter your name, your company or vessel and your e-mail address, and click **Get my licence**. Nothing is sent to you by e-mail: the license is offered on the page itself.
3. Click **Download licence (.lic)** and save the file.
4. Back in the application, click **Load License File...** and select it. Every licensed button unlocks immediately.

The application itself never connects to the internet — the button only opens your browser. On a computer with no connection, read the Machine ID from the License dialog (**Copy** puts it on the clipboard), open www.segd-monitor.com on any other computer or a phone, type the Machine ID into its licence form, and carry the `.lic` file across.

QUESTION	ANSWER
How long does it last?	Three months from the day it is issued. The License dialog shows the expiry date.
How is it renewed?	On the same page, with the same Machine ID and e-mail address, in the license's last week . The new license starts where the old one ends, so renewing early loses nothing. Load the new file as before.
And if I ask earlier?	The page gives the same license again and uses nothing up — the way to fetch a file that was lost. It is given only to the e-mail address it was issued to.
How many can I have?	The page shows how many of this computer's free licenses are used. Together they cover a year; after that, contact the vendor for a full license.
Which version do I need?	2.26.3 or later. An earlier version refuses a free license longer than a month, with <i>"This trial licence is longer than 38 days"</i> — install the current version and load the file again.

To obtain and load a full license from the vendor

1. Click  **License** in the Setup group of the toolbar. The button shows **FREE** or **PRO** according to the current state.
2. Read the **Machine ID** — a 16-character code such as A3F9C2E17B04D856 . It identifies this computer, contains no personal data, and cannot be converted back into any system information.
3. Send the Machine ID to the vendor together with the customer name that should appear on the license, and the modules and duration you are purchasing.
4. Receive your license file — a small text file with the extension `.lic` — and copy it onto the computer. E-mail attachment, USB stick or network share; the transport does not matter.
5. Click  **License** → **Load License File...** and select the `.lic` file. The licensed buttons unlock immediately and the dialog shows the customer name, enabled modules, source limit and expiry.

Each computer has its own Machine ID, so a license file from one machine will not validate on another. When ordering for several workstations, run the application once on each (FREE mode is enough), collect each Machine ID, and request one license per machine.

3.3 How the license works

- **Offline and self-contained.** The license file carries a cryptographic signature which the application verifies locally at every start. There is no internet check, no license server and no “phoning home”.
- **Machine-locked.** The Machine ID inside the license must match the computer. The file cannot be edited: any modification to the name, date, features or machine breaks the signature and the license is rejected.
- **Time-limited or perpetual.** Licenses carry an expiry date, or none. The status line shows the remaining days when fewer than 30 remain. Setting the system clock backwards does not extend a license — a rollback of more than 48 hours suspends it until the clock is corrected.
- **Feature-based.** Each license enables an explicit list of modules and a maximum source count, so tiers can be tailored — for example BBT-only, or everything unlimited.

3.4 Where the license is stored

When a license file is loaded it is copied into the per-user configuration directory and re-loaded automatically on every start, so you never load it twice:

SYSTEM	LICENSE LOCATION
Windows	%APPDATA%\XCorrSEG\license.lic (C:\Users\ <name>\appdata\roaming\xcorrseg\)< td=""></name>\appdata\roaming\xcorrseg\)<>
Linux	~/.config/XCorrSEG/license.lic
macOS	~/Library/Application Support/XCorrSEG/license.lic

The original `.lic` file you received stays untouched wherever you saved it — keep it as a backup. The application also accepts a `license.lic` placed directly next to the executable, which is convenient when IT prepares workstations centrally. Deleting the stored file returns the application to FREE mode; loading the backup restores it. Upgrading or reinstalling never touches the stored license.

The license belongs to the *user profile*, not the application folder. If a different operating-system user runs the application on the same machine, that user loads the `.lic` once as well.

3.5 Renewals and replacements

A renewal or an upgrade — more modules, more sources, a longer period — is simply a new `.lic` file for the same Machine ID. Load it over the old one and you are done. If a workstation is replaced, send the new machine's ID to the vendor for a replacement license.

4. Projects and the Project Database

4.1 The project concept

All work happens inside a **project**: one folder of your choosing holding a single SQLite database file plus every NPZ file the monitor records, the trace stacks, the session logs and the generated reports. The database stores the complete configuration — channel sets, source and gun layout, shot-point settings, interface state — and every decoded shot of every line. Graphs, viewers, bubble tests and reports all read from it.

Copying the project folder therefore backs up or transfers the entire survey. Nothing else is needed.

4.2 Creating, opening and switching projects

The welcome screen appears on every start. **Continue** — or simply **Enter** — reopens the last project; the cards beside it open another recent project, and **New project...** and **Open existing...** create or open one. Each recent project shows its last line and shot count once its database has been read. *Not found* means the file is no longer at that path, and nothing is created there; *Details unavailable* means its details could not be read — the drive did not answer in time, or the file could not be read as a project; clicking it tries again.




Figure 1. The welcome screen. Closing it exits the application.

The Project dialog can be reached at any time with  **Project...** in the Setup group.

Create NEW project

Project directory: **Browse...**

Database name:

Import settings from (optional): **Browse...**

Create Project

Open EXISTING project database

Browse for .db file...

Recent projects (double-click to open):

/Users/vlad/Documents/Test_seg_d_monitor/IARA_4D.db

Cancel

Figure 2. The Project dialog. The main window title always shows the active project.

To create a new project

1. Stop monitoring if it is running — the project cannot be switched during a session.
2. Click **Project...** in the Setup group.
3. Under **Create NEW project**, browse to the folder that will hold the survey.
4. Enter a database name, for example `Survey_2026_Block7`.
5. Optionally set **Import settings from** to a previous project's `.db` file, to start pre-configured with its channel sets and Source Configuration.
6. Click **Create Project**. The window title updates to the new project.

To open an existing project

1. Click **Project...** (with monitoring stopped).
2. Under **Open EXISTING project database**, click **Browse for .db file...**
3. Select the project's `.db` file and confirm. This is also how you open a project folder copied from another machine.

To start a new survey pre-configured from an old one

1. Create a new project as above.
2. Set **Import settings from** to the previous survey's database before clicking **Create Project**.
3. Open **Settings** → **Source Configuration** and check the imported array against what is actually deployed — the import copies the old layout verbatim (chapter 6).

4.3 What lives where

LOCATION	CONTENTS
<project>/<name>.db	The project database: settings, every shot of every line, bubble tests, theoretical k values.
<project>/npz/<SEG-D folder>/shots/	One NPZ per shot and source, named e.g. LINE001_SP001005_Source1.npz , with embedded metadata (line, SP, source, dt, windows, source file).
<project>/npz/<SEG-D folder>/trace_stacks/	Live trace-stack files, e.g. Stack_Source1_All_Channels.npy .
<project>/npz/<SEG-D folder>/bbt/	Decoded traces saved with bubble tests, enabling later re-processing without the raw SEG-D files.
<project>/line_reports/	End-of-Line PDF reports.
<project>/bbt_reports/	Bubble-test PDF reports.
<project>/logs/	One log file per monitoring session; the 30 most recent are kept (chapter 13).

Re-running a line replaces its data cleanly. When a line name is first seen after Start, that line's previous rows are deleted before the new shots are written. Other lines are never touched.

4.4 Projects on a network share

The project database is opened in WAL mode, which is what allows the viewers to read while the monitor writes. WAL needs a small shared-memory file alongside the database, and **SQLite cannot create that file on an SMB or NFS share**. On a network project folder the setting silently does not take effect, readers block the writer, and a busy line can lose shots to lock errors.

The application detects this and says so — both in the log and in a dialog at start, because it is a condition that costs data and that the operator can otherwise not see.

Keep the project folder on a local disk. The watched SEG-D folder may be a network share without any problem — it is only read. Copy the finished project to the share afterwards (chapter 14).

5. The Main Window

5.1 Layout

The main window is the monitoring cockpit. From top to bottom it holds the **command bar** (the toolbar groups), the collapsible **SEG-D folder panel**, the **monitoring view** with the similarity heatmap and the optional live-graphs panel, and the collapsible **log panel**.

This is the **Classic** layout, the one the window starts with, and the one sections 5.2 to 5.8 describe. The same tools, plots and log can be arranged two other ways, **Rail** and **Focus** (§5.9).




Figure 3. The main window before a session starts, with the SEG-D folder panel expanded. The heatmap area below fills as shots arrive.

5.2 The toolbar groups

The command bar is organised into nine labelled groups, left to right in the order of work: **Source**, **Reference**, **Monitoring**, **Status**, **Display**, **Inspect**, **Analysis**, **Setup** and **Log**. Source, Monitoring, Status and Log are always shown. Reference, Display, Inspect, Analysis and Setup are compact chips that pop open when clicked (§5.3).

A padlock (🔒) on a button means the module needs a license (§1.2).

Source

CONTROL	FUNCTION
SEG-D Folder ▾	Shows and hides the SEG-D folder panel. It collapses automatically when monitoring starts.
Summary chip	Always visible, showing the selected SEG-D folder and the reference Start will use (↑ or ↓) — or <i>no folder • no reference</i> . Click it to reopen the setup panel.

Reference

CONTROL	FUNCTION
Asc ↑ / Desc ↓	The order the SEG-D files are processed in (by file name) — on a line whose files carry the shot point, the line direction — and which reference is used. Locked from Start to Stop.
Reference per direction	Each direction can have a reference of its own: a line shot the other way often correlates better against a reference built from that direction's shots. A direction without one shows <i>same as ...</i> and shares the other's.
⋮ menu	Browse... an existing .npz, Generate from shots... (§7.2), View the file that direction uses, or Same as the other direction .

Monitoring

BUTTON	FUNCTION	SEE
 Start	Begin watching the SEG-D folder. Every arriving file is decoded, correlated, stored and shown.	7
 Stop	Stop watching and flush the database. The End-of-Line report is then offered for the line(s) of this session.	7, 12
 Graphs	Show or hide the live trend graphs beside the heatmap.	8

Status

CONTROL	FUNCTION	SEE
Session state	Idle , or the stage monitoring is at.	5.6
Files: N Masks: N	Files processed and nav masks seen in this session.	5.6
QC issues	The counters of nav-mask, partial-mask and missing-shot findings.	5.6

Display

CONTROL	FUNCTION
Color scale min / max	Similarity colour-mapping limits, 0...1. Narrow the range to make small differences visible.
Autocorr. window from / to	The time window (ms) of the autocorrelation comparison. The same window is applied online to every shot and is embedded in the recorded NPZ metadata.
Heatmap points / Full range	How many recent shots the heatmap shows, or the whole line.
 Neighbours	Watch your neighbours: tells a change of the source from one bad hydrophone, using the colour-scale minimum as its limit. Can be switched on and off while monitoring (§5.5.1).

Inspect

BUTTON	FUNCTION	SEE
 Headers	Open any SEG-D file and inspect its decoded headers: general headers, channel sets, extended and external blocks, and the per-trace G_FIRE gun data. Exports to CSV.	7
 Inspect NPZ	Quick look inside any NPZ file: arrays, shapes and metadata.	10
 Seismic	Wiggle and stack display of recorded NPZ files.	10
 Stacks	Live trace-stack windows, including  By Sub-arrays.	9

Analysis

BUTTON	FUNCTION	SEE
 Data	Filterable table of every stored shot, with CSV export.	10
 Database	Gun performance, line replay and per-source volume / P-B analysis.	10
 BBT	The independent Bubble Test window.	11
 SP Quality	Shot-point quality against the contractual acceptance rules, one line at a time.	10
 Line Report	Generate an End-of-Line report for any line already in the project database. Works offline, with no monitoring session.	12

Setup

BUTTON	FUNCTION	SEE
 Designer	Graphical editor of the source array geometry. Free.	6
 Settings	Channel Sets, Source Configuration and Shot Points.	6
 Project...	Create, open or switch projects.	4
 Manage Lines	Permanently delete a line's shot data and its NPZ files.	10.6
 Branding...	Choose up to three customer logos for this project's reports.	12
 License	Machine ID, license loading and current status. Shows FREE or PRO .	3
 Help	The built-in manual. Never license-gated. Also F1 .	—
Theme	Interface theme, applied to every open window and remembered.	5.7
Reduce motion	Windows appear in their final state at once, with no animation.	5.7
Layout	How the main window is arranged: Classic, Rail or Focus. Fixed between Start and Stop.	5.9

Log

CONTROL	FUNCTION	SEE
 Log ▾	Show or hide the scrolling processing log.	5.6
 Logs	Open the folder holding this project's session logs.	13
 Save Log	Save what is currently on screen to a text file.	13

5.3 Collapsible groups

Inspect, Analysis, Display and Setup are used far less often than Source and Monitoring, so each is shown as a compact chip rather than permanently occupying bar width. Clicking a chip pops its buttons open in a floating overlay above the window; opening one closes any other that was open. Each group remembers whether it was left open.

This is what keeps all seven groups on a single row. On a narrow window the toolbar additionally scales its font down rather than wrapping onto extra rows, so nothing is ever pushed out of reach.

5.4 The SEG-D folder panel

The SEG-D folder panel holds what you choose once per line:

CONTROL	MEANING
SEG-D Folder	The folder the monitor watches. Every new SEG-D file appearing here is decoded automatically while monitoring runs. NPZ output is grouped under this folder's name inside the project.
Automatic	Lines start, stop and report themselves. Beside the tick box are the quiet time after which a line is stopped, and the line letters: a new folder is started when its line name carries any one of them. Can be changed while monitoring (§7.8).

The panel **collapses automatically when monitoring starts**, giving its height to the heatmap. The summary chip in the Source group keeps showing the selected folder and the reference Start will use; click the chip, or  SEG-D Folder ▾, to bring the panel back.

5.5 The similarity heatmap

Each column is one shot and each row one channel, coloured by the similarity of that channel's signature to the reference. One heatmap is drawn per active source.




Figure 4. The similarity heatmap. Hovering reads out the value; right-clicking a cell opens the underlying traces.

To inspect the traces behind a heatmap cell

1. Right-click the cell.
2. Choose **Wiggle Stack** for that channel across files, **File Wiggle** for all channels of that one file, or **View Complete Headers** for the file's decoded headers.

5.5.1 Watch your neighbours

The heatmap colours a coefficient that is at or below the colour-scale minimum, and leaves the rest to the reader: was this hydrophone like that on the shots before, is it still like that, and did the hydrophones next to it move as well? With **Display** → **Neighbours** switched on, the application does that reading on every shot. It uses the coefficients already on the map and computes nothing of its own, so it works in Light Mode too.

STEP	WHAT IS CHECKED
1. The hydrophone itself, before	A coefficient at or below the minimum is compared with the same hydrophone's previous shots of the same source. If too few of them were also at or below the minimum, it is one isolated shot and nothing happens.
2. The hydrophone itself, after	Otherwise its next shots are counted. If the trouble does not hold, it is dropped.
3. Its neighbours	If it holds, the hydrophones next to it are looked at. A cluster with one hydrophone: the two nearest hydrophones. A cluster with two: its partner and the four nearest. Nearest means the next clusters along the same sub-array, in the order of the Source Configuration (§6.3); at the end of a string, the next ones on the one side.

A neighbour has *changed* when its usual value during the event lies outside everything it read on the shots before the event, by more than its own normal spread. It does not have to fall below the minimum: a neighbour that only moves is evidence too.




Figure 5. One hydrophone stays below the minimum and the hydrophones around it move with it: an alarm, outlined in black, steaming, with its bar and its two answers above the map. Lower down, one hydrophone is bad on its own: a note, outlined in dashed amber. The single red cell near the start is an isolated shot and is left alone.

RESULT	WHAT IT MEANS AND WHAT IS SHOWN
Alarm	The neighbours changed with the hydrophone: the source did something at that place. The area is outlined in black — the hydrophone and the neighbours that changed, from the first bad shot onward, growing while it lasts. A bar above the map names the source, the hydrophones and the shot points; the Status group shows Neighbours: N ; the log has a line.
Note	The neighbours did not change: one hydrophone, not the source — most often a sensor, a cable or a channel. A thin dashed amber outline on that row and a line in the log; no bar and no counter. The neighbours are looked at again on every later shot, so a note becomes an alarm if they start to move.
Nothing	An isolated shot, or trouble that did not hold.

An event ends when its hydrophones have read inside the limits for the look-back number of shots in a row. What follows is a new event.

An alarm calls until it is answered

While nobody has answered an alarm, the word **AIRLEAK** is sent in Morse through the machine's speakers — again and again, with a short pause, at the system's volume — and **steam** rises from the outlined area on the map. There are two answers, both on the bar above the map.

ANSWER	WHAT IT DOES
Noticed	Seen — keep monitoring. The sound and the steam stop. The alarm stays: the black outline, the bar (now reading <i>Noticed</i>) and the counter. It calls again only if it <i>spreads</i> — another hydrophone of the group goes bad, or another neighbour starts to move. When the event ends by itself it leaves the bar.
Ignore	Silent for as long as that event lasts. The outline turns grey and the bar and the counter clear. The application alarms again at the next new event: the same hydrophones after they recovered, or a different group of hydrophones of that source.

An alarm that ended before anyone answered goes on calling until it is answered. Stop does not silence it; switching **Neighbours** off or closing the application does. A note never sounds and never steams.

- **The sound** has a tick box in **Settings** → **Performance**, on by default. The application plays it with what the system already has. A machine with no audio player at all gets the plain system beep instead, and the log says so once.
- **The steam** follows **Reduce motion** (§5.7): with that on, or under remote desktop, there is no animation and the outline is the cue. The mouse works through the steam — hovering and right-clicking the map are unchanged.

Switching it on and off. The button works while monitoring; in the Rail and Focus layouts it is also beside Light, and the command palette has it by name. Switched on, the check starts at the next shot and nothing earlier counts — so an event already running at that moment can only be a note until its hydrophones have recovered once, because there are no shots before it to compare the neighbours with. Switched off, the bar and the outlines are cleared. The choice is stored with the project.

The numbers — how many shots to look back, how many to listen ahead, and the share of them that must be at or below the minimum — are in **Settings** → **Performance**; by default five shots, five shots and half of them. The limit itself is always the colour-scale minimum in force at the moment: move it and the check follows what the map shows.

This is information only. Nothing is stored, and it never reaches the End-of-Line report or the spec check.

5.6 The log panel and status labels

The status labels sit in the toolbar at all times, so they stay readable whether or not the full log text is expanded:

- **Session state** — *Idle*, or the current monitoring phase.
- **Files: N | Masks: N** — how many files have been processed and how many distinct nav masks have been seen.
- **QC issues** — a running count of detected problems.

The QC issue counters count files, not decode attempts. A file that is retried does not inflate them. What the log lines themselves mean is covered in §7.4.

5.7 Themes

The theme selector in the Setup group restyles every open window of the application immediately, and the choice is remembered per user — it is not part of the project, so it travels with the workstation rather than the survey.

Reduce motion, beside the theme selector, switches off the application's animations — the welcome screen then appears complete at once. It is meant for remote desktop and slow links and is remembered per user like the theme. In a Windows remote-desktop session animation is off whatever the box says.

5.8 Window behaviour

All windows — the main window and every dialog — are freely resizable and maximizable, and never open larger than the available desktop. On small displays the toolbar scales down and the parameter rows wrap instead of forcing a wide window.

5.9 Layouts and the command palette

Setup → **Layout** arranges the main window one of three ways. Every layout has the same tools, the same plots and the same log; only where they sit differs, and each tool opens the same window whichever layout it is started from. The choice is remembered per user, like the theme, and applies at once. It cannot be changed between Start and Stop.

LAYOUT	ARRANGEMENT
Classic	The command bar across the top, as described in the sections above. The window always opens the first time in this layout.
Rail	The tools as a menu down the left side. A header band names the line (the project, until a line is running) and the session state, and holds Start, Stop and Light. The plots are in the centre; Session, QC and the log are in a column on the right; a status line runs along the bottom.
Focus	One slim bar across the top with Start, Stop and Light, and the plots over nearly the whole window. A small dock under them holds Graphs, Stacks, Seismic, SP quality and Line report; the same column is on the right.




Figure 6. The Rail layout before a session starts: the tools down the left side, the header band, and Session, QC and Log on the right.




Figure 7. The Focus layout: one bar across the top, and the dock of the five most-used tools under the plots.

In Rail and Focus:

- The **SEG-D folder**, **Reference** and **Display** buttons open the same panels as in Classic ([§5.3](#), [§5.4](#)).
- **Setup** — at the foot of the menu in Rail, the gear at the right of the bar in Focus — opens the Setup group with every Setup tool, the theme and the Layout box. The way back to another layout is therefore always in the same place.
- The **Session** and **QC** panels show what the Status group shows in Classic, one figure to a row; the log is always open.
- A tool that needs a licence shows a padlock, and one that cannot run at the moment is greyed, exactly as its Classic button is.
- In a narrow window the Rail menu shows icons only; hover over one for its name.

The command palette

Ctrl + K (**⌘ K** on macOS) opens a list of every tool by name, in any layout. Type a few letters to narrow the list, use **↑** and **↓** to move, **Enter** to run the selected tool and **Esc** to close. A tool that cannot run at the moment is listed as not available. The layouts and the themes are in the list as well, so a layout can be changed without opening Setup.

6. Configuring the Source Array

This is the most important setup step in the application. The source configuration tells the monitor which guns exist, which sub-array each belongs to, and which hydrophone channel records each one. Everything downstream — the graphs, the stacks, the bubble test, the missing-shot detection and every report — is computed from it. If it does not match the deployed array, the data will look wrong in ways that are hard to trace back.

6.1 The model, and how it reaches the app

There are two places to work: the  **Designer**, which builds a layout graphically, and  **Settings** → **Source Configuration**, which is the authoritative table the application actually uses. The normal workflow is to rough out the array in the Designer, push it across, then fine-tune in the table.

DESIGNER	BECOMES	SOURCE CONFIGURATION
Array	→	Source
Sub-array label	→	Subarray
Cluster	→	(grouping only — its guns share one NFH channel)
Gun	→	Gun, with its volume
Hydrophone	→	NFH Channels

Source, sub-array and gun IDs are arbitrary integers. They need not be consecutive and need not start at 1 — use whatever numbering the vessel uses.

6.2 Building an array in the Designer




Figure 8. The Designer with four sub-arrays of two clusters, three guns per cluster. The nav-mask readout on the right updates as the configuration changes.

To build a new source array from scratch

1. Open **Designer** from the Setup group.
2. In **Array Configuration**, set **Sub-arrays** to the number of strings in one source (1–16).
3. Set **Clusters / Sub-array** (1–10) and **Guns / Cluster** (1–4) to match the deployed array.
4. Set **Gun Volume** to the volume in cu.in. that new guns should be created with. Individual guns can be changed afterwards.
5. Set **Hydrophones / Cluster** — 1 or 2. Use 2 only where the cluster really carries two NFH channels; the BBT window's per-file hydrophone selector depends on this being right.
6. Click **Create New Array**. The array is drawn on the canvas and added to the **Current Array** list.
7. Click a cluster on the canvas to select it, then **Fill Selected Cluster** to populate it with guns at the configured volume. **Clear Selected Cluster** undoes it. Several clusters can be selected at once.
8. Repeat steps 2–7 for each further source.
9. In **Navigation Mask Configuration**, set **Number of arrays** and tick the sub-arrays that are active. The mask readout below updates live — check it against what the navigation system will actually send ([§6.4](#)).
10. Click **Apply to Main App**.
11. Open **Settings** → **Source Configuration** and verify every gun's Active flag and NFH channels. The Designer assigns one channel per cluster; per-gun corrections are made here.
12. Click **OK**. The configuration is stored in the project database, so it travels with the project.


To duplicate a source for a mirrored array

1. Select the source to copy in **Current Array**.
2. Click **Duplicate Current Array**. A new array is added with the same geometry and fresh gun numbers.
3. Adjust its gun volumes or hydrophone assignments as needed, then  **Apply to Main App**.

The Designer can also write the layout out as files: **Export Gun Map TXT**, **Export Navigation Masks TXT**, **Export Arrays to PNG** (the schematic used in bubble-test reports is drawn from this layout), and **Export Config JSON / Load Config JSON** for the Designer's own format.

6.3 The Source Configuration tab

This tree is what the application actually reads. Each row is a Source, a Subarray or a Gun, with an **Active** checkbox, the gun's **NFH Channels** and its **Volume**.



Define sources, their subarrays, and the guns in each subarray.
Uncheck 'Active' to exclude a subarray or gun from nav-mask detection and from every analysis - inactive subarrays/guns never contribute a trace
'NFH Channels' are the trace/channel number(s) (in the active channel set) that monitor a gun - enter one or more, e.g. '1' or '1, 2' for a gun with 2

Type	ID	Active	NFH Channels	Volume
Source	1	☒		
Subarray	1	☒		
Gun	1	☒	11, 12	45
Gun	4	☒	13, 14	70
Subarray	2	☒		
Gun	2	☒	11, 12	45
Gun	5	☒	15, 16	100
Subarray	3	☒		
Gun	3	☒	13, 14	70
Gun	6	☒	15, 16	150
Source	2	☒		
Subarray	1	☒		
Gun	7	☒	21, 22	45
Gun	10	☒	23, 24	70
Subarray	2	☒		
Gun	8	☒	21, 22	45
Gun	11	☒	25, 26	100
Subarray	3	☒		
Gun	9	☒	23, 24	70
Gun	12	☒	25, 26	150
Source	3	☒		
Subarray	1	☒		
Gun	13	☒	31, 32	45
Gun	16	☒	33, 34	70
Subarray	2	☒		
Gun	14	☒	31, 32	45

Buttons: Add Source, Add Subarray, Add Gun, Remove Selected, Auto-number NFH Channels, Load Source Config JSON, Save Source Config JSON, Activate All, Deactivate All, Test Nav Masks, OK, Cancel

Figure 9. Source Configuration — the single source of truth for the array layout.

Unchecking Active removes that sub-array or gun from everything — nav-mask detection and every computation. An inactive gun never contributes a trace. This is how you exclude a gun that has been recovered or is known bad, but it also means an accidentally unticked box silently drops data.

To add a sub-array to an existing source

1. Open  **Settings** → **Source Configuration**.
2. Select the **Source** row it belongs to.
3. Click **Add Subarray**. It is added under that source.
4. Edit its **ID** so it matches the sub-array's real number — this is what determines its nav-mask bit ([§6.4](#)).
5. Select the new sub-array and click **Add Gun** once per gun.
6. Fill in each gun's **NFH Channels** and **Volume**, then click **OK**.

To assign hydrophone channels to a cluster

1. Select a gun row and click in its **NFH Channels** cell.
2. Enter the trace/channel number in the active channel set that records this gun — for example 11.
3. For a cluster with two hydrophones, enter both separated by a comma: 11, 12. The BBT window's per-file hydrophone selector offers exactly these.
4. Guns in the same cluster normally share the same channel or channels — enter the same value on each of them.
5. Click **OK**.

To number every gun's channels at once

1. Click **Auto-number NFH Channels**.
2. Sequential channel numbers are assigned to every gun in tree order.
3. Check the result against the deployed cabling before clicking **OK** — this is a convenience for a regular array, not a substitute for knowing the wiring.

To change gun volumes after a swap

1. Open  **Settings** → **Source Configuration**.
2. Click the **Volume** cell of the affected gun and enter the new volume in cu.in.
3. Click **OK**. Ideal-volume figures in the graphs and in the End-of-Line report follow the new value from the next shot onward.

To take a gun out of service mid-survey

1. Open  **Settings** → **Source Configuration**.
2. Untick **Active** on that gun's row.
3. Click **OK**. The gun stops contributing to nav-mask detection, the graphs and the reports.
4. Re-tick it when the gun is back. **Activate All** and **Deactivate All** set every row at once.

6.4 Nav masks

The SEG-D external header carries a **gun nav mask**: a bit field saying which sub-arrays fired on that shot. The mapping is fixed and simple:

Sub-array N maps to bit $(N-1)$ of the nav mask. Sub-array 1 is bit 0 (value `0x0001`), sub-array 2 is bit 1 (`0x0002`), sub-array 3 is bit 2 (`0x0004`), and so on. A source whose sub-arrays are 1–4 therefore has mask `0x000F`.

This is how the application knows which source a shot belongs to. If the mask a shot carries matches no configured source, the file cannot be attributed and is reported as unreadable rather than being silently dropped — see [§7.4](#).

To verify the nav masks before a line

1. Open  **Settings** → **Source Configuration**.
2. Click  **Test Nav Masks**. The masks are computed from the current, still-unsaved tree — you do not have to click OK first.
3. Compare each source's computed mask with what the navigation system is configured to send.
4. If they disagree, correct the sub-array **IDs** (not the gun numbers) — the ID is what sets the bit.

6.5 Worked example: 3 sources × 4 sub-arrays × 6 guns

A three-source array, each source with four sub-arrays of two clusters, three guns per cluster and one hydrophone per cluster.

In the Designer: Sub-arrays 4, Clusters / Sub-array 2, Guns / Cluster 3, Hydrophones / Cluster 1, Number of arrays 3. Create New Array, select all clusters, Fill Selected Cluster, then repeat for the second and third source — or use Duplicate Current Array twice.

That yields 24 guns per source, 72 in total. Applying it to the main app produces a Source Configuration of this shape:

SOURCE	SUBARRAY	GUNS	NFH CHANNELS
1	1	1–6	11, 12
1	2	7–12	13, 14
1	3	13–18	15, 16
1	4	19–24	17, 18
2	1–4	25–48	21...28
3	1–4	49–72	31...38

And these nav masks:

SUB-ARRAYS ACTIVE	BITS	MASK
1, 2, 3, 4	0, 1, 2, 3	<code>0x000F</code>
1, 2 only	0, 1	<code>0x0003</code>
3, 4 only	2, 3	<code>0x000C</code>

During monitoring, a shot arriving with mask `0x000F` is attributed to the source whose four sub-arrays are all active, its 24 guns are read from the G_FIRE headers, and one shot-level summary row plus 24 per-gun rows are written to the database. A shot arriving with a mask that matches no configured source produces a single warning per file and is counted as unreadable.

6.6 Channel Sets

The Channel Sets tab defines which SEG-D channel sets exist and which one is **active for analysis**. Each set records its channel count, record length in milliseconds, and sample interval.

Channel Sets Source Configuration Shot Points Performance Far Field Hydrophone QC

Active set for analysis: Set 4 ▼

The active set feeds the similarity heatmaps, P/B, the trace stacks, the per-shot NPZ files and the Bubble Test. Far Field and Hydrophone QC read it too, unless their own tab names another set.

Set	Channels	Record (ms)	dt (ms)
1	78	99,2	0,100
2	90	512,0	0,100
3	1	512,0	0,100
4	90	10240,0	0,500
5	0	0,0	0,000

OK Cancel

Figure 10. Channel Sets. The active set is the one every NFH-channel number in the Source Configuration refers to.

NFH channel numbers in the Source Configuration are interpreted in the **active** channel set. Changing the active set changes what every gun's channel number means.

6.7 Shot Points and the spec check

Channel SetsSource ConfigurationShot PointsPerformanceFar FieldHydrophone QC

Shot point increment used to detect missing shots (both live monitoring alerts and the End of Line report). Direction (increasing/decreasing) is detected

SP increment: 1,00

☐ Enable bad-shot spec check (End of Line report)

Spec check criteria and limits

A shot point counts as BAD when this many of its guns fail a per-gun criterion, or when a shot-level criterion fails:

☒ Source fired nothing ☒ Timing ☒ Depth ☒ Pressure ☒ Nav mask

Shot is bad when this many guns fail: 1 (nav mask mismatch always marks the shot bad)

Max consecutive bad shots per source: 4

Window 1 - max bad: 6 within this many shots: 12

Window 2 - max bad: 8 within this many shots: 20

Max bad shots overall (%): 3,00

Minimum acceptable range length (shots): 80

Gun firing synchronisation

☒ Time-break sync sample interval (ms): 2,0000 ☐ Count as bad shot

☒ Firing spread first to last gun, at most (ms): 1,00 ☐ Count as bad shot

Ticked rules are shown on the live graphs and in the End of Line report. "Count as bad shot" also makes the rule a Spec Check criterion (spec check enabled) - one gun is enough, whatever "this many guns fail" says.

OKCancel

Figure 11. Shot Points. The spec check is off by default; when enabled it adds a per-line verdict page to the End-of-Line report.

SP increment is the nominal spacing between shot points. **Detect from data** reads it from a line already in the database.

Missing-shot detection does not use this setting. The interval is always derived from the line's own shot-point transitions, and where a line cannot say, nothing is reported rather than a guess. A wrong value here therefore cannot produce phantom missing shots. Direction — increasing or decreasing SP — is detected automatically too.

Enable bad-shot spec check turns on contractual line acceptance. When enabled, the criteria below it become editable:

SETTING	MEANING
Misfire / Timing / Depth / Pressure / Nav mask	Which per-gun failures make a shot count as bad. A nav-mask mismatch always marks the shot bad regardless.
Guns failing per shot	How many guns must fail before the shot itself counts as bad.

Gun firing synchronisation	A group of its own, working with the spec check on or off. Time-break sync measures every gun from the time break (its aim point) against half the seismic recording's sample interval - enter it from the recording system's setup, not a near-field hydrophone set's. Firing spread measures each shot's first-to-last gun time against the contract's limit. A ticked rule is shown on the live graphs and in the report as information; Count as bad shot (spec check enabled) also makes it a bad-shot criterion, one gun being enough. A deliberately staggered (beam-steered) array exceeds any small spread limit on every shot.
Max consecutive bad shots per source	The run-length rule.
Window 1 and Window 2	Two sliding-window rules: at most n bad shots within m consecutive shots.
Max bad shots overall (%)	The whole-line percentage limit.
Minimum acceptable range length	How long a clean stretch must be to count as acceptable.

The live check during acquisition reports the run-length and sliding-window rules, but **not** the whole-line percentage — a partial line's percentage is not the line's. One bad shot in the first 32 is 3.1%, over a 3% limit, on a line that finishes at 2.0%. The percentage rule is therefore evaluated only at the end, in the report.

6.8 Carrying a configuration to another survey

To reuse a configuration on a new project

1. *Either* create the new project with **Import settings from** pointing at the old project's database ([§4.2](#)) — this copies the channel sets and the Source Configuration together.
2. *Or* open the old project, go to **Source Configuration** and click  **Save Source Config JSON**; then open the new project and click  **Load Source Config JSON**.
3. Either way, check the result against the deployed array before starting a line — an imported configuration is the *old* survey's, not this one's.

6.9 Far Field

The far-field tab switches on the far-field signature QC described in [§8.4](#), [§12.7](#) and [Appendix F](#). It is **off by default**, and it runs in normal mode only — Light Mode never reads the traces it needs.

Channel Sets
Source Configuration
Shot Points
Performance
Far Field
Hydrophone QC

☒ Compute each shot's far-field signature (normal mode only)

Signal
Channels
Geometry
Targets
QC

Hydrophone to use:

Odd channel of the pair

Hydrophone per cluster (blank = the default above):

Source	Cluster	NFH channels	Use
1	1	11, 12	Default — ch 11
1	2	13, 14	Default — ch 13
1	3	15, 16	Even — ch 16
2	1	21, 22	Default — ch 21
2	2	23, 24	Default — ch 23
2	3	25, 26	Default — ch 25
3	1	31, 32	Default — ch 31
3	2	33, 34	Default — ch 33
3	3	35, 36	Default — ch 35

All 9 clusters carry two hydrophones.
How far each hydrophone hangs above its cluster is on the Geometry page.

OK
Cancel

Figure 12. The Far Field tab, configured for a 3120 cu.in source: gun positions imported from the array model, and the model's own figures entered as targets.

SETTING	WHAT IT IS
Method	Notional sources reconstructs what each cluster radiated and rebuilds a true far-field signature in bar·m; it needs gun positions. Summation merely adds the hydrophone traces — no geometry needed, but the amplitudes are not bar·m and cannot be compared with a model.
Window	How much of each NFH record is used, from its start. 512 ms is the whole record on a typical Gunlink NFH channel set.
NFH scale	Bar per recorded unit. 1.0 when the SEG-D already carries bar, which is the usual case.
Depth check	Measure each cluster's depth from the ghost notch in its own hydrophone trace, independently of the gun's depth transducer. Not a switch: it runs whenever the far field is on, reading the window the far field has already read, so it costs no extra disk access. Reported per cluster, not per gun , as a line median over every shot, and shown only against a depth baseline (below).
Depth tolerance	How far a cluster's depth may drift from its transducers since the depth baseline before the report highlights its row in light yellow. Fixed at 0.25 m, not in the Settings dialog. Nothing else follows from it: no finding and no cover verdict.
Take depth baseline	The ghost notch reads each cluster's depth with a bias from the array itself, so it is only shown against a depth baseline : a line you nominate, whose depth sensors you trust. It takes that line's sensors as truth, so choosing the line is the assertion. Each hydrophone's offset is measured from its own readings (one with fewer than 20 gets none) and each source's firing configuration is recorded; every line then shows drift since the baseline, on shots fired the same way. Nothing is applied until you confirm it — Apply and Cancel, Cancel the default. Clear discards the baseline. Until one is taken the depth check shows the sensors only.
Direct arrivals	Aligned (the default) sums every cluster's direct arrival at its own time zero, with its ghost $2z/c$ after it — how GunLink computes its far field. Delayed by depth delays each by its depth below the deepest cluster, as an array model presents it.
Signature polarity	Which way up the signature is published. SEG marine (primary negative) is the default and is the Bubble Test's convention — the SEG standard for marine seismic records a compression as a negative number on a hydrophone, so the primary is the strongest negative excursion, the surface ghost is positive and a bubble negative again. Positive draws it the other way up, which is how an array model usually plots it, for a side-by-side comparison. As recorded touches nothing. <i>No reported figure changes under any of them</i> : zero-to-peak is the largest absolute excursion, peak-to-peak a span, the spectrum a magnitude and the notch a frequency. Only the stored and plotted waveform is re-signed — which matters if the signature is used as a deconvolution operator.
Hydrophone to use	A cluster normally carries two hydrophones and Gunlink computes the far field from one of them, not from their average. Choose the odd channel of each pair, the even one, or both averaged. This is the default for every cluster of every source; the table below overrides it per cluster.
NFH height above cluster, odd and even	How far each hydrophone sits above its cluster's centre. The two are vertically separated — usually an internal and an external unit — so each has its own box, and a cluster is solved with the height of the hydrophone it reads. The computed amplitude scales with this distance, so it matters.
Speed of sound, sea-surface reflection	Take both from the recording system's setup — the values it computes its far field with. A coefficient quoted as a positive number for an inverting surface (for example 0.7) is entered negative. The defaults, 1500 m/s and -1 , are placeholders.

Gun positions	Import gun positions (CSV) reads a <code>gun,x,y,z</code> file in metres, <code>z</code> the gun depth, positive down. The gun table of the source's model report is exactly this.
Hydrophone per cluster	One row per cluster of every source, naming its real channel numbers. <i>Default</i> names the channel it resolves to on that cluster; set a row only where that cluster differs. A cluster is remembered by its channels, not by its position in the list, so switching a gun off does not move anyone's choice. A cluster carrying a single hydrophone has nothing to choose: its row is locked and names that channel, which is used whatever the default says. The note under the table says how many clusters are in that case.
Targets	A target and tolerance per metric, in per cent or in the metric's own units. Leave a Target box empty to skip that metric. Untick Report to leave a metric out of the report, the Stop log and the target checks altogether.
QC	Notional QC, optional and off until ticked — each shot PASS, WARNING or FAIL with its reasons. A cluster with no live hydrophone and the stability of the solve are always checked; the validation hydrophones (judged on shape) have their own tick box. Two heat maps in the End of Line report, each ticked on its own with its colour scale: every cluster's energy against its own line median, and its bubble period against the clusters of the same volume. Also whether FAILED shots are left out of the line's statistics. The defaults are a starting point, to be set from your own lines. See the Far-Field Guide.

If the hydrophone a cluster is set to use is dead on a shot, the app uses the other one of the pair — with that one's height — and says so once per source in the log. It does not leave the cluster silent: a notional source with no recording is worse than the wrong one of two.

If every cluster carries one hydrophone, set both height boxes to the same value. With one hydrophone per gun, consecutive channels are neighbouring guns rather than an internal/external pair, so splitting the height by odd and even would put every other cluster at the wrong depth. The note under the cluster table says which case the array is in.

There is a separate, fuller guide to this feature: *SEG-D Monitor — Far-Field Signature Guide* (`docs/SEG-D_Monitor_FarField_Guide.pdf`). It covers the files to gather, choosing the channel set, which hydrophone to read, what may carry a verdict, a worked example with real figures, and a troubleshooting table. This section is the summary.

Set targets only for what you actually want checked. Absolute levels depend on the hydrophone geometry and on filtering ([Appendix F](#)), so a level that disagrees with a model is a calibration question before it is a source problem. Shot-to-shot comparison does not depend on any of that.

7. Monitoring a Line

7.1 Preparation checklist

Before pressing Start

1. Open the correct project — the window title confirms it ([chapter 4](#)).
2. Check  **Settings**: the active channel set and the Source Configuration must match the deployed array ([chapter 6](#)). Click  **Test Nav Masks** and confirm the masks are what navigation will send.
3. Select the **SEG-D Folder** the recording system writes into.
4. Choose the direction, **Asc** ↑ or **Desc** ↓, and select or generate its reference under **Reference** ▾ ([§7.2](#)).
5. Set the autocorrelation window and the heatmap colour limits in the Display group.
6. Press  **Start**.

7.2 Generating a reference

The reference is the "golden" signature every shot is correlated against. **Generate from shots...** builds one by stacking a set of known-good SEG-D files.




Figure 13. Building a reference from known-good shots, and the Reference Viewer showing the result.

To generate a reference

1. Open **Reference** ▾, and in the row of the direction the reference is for, click ⋮ → **Generate from shots....**
2. Choose the **Input folder** holding known-good shots.
3. Check the **Output .npz** path — it defaults into the project's npz folder.
4. Set the autocorrelation window. The dialog reports which channels will be included and which nav masks it recognised.
5. Click **Generate**. Progress is shown, and the finished file is selected automatically.
6. Click **View Graph** (or ⋮ → **View** in the Reference group) to plot the reference traces and their autocorrelations before committing to it.

For a line shot the other way, build a second reference from that direction's shots, from that direction's row: each direction then uses its own, and a direction without one shares the other's.

The scan reports "*N/total files matched a known nav mask and contributed to the reference*". If that number is much lower than the number of files you selected, the Source Configuration's nav masks do not match the data — fix that first ([§6.4](#)).

7.3 What happens during monitoring

Every file appearing in the watched folder is queued, decoded, correlated against the reference, appended to the heatmap, streamed to the live graphs, stacked into the trace stacks, written as an NPZ, and written to the project database in the background.

Two properties of that pipeline are worth understanding, because they explain what you see in the log:

- **Files are processed one at a time, in shot-point order.** They are held in a single ordered queue, so a burst of files arriving while an earlier batch is still being decoded cannot be processed out of sequence.
- **The folder watcher is an optimisation, not the guarantee.** Correctness comes from a **reconciliation sweep** that re-scans the folder every 30 seconds and re-queues anything present on disk but not yet ingested. If the watcher misses a file, the sweep picks it up within the minute.

Misfired (non-triggered) guns are stored as well, so end-of-line problem detection sees them; the live trend graphs plot the guns that fired.

Depths decoded from the G_FIRE headers are converted from decimetres to metres once, at ingestion, so the graphs, the viewers and the reports all show metres consistently.

7.4 Reading the log

The log is the application's primary diagnostic surface, and the compiled build has no console — so what it says is all you get. Expand it with  **Log** ▾. Everything below is also written to the session log file ([chapter 13](#)).

WHAT YOU SEE	WHAT IT MEANS
[12/640] file.segd: nav mask 0x000F → Source 1 ✓	Normal. The file decoded, its mask resolved to a source, and its rows were written.

 Reconciliation sweep re-queued N file(s)	The 30-second re-scan found files on disk that had not been ingested and put them back in the queue. Occasional messages are normal; a steady stream means the watcher is not keeping up, which the sweep is compensating for.
 Drive stalled – folder listing ...	The filesystem did not answer. The application does not conclude that files are missing; it backs the sweep off from 30 s to 60 s to 120 s so it stops adding load to a drive that is already struggling, and resets on the first clean scan.
 Drive stalled – N file(s) unreadable this session	Files are present but could not be read. They are retried indefinitely, because on a rig the file is there and will eventually be readable. Check the mount; no shot is lost by waiting.
 Nav mask mismatch: file.segd – nav mask matches no configured source	The header decoded cleanly but named a source that is not configured. Retrying cannot fix this , so the file is rejected and the message is said once per file per session . Correct the Source Configuration (chapter 6). The shot is counted as <i>unreadable</i> , not <i>missing</i> .
 Source Configuration has N source(s) but computed 0 active nav masks	Every sub-array is inactive, or the sub-array IDs produce no bits. Nothing will be attributed to any source until this is fixed.
 N carried-over unreadable file(s) not re-checked	At Start the application re-probes files that failed in a previous session, within a 15-second budget. Whatever it did not reach keeps its previous verdict and the reconciliation sweep revisits it within the minute.
WARNING: ... journal mode ...	The project database is not really in WAL mode — the project is on a network share. See §4.4 .

A missing shot and an unreadable file are not the same thing, and the report never conflates them. A missing shot is a claim about the source — the gun did not fire. An unreadable file is an application or drive problem. If a file could not be read and its shot point could not be recovered from its headers, missing-shot detection is suppressed *for that line* and the report says so, rather than reporting a gap it cannot substantiate.

7.5 Inspecting headers

To inspect a file's decoded headers

1. Click  **Headers** in the Inspect group, or right-click a heatmap cell and choose **View Complete Headers**.
2. Select the SEG-D file.
3. Expand the tree: general headers, channel-set headers, extended and external blocks, and the per-trace G_FIRE gun data (depth, pressure, timing, volume).
4. Click **Export Headers to CSV** to save the decoded values.

This is the first thing to check when a file will not attribute to a source: the external header's nav mask is visible here, and can be compared directly against  **Test Nav Masks**.

7.6 Stopping

 **Stop** stops the watcher and flushes the database; the shot count is logged. If the End-of-Line report module is licensed, the report is then offered for the line or lines of this session ([chapter 12](#)). If it is not licensed the offer is skipped silently — monitoring itself is unaffected.

Recorded data, stacks, NPZ files and the session log all remain in the project.

7.7 Light Mode and the performance settings

 **Light** in the Monitoring group runs the cross-correlation and nothing else. It decodes the headers, reads the autocorrelation window and computes the similarity; it skips the full-trace read, the P/B ratio, the far field, the trace stacks and the per-shot NPZ file. The similarity values are identical to a full run — the mode changes what runs, never what is computed.

Use it when a line is arriving faster than the app can keep up with, or when cross-correlation is the only QC that line needs. What you give up, for that line, is **all gun QC**: no depth, pressure, timing, volume, P/B or far-field history, no trace stacks, no per-shot NPZ files and no End-of-Line report. The Graphs, Data, Database, SP Quality and Line Report buttons are greyed out while it is on, and the status reads **MONITORING · LIGHT**.

The toggle is **locked while monitoring runs** — the mode must not change in the middle of a line — and starting a line in Light Mode asks for confirmation once per project.

 **Settings** → **Performance** holds two more controls:

SETTING	EFFECT
Offer an End of Line report when monitoring stops (<i>on</i>)	Turn it off and  Stop skips the report prompt, the comments dialog and the cross-correlation snapshot they need. The database is still flushed, so no shot data is lost and  Line Report can still produce the report later.
Live graphs refresh every <i>N</i> seconds (2 s)	The longest a new shot can take to appear in the trend graphs. Lower is more responsive and costs more CPU.

Stop always waits for the processor to finish the file it is on and for every recorded shot to reach the database — that is what keeps the line's data intact — and it prints how long it took, phase by phase, in the log:  Stop took 3.10 s — processor 0.40 s · database flush 2.50 s · shot count 0.01 s · similarity save 0.19 s. If a Stop ever feels slow, that line says which part to look at.

7.8 Automatic mode

With **Automatic** ticked in the SEG-D folder panel ([§5.4](#)), lines start, stop and report themselves. It is off by default, is stored with the project, and can be switched while monitoring.

WHEN	WHAT THE APPLICATION DOES
A line is running and no new file has arrived for the quiet time	Stops the line and makes its End-of-Line report. The quiet time is 5 minutes unless you change it in the panel.
A report is being made	Waits for it. A line that begins meanwhile is caught up from its first file.
It is waiting, and a new line folder has three SEG-D files	Makes that folder the SEG-D folder, sets the direction, and starts monitoring. Three files are what it takes to tell which way the line runs.
It is waiting, and files arrive again in the folder just finished	Resumes that folder. The line is run again from its first file, so its data and its report are replaced by complete ones when it ends again.

What counts as quiet

The clock runs only when everything received has been processed and the drive is answering. A network drive that has stalled is not "no new files", and never ends a line; the quiet time is counted again from the moment the drive recovers. A line that has not given a single shot yet is not stopped for being quiet.

Which folders are lines

The folder that holds the line folders is watched: the one above the SEG-D folder selected — or the SEG-D folder itself, if what you selected there is the main folder where the lines are kept. The panel says which folder that is.

A folder in it is taken only if it is named like a line — digits, one letter, digits — and that letter is **any one** of the line letters shown in the panel, N, P, I and R unless you change them. So 1001P0001, 1002N0002, 1003I0003 and 1004R0004 are each started; a test line such as 1005T0005, an approach folder, or a folder named any other way is ignored. The letters are a list to choose from, not a word the name has to contain.

Old lines are never started. The line folders that are already there when the mode first looks are left alone. When you switch it on while a line is running, press **Start** for that line yourself; the lines after it follow by themselves.

Direction and reference

For a line it starts, the application reads the **shot point in the headers of the folder's first three files**. A shot point that **rises** as the line is shot is **Asc** ↑, one that **falls** is **Desc** ↓, and that direction's reference is then used ([§7.2](#)). The order of the shots is taken from the headers as well — the time of each shot — so it does not depend on how the files are named. Only the headers are read, never the traces.

If the headers cannot say — they could not be read, or carry no usable time — it keeps the direction selected, and the log says to check it. A line it cannot start — there is no reference for that direction, say — is written in the log and tried again a minute later.

Set a reference for each direction before leaving the lines to start by themselves. A line whose direction has no reference cannot be started automatically.

Nothing waits for a click

BY HAND	IN AUTOMATIC MODE
"Produce an End-of-Line report?"	Not asked.
Files that could not be read: Retry, Generate anyway, Cancel	Generated anyway; the log says how many files.
The comments box	Not shown. Comments are added afterwards (§12.3).
The progress window	None.
"Report created"	One line in the log, with the path.
The Light Mode question and the SP-increment box at Start	Not asked: Light Mode is used as set, and the SP increment is left as it is.

There is no report in Light Mode, or when the report is switched off in Settings → Performance.

Working alongside it

- **Start** and **Stop** work as always, with their questions. A line you stop yourself is left alone: automatic mode does not restart that folder, and waits for the next new one.
- While any dialog is open — Settings, a report's comments — automatic mode does nothing, and carries on when it is closed.
- Status reads **MONITORING · AUTO**, **AUTO · WAITING FOR A LINE** or **AUTO · REPORT**; the panel says which folder is being watched; every decision is a line in the log.

8. Live Trend Graphs

8.1 What the graphs show

The graphs panel shows per-shot trends while the line is being acquired: gun depth, pressure, firing time and timing error per gun, and fired volume and Peak/Bubble ratio per source, with the configured limits drawn as boundary lines. One panel is stacked per source, full window width each.



Figure 14. Live trend graphs — gun depth against shot point for three sources, with the tolerance band shaded and a gun drifting out of tolerance flagged on source 3.

8.2 Using the panel

CONTROL	FUNCTION
Metric	Switches every source panel together between Depth, Pressure, Timing, Volume, P/B, Summary, Far field and Firing spread. While the time-break sync rule is checked, Timing draws its window as purple dashed lines and rings the readings outside it; Firing spread plots each shot's first-to-last gun time against its limit.
View	Switches the Depth, Pressure and Timing graphs between Lines (one trace per gun) and Heat map (one row per gun, one column per shot point). The heat map uses the End of Line report's colours: blue-grey inside the limits, deeper as a reading nears its limit, red only outside the limits, grey where a gun has no reading. It covers the same shots and guns as Lines (Points, Full Range, Select Guns); hover a cell for gun, SP and value, double-click it for the shot detail. A metric whose SEG-D header carries no limit cannot be coloured and says so. The other metrics are not affected, and the choice is not kept between sessions.
Select Guns	Restrict the plot to chosen guns — useful when 24 traces per source is too many to read.
Full Range	Plot the whole line rather than the last <i>N</i> shots.
Points	How many recent shots to show when Full Range is off.

When the panel is opened it is rebuilt from the project database, so data recorded while it was closed appears immediately. It shows **only the lines of the current session** — a project holds every line ever monitored, and mixing them would put unrelated shot points in the same plot.

8.3 The gun × SP heatmap

Right-click a Depth, Pressure or Timing graph to open a heatmap of that metric for that source: one row per gun, one column per shot point.




Figure 15. Depth heatmap of one source. Gun 14 runs about a metre deep for a stretch of the line — the kind of pattern a per-gun trend line hides among 72 others.

Colour is each reading's deviation from expected, **scaled to the tolerance the graph itself draws**: white on nominal, fully blue or red exactly at the limit, deepening beyond it. So a drifting gun reads as a horizontal stripe and a bad stretch of line as a vertical one. Grey means that gun has no shot at that SP; it never reads as a good value.

The window covers the same shots the graph is showing (all of them with **Full Range** ticked), shows every gun including ones hidden on the graph, and keeps up with acquisition. Switch metric inside it, hover a cell for gun / SP / value, double-click a cell for the full shot detail, and use **Fit** to return to the whole line after zooming.

Only depth, pressure and timing have heatmaps. Volume and P/B have no per-gun tolerance band, so a deviation map of them would not mean anything.

8.4 The far-field page

With far-field QC switched on ([§6.9](#)), the Metric list gains a  **Far field** page: the far-field peak-to-peak of every shot, per source.




Figure 16. Far-field peak-to-peak along a line. The green band is the configured target and its tolerance; red marks shots that failed a target check. The dip is one source losing output for about fifteen shots.

What the app computes for each shot is the signature the array would produce far below itself, rebuilt from the near-field hydrophones — the same quantity an array model predicts, and the reason it can be compared with one. [Appendix F](#) sets out the mathematics. The figures behind the page are stored per shot, so the End-of-Line report can show the line as a whole ([§12.7](#)).

9. Trace Stacks

9.1 The Stacks window

 **Stacks** shows the live all-channels stack per source exactly as accumulated online: the same per-channel window and sample interval as during monitoring, taken from the stack's own metadata, with channel boundaries marked, real shot-point numbers on the horizontal axis, and a gain slider.



Figure 17. Live trace stacks, three sources side by side.

9.2 By Sub-arrays

 **By Sub-arrays** opens every sub-array of the selected source in one window, each in its own dockable panel. Panels can be dragged side by side, stacked vertically, tabbed, or floated as separate windows, and they share the source selector, gain and colour scheme. The sections are sliced from the existing all-channels stacks, so this changes nothing about online processing.




Figure 18. Sub-array stacks in a user-arranged layout. Colour schemes include Gray, Gray inverted, Seismic, Polarity, RdBu, Viridis, Bone, Copper and a Petrel-like map.

9.3 Detaching, zooming and fitting

To work with one sub-array in its own window

1. Drag the panel's title bar out of the window. It becomes a floating window with its own toolbar.
2. The remaining panels **keep refreshing** — detaching one does not pause the others.
3. Zoom with the toolbar's zoom tool. The zoom is preserved across refreshes, and is anchored to **shot points**, so it stays over the same part of the line even under "Traces: Last N", which re-slices the data every shot.
4. Click **Fit** to return to the full view.
5. Drag the window back onto the main one to re-dock it.


Use **Fit** rather than the matplotlib toolbar's Home button to get back to the whole line. The axes are cleared on every refresh, so Home's history no longer refers to the current data.

10. Viewers

All four viewers are read-only. They never write to the project.

10.1 Data Viewer

 **Data** is a filterable table of every stored row in the project.



The screenshot shows the Data Viewer interface. At the top, there are filter controls: 'Line' (set to 'all'), 'Source' (set to 'all'), 'Gun' (set to 'all'), and 'Guns' (set to 'All guns'). There is a checkbox for 'Include shot-level rows (gun 0)' which is checked. Buttons for 'Apply', 'Refresh filters', and 'Export CSV...' are present. Below the filters is a table with the following columns: line_name, shot_point, source_number, source_id, source, gun_number, sub_string, cluster_number, firing_time, depth, pressure, nominal_depth, and nominal_pressure. The table contains 23 rows of data, showing various shot points and gun numbers. At the bottom, there are navigation buttons 'Prev' and 'Next', a row count 'Rows 1-500 of 12160', and a 'Close' button.

	line_name	shot_point	source_number	source_id	source	gun_number	sub_string	cluster_number	firing_time	depth	pressure	nominal_depth	nominal_pressure
1	SEQ-0312-...	1001.0	1		Src 1	0							
2	SEQ-0312-...	1001.0	1		Src 1	1	1			6.169301648...	1951.2241202...	6.0	
3	SEQ-0312-...	1001.0	1		Src 1	2	2			6.0230112725...	1992.093935...	6.0	
4	SEQ-0312-...	1001.0	1		Src 1	3	3			5.846452093...	2021.984949...	6.0	
5	SEQ-0312-...	1001.0	1		Src 1	4	1			6.011885525...	2028.1810301...	6.0	
6	SEQ-0312-...	1001.0	1		Src 1	5	2			5.845327356...	2009.218769...	6.0	
7	SEQ-0312-...	1001.0	1		Src 1	6	3			6.1581210542...	1998.751852...	6.0	
8	SEQ-0312-...	1001.0	1		Src 1	7	1			0.0	0.0	6.0	
9	SEQ-0312-...	1001.0	1		Src 1	8	2			0.0	0.0	6.0	
10	SEQ-0312-...	1001.0	1		Src 1	9	3			0.0	0.0	6.0	
11	SEQ-0312-...	1001.0	1		Src 1	10	1			0.0	0.0	6.0	
12	SEQ-0312-...	1001.0	1		Src 1	11	2			0.0	0.0	6.0	
13	SEQ-0312-...	1001.0	1		Src 1	12	3			0.0	0.0	6.0	
14	SEQ-0312-...	1001.0	1		Src 1	13	1			0.0	0.0	6.0	
15	SEQ-0312-...	1001.0	1		Src 1	14	2			0.0	0.0	6.0	
16	SEQ-0312-...	1001.0	1		Src 1	15	3			0.0	0.0	6.0	
17	SEQ-0312-...	1001.0	1		Src 1	16	1			0.0	0.0	6.0	
18	SEQ-0312-...	1001.0	1		Src 1	17	2			0.0	0.0	6.0	
19	SEQ-0312-...	1001.0	1		Src 1	18	3			0.0	0.0	6.0	
20	SEQ-0312-...	1002.0	2		Src 2	0							
21	SEQ-0312-...	1002.0	2		Src 2	1	1			0.0	0.0	6.0	
22	SEQ-0312-...	1002.0	2		Src 2	2	2			0.0	0.0	6.0	
23	SEQ-0312-...	1002.0	2		Src 2	3	3			0.0	0.0	6.0	

Figure 19. Data Viewer — the raw project data, filterable and exportable.

To export a filtered subset to CSV

1. Set the **Line**, **Source** and **Gun** filters.
2. Untick **Include shot-level rows (gun 0)** if you want only per-gun rows. Gun 0 is the shot-level summary row, carrying volume and P/B for the whole source.
3. Click **Apply**.
4. Click **Export CSV....** The current filter is exported, not just the visible page.

Large projects are paged —  **Prev** and **Next**  move through them, and the row count is shown alongside.

10.2 Database Viewer

 **Database** has three tabs.

Gun performance

One gun's depth, pressure, firing time and error across the whole project, with line boundaries marked and average / σ statistics.



Figure 20. Gun performance across the project.

Line replay

Replays a line shot by shot with play/pause, adjustable speed and a scrub slider: per-gun depth and firing-time bars at the current shot point, errors highlighted, plus each gun's depth history up to that shot.




Figure 21. Replaying a line shot by shot.

Sources

Fired versus ideal volume and Peak/Bubble ratio per source.



Figure 22. Per-source volume and P/B ratio.

10.3 Seismic Viewer

 **Seismic** opens on the project's NPZ folder for the currently selected SEG-D folder, and can browse elsewhere. It has two display modes:

- **Shot points** — step through recorded shots with  /  or the arrow keys; traces are drawn as a filled wiggle display with a gain slider. The **Source** list (it starts on *All sources*) shows one source's shots only, so  /  step between that source's neighbouring shots; the source is read from each file's name, and the place on the line is kept when you change source. The same list narrows the trace stacks. **Use online AC window** pre-fills the display window with the monitoring autocorrelation window.
- **Trace stacks** — displays a stack applying the same window as online monitoring, from the stack's own metadata, with channel boundaries, real SP numbers, gain and the colour schemes listed in [§9.2](#).




Figure 23. Seismic Viewer stepping through recorded shots.

10.4 SP Quality

 **SP Quality** shows the contractual acceptance picture for one line: a horizontal track per source over a shared shot-point axis, red markers on bad shots, red spans over violations, green spans over acceptable ranges, and a selectable text summary of those ranges and the rule verdicts.




Figure 24. SP Quality. The header states the verdict and the bad shot count; the summary lists each rule with its limit.

To check a line against the acceptance criteria

1. Configure the criteria first in  **Settings** → **Shot Points** (§6.7). Until then the dialog notes that it is using default thresholds.
2. Open  **SP Quality** from the Analysis group.
3. Select the **Line**.
4. Read the verdict in the header and the per-rule results in the summary below.
5. Click **Refresh** after more shots have been acquired — this is a snapshot, not a live view.

On a multi-source line the green acceptable-range spans are not attributed per source, so where two sources have different acceptable stretches the spans can overlap indistinguishably. The dialog deliberately draws exactly what the acceptance engine reports, so that it can never disagree with the report.

10.5 The Refresh badge

Data Viewer, Database Viewer and SP Quality are **snapshots**. When new shots are committed while one is open, it shows a *new data* badge rather than redrawing itself.

This is deliberate. Each of these views recomputes a whole line, and doing that on every commit — roughly once a second during a busy line — is what made them slow. Click **Refresh** when you want the newer data. The live views (the heatmap, the graphs and the stacks) update on their own and need no such step.

10.6 Manage Lines

 **Manage Lines** lists every line in the project with its shot, row and NPZ-file counts, and permanently deletes the ones you select.




Figure 25. Manage Lines. The counts are read from the project database, so the confirmation quotes the real volume at risk.

Deleting a line removes its shot rows, its far-field rows and its per-shot NPZ files, and **cannot be undone** — the confirmation quotes the counts first. Lines are deleted one at a time, so if one fails the others are still

removed. A delete that cannot complete straight away (a busy database on a share) continues in the background, and that line's new shots are held back until it finishes, so a re-run can never mix old and new data.

The Data and Database viewers are read-only by design; deleting lives here instead, so that guarantee stays true.

11. Bubble Test (BBT)

The BBT window is independent of the main window — it stays open on its own and can be used while monitoring runs. It implements the complete bubble-test workflow: decoding, picking, evaluation against the Rayleigh–Willis relation, Pass/Fail against a configurable limit, storage in the project database, and a multi-page PDF report.




Figure 26. The Bubble Test window: file table, parameters, results and the trace plot.

11.1 Selecting files

To load a bubble test

1. Click **Select folder...** and choose the folder holding the bubble-test SEG-D files. They have the same structure and headers as production files.
2. Click **Scan files**. Every file is listed with its active gun, decoded from the G_FIRE headers.
3. Untick **Use** on any file to exclude it. By default the last *N* shots of each gun are ticked.
4. Where a gun or cluster has two hydrophones in the Source Configuration, choose which one to process in that file's **Hydrophone** selector.
5. Check the **BBT name** and **Date**. The name is decoded from the headers as FFID + T + HHMM (e.g. 3029T1805) and stays editable; the date comes from the header DateTime.

11.2 Parameters

PARAMETER	MEANING
Bubble pick window	From / To (ms), Min peak separation and Polarity: the time window searched for the primary and bubble peaks, with a minimum peak-separation guard. Saved with the project and shared with Settings → Hydrophone QC (§12.8) — change it in either place and both follow.
Limit (Δ%)	Pass/Fail threshold: a gun passes when $ \Delta\% \leq \text{limit}$.

Ideal Depth [m]	The nominal gun depth. Together with Ideal Pressure it defines each gun's theoretical k on its first test (§11.4). Locked for the project once the first k is saved.
Ideal Pressure [psi]	The nominal firing pressure, converted internally to bar (1 bar = 14.5038 psi).
Polarity	SEG marine polarity convention, respected when picking: a compression records negative on a hydrophone, so the primary and bubble are negative excursions. The far-field signature (§6.9) follows the same convention, so the two reports for one source cannot disagree.
Good shots (last N)	How many of each gun's most recent shots are ticked by default.
Seq from	How the sequence number is taken from the line name — Auto, or the first/last 3 or 4 characters. This drives both the report's Seq column and every line list in the application, so they can never disagree.

11.3 Picking and manual correction

Within the processing window the **primary** is the strongest positive peak; the **bubble** is the strongest local maximum at least the minimum separation after it, with light 1 ms smoothing to avoid ripple picks. Selecting a result row plots that gun's trace with both markers.




Figure 27. Manual bubble re-pick: clicking the plot snaps the marker to the nearest local maximum and recomputes that gun instantly.

To correct a mis-picked bubble

1. Select the gun's row in the results table so its trace is plotted.
2. Choose the shot with the **Shot (FFID)** selector if the gun has several.
3. Click on the plot at the correct bubble peak. The marker snaps to the nearest local maximum and that gun's numbers are recomputed immediately.

11.4 The computation

All evaluation follows the Rayleigh–Willis relation:

QUANTITY	FORMULA
Pressure conversion	$P_{\text{bar}} = P_{\text{psi}} / 14.5038$
Measured k (QC only)	$k_{\text{measured}} = T_{\text{meas}} \cdot (1 + 0.1 \cdot D)^{(5/6)} / (P_{\text{bar}} \cdot V)^{(1/3)}$
Computed bubble time	$T_{\text{calc}} = (P_{\text{bar}} \cdot V)^{(1/3)} / (1 + 0.1 \cdot D)^{(5/6)} \cdot k_{\text{used}}$
Deviation	$\Delta T = T_{\text{calc}} - T_{\text{meas}}; \Delta\% = \Delta T / T_{\text{calc}} \cdot 100$
Verdict	Pass if $ \Delta\% \leq \text{limit}$

Theoretical k is computed once per gun, from ideal conditions. On a gun's first bubble test its theoretical k is derived from that test's measured bubble period normalised to the Ideal Depth and Ideal Pressure:

$$k_{\text{theoretical}} = T_{\text{meas}} \cdot (1 + 0.1 \cdot D_{\text{ideal}})^{(5/6)} / (P_{\text{ideal_bar}} \cdot V)^{(1/3)}$$

The stored k is then fixed for the life of the project, and every test — including the first — evaluates T_{calc} with the gun's own *measured* depth and pressure. The first test therefore already yields a meaningful deviation rather than a trivial zero.

11.5 Saving, comments and saved tests

A comment can be attached to the test as a whole and edited per gun directly in the results table. **Save to DB** writes one row per gun, keyed by test name plus gun; the decoded traces are stored with the test so it can be re-processed later without the raw SEG-D files. The **Saved tests** tab browses stored tests filtered by test or gun number, with download and re-process options.

11.6 The BBT report

Report produces an A4 PDF in `<project>/bbt_reports/`:

1. **Title page** — client, survey, contractor and vessel, with the test dashboard.
2. **Gun Data Overview** — measured versus computed bubble time, $\Delta\%$ against the limit lines, and Δms , for every gun in the project at its most recent values. Guns tested in *this* report are highlighted; historical guns are faded.
3. **Bubble Period vs Chamber Volume** — the same all-guns dataset against the cube root of volume and against plain volume.
4. **Per-gun pages** — the hydrophone trace overlay of that gun's shots, Pass/Fail status, and its position highlighted in the array schematic drawn from the Designer layout.
5. **Cumulative Data Table** — every saved bubble test in the project sorted by gun and line, with the current test's rows shaded and failures marked.

11.7 Choosing k

k is the constant in the relation every gun is checked against ([§11.4](#)). It is not one number for every survey: it depends on the gun type and design, which set how fast the air is released, and on the water — temperature, sound velocity and salinity — and the result is very sensitive to it. The  beside **Default k** in the Bubble Test window shows this guidance in the application.

At the start of a project, before a k has been computed

1. Leave **Default k** at the customary figure for the gun type, or enter the manufacturer's. That is a global k: with it the calculated bubble period is only as good as about ± 50 cm of gun depth, so do not use it for an absolute depth check.
2. Set **Ideal Depth** and **Ideal Pressure** to what the array is meant to run at, make sure the guns really are at that depth and pressure, and run the first Bubble Test.
3. Read the **k meas** column. For guns of one type the values should agree closely. Take their average, leaving out the extremes: that is the project's k, and it holds for every gun to about ± 5 cm of depth.
4. A gun whose k stands apart from the rest has something wrong with its volume or its depth. Check the gun before trusting its k.

For guns of one type and design the bubble period should rise in a straight line with the cube root of the gun volume; a gun off that line is the same warning as a k that stands apart.

On a gun's first Bubble Test its k is computed from the measured bubble period at the Ideal Depth and Ideal Pressure and stored with the project; every later test of that gun is judged against that stored value. The first test therefore fixes k for the rest of the project.  **Edit k Values...** shows the stored table: use it to replace a k that came from a bad first test — with the average from step 3, say — or to set k before the first test.

12. Reports

12.1 The End-of-Line report

The End-of-Line (EOL) report is offered automatically when  **Stop** is pressed, for the line or lines of the session just finished, and is saved into <project>/line_reports/.

End of Line Report

Line SEQ-0312-P1804

Further investigation suggested

10 finding(s) on this line may warrant further investigation before the next line - see the summary below.



640

TOTAL SHOTS

18

ACTIVE GUNS

1h 46m

TIME TO SHOOT LINE

CLIENT	SURVEY	VESSEL	REPORT DATE
Demo Client	Demo Survey	Demo Vessel	2026-10-02 08:59
FIRST SHOT	LAST SHOT	SOURCES	LINES
06:00:00 (SP 1001)	07:46:30 (SP 1640)	Src 1, Src 2, Src 3	SEQ-0312-P1804

SUMMARY OF FINDINGS

SEVERITY	SUBJECT	CATEGORY	DETAILS	COUNT
CRITICAL	Source 2	Pressure	Pressure outside tolerance on 36 shot(s)	36
CRITICAL	Gun 14	Depth	Depth outside tolerance on 12 shot(s) (worst 7.44 m vs nominal 6.00 m at SP 1204)	12
CRITICAL	Gun 9	Misfire	Did not trigger on 9 shot(s)	9
MINOR	Gun 5	Timing	Exceeded timing tolerance on 4 shot(s) (worst +1.37 ms at SP 1604)	4
MINOR	Source 2	Volume	Fired volume deviated >5% from ideal on 3 shot(s) (worst 5.2% at SP 1014)	3
MINOR	Source 3	Volume	Fired volume deviated >5% from ideal on 3 shot(s) (worst 5.2% at SP 1189)	3
MINOR	Source 1	Volume	Fired volume deviated >5% from ideal on 2 shot(s) (worst 5.2% at SP 1160)	2
MINOR	Gun 16	Timing	Exceeded timing tolerance on 1 shot(s) (worst +1.04 ms at SP 1591)	1
MINOR	Gun 17	Timing	Exceeded timing tolerance on 1 shot(s) (worst +1.14 ms at SP 1597)	1
MINOR	Gun 11	Auto-fire	Auto-fire detected on 1 shot(s)	1

Figure 28. The report cover: verdict banner, KPI cards, line metadata and a ranked summary of findings.

The report contains:


1. **Cover** — an *All clear / Minor issues / Critical issues found* banner, KPI cards (total shots, active guns, time to shoot the line), the client / survey / vessel metadata, first and last shot, and **every finding** — automatically detected timing, depth, pressure, volume, P/B deviation, misfire and auto-fire events — ranked with the most severe first, each with its count. A long list continues on *Summary of Findings (continued)* pages straight after the cover; there is no separate problems page. A line marked NTBP ([§12.3](#)) carries an **NTBP - NOT TO BE PROCESSED** strip above the banner.
2. **Timing performance** as a heat map per source: one row per gun, one column per shot point.
3. **Gun firing synchronisation**, when a firing rule is checked: each shot's gun furthest from the time break, and its first-to-last spread, against their limits.
4. **Depth and pressure**, heat maps of the same kind.
5. **Volume and Peak/Bubble ratio** per source. The volume chart plots **cu.in.**, not percentage of ideal — a percentage is a flat line on a healthy line and shows nothing.
6. **Per-gun scorecard** — each gun's averages, with *SEE FINDINGS* on a gun the findings name.
7. **Spec-check verdict**, when the spec check is enabled ([§6.7](#)).

Timing Performance by Source

3 source(s)
Line **SEQ-0312-P1804** · 2026-10-02 08:59


Source 1 **OUT OF TOLERANCE**

4 of 1284 readings outside limits · worst +0.37 ms (G5, SP 1604)




Source 2

all 1269 readings inside limits · nearest a limit: 86 % of it (G8, SP 1173)



Source 3 **OUT OF TOLERANCE**

2 of 1278 readings outside limits · worst +0.14 ms (G17, SP 1597)



nominal limit 2 x limit

Blue-grey: inside the limits, deeper as a reading nears its limit. Red: outside the limits - the readings counted on the cover and in the Spec Check - deeper further out, darkest at twice the limit distance. White: no reading (the gun did not fire, or reported no value).

Figure 29. Timing as heat maps, one row per gun: blue-grey inside the limits, deeper toward them; red only where a reading broke its limit. Here one gun drifts toward its limit and then breaks it.

Pages and blocks that would only say "nothing is wrong" are omitted. The cover's verdict banner is the single place that states an all-clear.

Reading the timing, depth and pressure heat maps. Each cell is one gun on one shot. **Red means outside the limits, and nothing else:** exactly the readings the findings and the Spec Check count, judged against the limits in that shot's own SEG-D header, deeper red the further out. Inside the limits a cell is blue-grey, pale at nominal and deeper as the reading nears its limit, so a gun drifting toward a limit shows before it breaks it. White is no reading — the gun did not fire. A gun going wrong reads as a horizontal stripe, a bad stretch of line as a vertical one. Above each map: how many readings fell outside, and the worst with its gun and shot point — or, on a clean source, how near any reading came to a limit. Depth is measured from nominal against the shallow or the deep tolerance, whichever side the reading is on; pressure from the middle of the low-high band.

12.2 Generating a report for any stored line

To produce a report without a monitoring session

1. Open the project holding the line.
2. Click  **Line Report** in the Analysis group.
3. Choose the line or lines.
4. Fill in the report metadata and comments if prompted, then confirm.

This works entirely offline from what is already in the database, so a report can be regenerated at any time — after changing the branding, for example, or to reissue one for a client.

It also works **while monitoring**, for every line except the one being recorded: that line is left out of the list until it has stopped.

One report per line. When a new report has been written, the older reports of the same line are removed from `line_reports`. An older one that is open in a viewer cannot be removed then; the log says so, and it goes with the next report.

12.3 Comments

Before the report is generated you are offered a Comments box. Text entered here appears in the report. It supports bold, italic and colour, and is spell-checked against an offline English dictionary — no network is involved.

Each line keeps its own comments. Reopen a line's report and its comments are in the box; a line that has none yet opens with the last text used. This is how a report made in automatic mode ([§7.8](#)) gets its comments: click  **Line Report**, choose the line, write, and confirm. The report is made again with them and replaces the one without.

The same box has a **Line status** row for each line of the report. Tick **NTBP - not to be processed** and, if you wish, type a short reason; **OK** stores it with the line, and the report's first page then carries an **NTBP - NOT TO BE PROCESSED** strip with the reason, above the verdict; in a report over several lines the strip names each NTBP line. The verdict and the findings are unchanged. Untick it in a later report to clear it. Deleting the line in Manage Lines removes its status; re-running the line keeps it.

12.4 Branding

Optional - add up to 3 logos (your company, the client, the contractor/vessel) to brand this project's reports. Leave empty and reports simply show no customer logos.

Recommended: PNG with a transparent background, about 400x160 px (a ~2.5:1 landscape lockup). Square badges work too - every logo is scaled to a uniform height and placed under the line/test name in the report header.

The "Generated by SEG-D Monitor" mark always remains on every page for attribution.

Logo 1: — none — Choose... Clear

Logo 2: — none — Choose... Clear

Logo 3: — none — Choose... Clear

Save Close

Figure 30. Up to three customer logos per project.

To add customer logos to this project's reports

1. Click  **Branding...** in the Setup group.
2. Click **Choose...** beside Logo 1, 2 or 3 and select an image. A PNG with a transparent background at roughly 400x160 px works best; square badges are fine too.
3. Click **Save**. The logos appear in the header of every report this project generates.
4. **Clear** removes a logo; leaving all three empty simply produces reports with no customer logo.

The "Generated by SEG-D Monitor" mark always remains on every page for attribution.

12.5 Missing shots and unreadable files

The report distinguishes two things that look similar and mean opposite things:

REPORTED AS	MEANING
Missing shot	A shot point that should exist on the line's grid and does not. This is a statement about the source .
Unreadable file	A file that is present on disk but which the application could not read or could not attribute to a source. This is a statement about the application or the drive .

Where missing shots are reported. Among the **findings** on the cover, as a critical row naming how many shots are missing, across how many gaps, and at which shot points (a long list is truncated, but the count beside it is always exact). There is no separate page for them: one repeated the other, and the shot points are on the finding.

Where a file could not be read but its shot point was recovered from its headers, that SP is subtracted from the gap it falls in, so it is not double-counted as missing. Where the shot point could not be recovered at all, the file cannot be placed on the grid — and rather than report a gap it cannot substantiate, the report **suppresses missing-shot detection for that line** and states that the count is unverified, with the number of unreadable files.

12.6 How reports are rendered

Reports are built as HTML and rendered to PDF, trying an embedded browser engine first and falling back to a second renderer if it is unavailable in your build. **A log message about the fallback is normal, not an error** — the content and layout are identical either way.

Rendering a long report takes time. A progress dialog shows the current stage and can be cancelled, and there is a hard timeout so a stuck render can never hang the application.

12.7 The far-field pages

When a line has far-field data ([§6.9](#)), the report gains pages for it, each carrying one kind of thing:

PAGE	CARRIES
Far-Field Signature	The KPIs and every metric against its target, per source.
Levels along the line	Peak-to-peak and zero-to-peak against shot point, with their target bands, and the assumptions the figures rest on.
Median Signature & Spectrum	The line's median waveform and its amplitude spectrum, with the ghost notch marked.
Depth from Ghost Notch	Every cluster with its gun numbers, sensor depth against notch depth, whatever the shot count. Information only: it raises no problem, reaches no cover verdict and never counts a shot point bad. A row is highlighted when the two depths differ by more than 0.25 m. Appears only when the depth check is on.
Shots Outside Target	The shots behind the failing-target count. Appears only when something fell outside one.
Notional QC	Only where notional QC was on: PASS, WARNING and FAIL counts with their reasons per source, then one row per cluster, its bubble period beside its same-volume peers' and its guns' Bubble Test. Information only. FAILED shots are left out of the other far-field pages' figures unless the QC page of Settings says otherwise.

A line without far-field data gets the report it always got.

Far-Field Signature

 SEQ-0312-P1804
 Line **SEQ-0312-P1804** · 2026-10-02 08:59

120 SHOTS WITH A FAR FIELD		8 SHOTS FAILING A TARGET		Notional sources METHOD
METRIC	TARGET	SOURCE 3 (ID 1)	SOURCE 3 (ID 2)	SOURCE 3 (ID 3)
Peak-to-peak (bar-m)	60 ±5 %	59.0 (57.2–60.9) -1.7% vs target 38 of 40 shots within 2 outside	58.9 (49.3–60.8) -1.8% vs target 34 of 40 shots within 6 outside	59.2 (57.6–60.5) -1.4% vs target 40 of 40 shots within
Zero-to-peak (bar-m)	–	35.6 (34.6–36.8)	35.6 (29.8–36.7)	35.7 (34.8–36.5)
RMS, 2 s window (bar-m)	–	1.96 (1.90–2.02)	1.96 (1.64–2.02)	1.97 (1.92–2.01)
Peak frequency (Hz)	–	73.0 (72.4–73.6)	72.8 (72.2–73.3)	73.0 (72.6–73.4)
–3 dB bandwidth (Hz)	–	45.0 (45.0–45.0)	45.0 (45.0–45.0)	45.0 (45.0–45.0)
–6 dB bandwidth (Hz)	–	62.0 (62.0–62.0)	62.0 (62.0–62.0)	62.0 (62.0–62.0)
10–70 Hz maximum (dB)	–	206.4 (206.4–206.4)	206.4 (206.4–206.4)	206.4 (206.4–206.4)
10–70 Hz average (dB)	–	203.9 (203.9–203.9)	203.9 (203.9–203.9)	203.9 (203.9–203.9)
10–70 Hz ripple (dB)	–	7.6 (7.6–7.6)	7.6 (7.6–7.6)	7.6 (7.6–7.6)
Ghost notch (Hz)	95.1 ±5	94.0 (94.0–94.0) -1.1 vs target 40 of 40 shots within	94.0 (94.0–94.0) -1.1 vs target 40 of 40 shots within	94.0 (94.0–94.0) -1.1 vs target 40 of 40 shots within

Line median (5th–95th percentile), then how many shots met the target. A row is shaded when the MEDIAN falls outside its target. Measured at the channel set's own sample rate. A level that disagrees with a model is a calibration question before it is a source fault: it scales with the hydrophone distance, which is a mounting dimension. P/B and bubble period are not offered as model targets: a tuned, clustered array has no clear bubble.

Figure 31. Every metric against its target: the line median, the 5th–95th percentile spread, and the share of shots inside tolerance, one column per source. A row is shaded when its median falls outside the target.


Far-Field Signature — Levels along the line

SEQ-0312-P1804
Line SEQ-0312-P1804 · 2026-10-02 08:59

Far-field peak-to-peak along the line


green = target, dashed = tolerance

● Source 3 (id 1) ● Source 3 (id 2) ● Source 3 (id 3)



Far-field zero-to-peak along the line

● Source 3 (id 1) ● Source 3 (id 2) ● Source 3 (id 3)



Method Notional sources · the active channel set · 512 ms window · odd hydrophone

Geometry hydrophone 1 m (odd) / 1 m (even) above cluster · NFH scale 1 bar per unit · gun depths measured per shot

Sea c 1500 m/s · reflection -1 (a flat sea, the ideal)

Output signature published with SEG marine polarity (a compression records negative, so the primary is negative and the surface ghost positive) - as in the Bubble Test · vertical, at infinite distance · direct arrivals aligned · measured at the set's own rate, stored at 2 ms

Absolute levels depend on the hydrophone height and on filtering - compare with a model filtered the same way.

Figure 32. The median signature of the line and its amplitude spectrum, with the ghost notch marked. The notch position is a direct check of the gun depths the shots were fired at.

The assumptions the figures rest on — method, NFH scale and height, sound speed, reflection — are printed under the table, because the absolute level depends on them.

Shots Outside Target. The first page's *shots failing a target* figure is a count; this page is the shots behind it, straight after the graphs. One row per target missed, naming the shot point, the source, the metric, what was measured, the target, and the deviation from it. A shot that missed several targets is listed once with a row per miss. The page appears only when there is something on it — no target set, or nothing outside one, and it is absent.

Read a *level* miss (peak-to-peak, zero-to-peak, RMS) differently from a *shape* miss. A level scales with the hydrophone mounting distance, so a whole line sitting below its target is a calibration question before it is a source fault. The ghost notch is geometry and has no such excuse: a notch away from its expected frequency means the guns were not at the depth the model assumed.

12.8 The hydrophone QC pages

Settings → *Hydrophone QC* (off by default) measures, on every shot in normal mode, **each hydrophone of the active guns** — every NFH channel an active gun lists in the Source Configuration, both of a pair — on its raw near-field trace: peak-to-peak, zero-to-peak and bubble period. Light Mode measures nothing, since it never reads the traces.

SETTING	MEANING
Channel set	<i>Active set</i> reuses the traces already read for P/B and costs no extra read; another set is read over the window only.
Bubble pick window	The Bubble Test's window (§11.2) — one project setting shown in both places. A window running past the channel set's record is measured to the end of the record, and the log says so once.
Heat-map scales	Full colour at $\pm$ this much from each hydrophone's own line median, one per metric (for example $\pm 20\%$ for the levels).

How the bubble is picked. The primary is the strongest excursion of the chosen polarity, exactly as in the Bubble Test. The bubble is the crest of the *first* excursion reaching half the largest one after the minimum separation — not the largest. A Bubble Test fires one gun alone; in production every gun fires, and a hydrophone also hears its neighbours' bubbles, later and often larger than its own gun's.

In the report, for a line that has these data: per source, one page of three heat maps — rows the hydrophones in channel order, columns the shot points — coloured by each shot's difference from that hydrophone's own line median: white at the median, full blue or amber at the scale, grey where there is no value (a dead channel, or no bubble inside the window). A gun drifting or dropping out reads as a horizontal stripe; a bad stretch of line as a vertical one. A table of each hydrophone's guns, shot count and medians follows. Levels are in the recording's own units. The page header names the window the line was measured in.

Hydrophone QC is **information only**: it judges no shot, raises no problem, reaches no cover verdict and never counts a shot point bad in the Spec Check.

13. Session Logs and Diagnostics

13.1 The session log file

Every time you press  **Start**, a new log file is opened at

```
<project>/logs/segd-monitor_<date>--<time>.log
```

and every message from the on-screen log is written to it, together with a header recording the application version, the watched folder, the database journal mode, the PDF renderer in use and the fontconfig path. The 30 most recent session logs are kept.

BUTTON	FUNCTION
 Logs	Opens the folder holding the session logs.
 Save Log	Writes what is currently on screen to a text file of your choosing.

The on-screen log is capped and shows the recent tail; the file holds the whole session.

13.2 Reporting a problem

To report a problem usefully

1. Note the exact button or action, and the text of any error dialog.
2. Click  **Logs** and take the session log for the run in question.
3. Send both. If the problem concerns a specific line, say which one and which shot points.

Always send the session log file, not just a screenshot of the log panel. The file carries the startup header — version, paths, journal mode, renderer — which is exactly the information the on-screen log does not repeat, and which usually identifies the problem immediately.

14. Backup, Transfer and Multi-Workstation Use

- **Backup is a folder copy.** The project folder contains the database, every NPZ, the stacks, the logs and the reports. Nothing else is needed to restore or move a survey.
- **Transfer to another machine:** copy the project folder and open its database with  **Project....** The license is per machine — the target needs its own license for the professional modules, though its FREE monitor opens any project.
- **Starting a new survey pre-configured:** create a new project and use *Import settings from* to copy the channel sets and Source Configuration from the previous survey ([§4.2](#)).
- **Simultaneous access:** the viewers can read while the monitor writes *on the same machine*. Do not have two machines write into the same project folder over a network share, and see [§4.4](#) before putting a live project on one at all.

15. Troubleshooting and FAQ

15.1 Licensing

SYMPTOM	RESOLUTION
A button shows  and reports a licensed feature	The license is missing, expired, or does not include that module. Open  License to check; §3.2 describes obtaining one.
“This license is locked to machine ...”	The <code>.lic</code> was issued for a different computer. Send <i>this</i> machine's ID (shown in the same dialog) for a replacement.
“License signature is invalid”	The file was modified, truncated in transfer, or belongs to a different product build. Re-copy the original exactly as received; if it persists, request a re-issue.
“License expired on ...”	Request a renewal — a new <code>.lic</code> for the same Machine ID — and load it over the old one.
“System clock inconsistency detected”	The system date was moved backwards by more than 48 hours. Correct the clock; the license resumes automatically.
License asked for again after moving the application folder	The application folder does not store the license — the user profile does (§3.4). A different OS user loads the <code>.lic</code> once as well.

15.2 Monitoring and ingest

SYMPTOM	RESOLUTION
No files are picked up	Verify the SEG-D folder path, that the recording system is writing into it, and that the files carry the expected extension. The reconciliation sweep re-scans every 30 s, so a file present on disk should appear within a minute.
Files stop appearing; the log shows repeated retries or “Drive stalled”	The drive or share is not answering. The application retries indefinitely and recovers on its own — no shot is lost by waiting. Check the mount. See §7.4 .
A file is reported once as a nav-mask mismatch and never retried	Its header names a source that is not configured, and retrying cannot change that. Fix the Source Configuration (chapter 6) and re-run the line.
Every shot reports a nav-mask mismatch	The sub-array IDs do not match what navigation sends. Use  Test Nav Masks (§6.4).
The application is slow to start or stop monitoring	The project database is on a network share and is not really in WAL mode. Move the project to a local disk (§4.4).
Heatmap appears but the graphs are empty	The graphs show only the current session's lines. If you have just started, they fill as shots arrive.

15.3 Reports

SYMPTOM	RESOLUTION
Report generation fails; the log mentions pango, GLib, or a library that will not load	The graphics libraries the PDF renderer needs were not found. Set the environment variable <code>SEGD_GTK_DIR</code> to the folder holding them and restart.

Windows: “Fontconfig error: Cannot load default config file”	The build is missing its bundled font configuration. Set <code>SEG_D_FONTCONFIG_DIR</code> , or obtain a build that includes it. The startup header in the session log reports the fontconfig path in use.
A message about falling back to a second PDF renderer	Normal, not an error (§12.6).
The report says missing shots could not be verified	A file could not be read and its shot point could not be recovered, so detection was suppressed for that line rather than reporting a guess (§12.5).
No report is offered on Stop	The End-of-Line module is not licensed; the offer is skipped silently. Monitoring and the stored data are unaffected, and the report can be generated later once licensed.

15.4 Displays and viewers

SYMPTOM	RESOLUTION
A viewer shows a “new data” badge but no new rows	These are snapshots by design — click Refresh (§10.5).
A sub-array panel seems frozen	Detaching a panel does not pause the others; if a panel is not updating, check that the stack file is still being written. Use  Fit to return to the full view (§9.3).
A window will not reopen	Most secondary windows are singletons. If one seems stuck, close it from its own window controls before reopening it.
The toolbar is cramped on a small screen	The toolbar scales its font down automatically; the Inspect, Analysis, Display and Setup groups also collapse to chips (§5.3).

Appendix A. Folder and File Reference

PATH	CONTENTS
<application folder>	The compiled application. Replaceable as a whole on upgrade; stores nothing of yours.
<application folder>/license.lic	Optional: a license placed here is loaded automatically, which suits centrally prepared workstations.
Windows: %APPDATA%\XCorrSEGD\	Per-user: the active license and the clock-guard state.
Linux: ~/.config/XCorrSEGD/	Per-user: the active license and the clock-guard state.
macOS: ~/Library/Application Support/XCorrSEGD/	Per-user: the active license and the clock-guard state.
<project>/	Everything belonging to one survey — see §4.3 for the full layout, including <code>logs/</code> , <code>npz/</code> , <code>line_reports/</code> and <code>bbt_reports/</code> .

Environment variables

VARIABLE	EFFECT
SEGD_GTK_DIR	Overrides where the PDF renderer looks for its graphics libraries, on every platform.
SEGD_FONTCONFIG_DIR	Overrides the font-configuration directory. Windows only.

Appendix B. Bubble-Test Formula Summary

QUANTITY	FORMULA / MEANING
Pressure conversion	$P_{\text{bar}} = P_{\text{psi}} / 14.5038$
Measured k (QC only)	$k_{\text{measured}} = T_{\text{meas}} \cdot (1 + 0.1 \cdot D_{\text{meas}})^{(5/6)} / (P_{\text{meas_bar}} \cdot V)^{(1/3)}$
Theoretical k (first test of each gun, then fixed)	$k = T_{\text{meas}} \cdot (1 + 0.1 \cdot D_{\text{ideal}})^{(5/6)} / (P_{\text{ideal_bar}} \cdot V)^{(1/3)}$
Computed bubble time	$T_{\text{calc}} = (P_{\text{meas_bar}} \cdot V)^{(1/3)} / (1 + 0.1 \cdot D_{\text{meas}})^{(5/6)} \cdot k_{\text{used}}$
Deviation	$\Delta T = T_{\text{calc}} - T_{\text{meas}}; \Delta\% = \Delta T / T_{\text{calc}} \cdot 100$
Verdict	Pass if $ \Delta\% \leq \text{limit}$

T = bubble period [ms]; D = depth [m]; P = pressure; V = chamber volume [cu.in.]; k_{used} = the gun's stored theoretical k , falling back to the project average and then to the manual constant.

Appendix C. Spec-Check Rules

Configured in  **Settings** → **Shot Points** ([§6.7](#)); off by default. A shot counts as **bad** when at least the configured number of its guns fail any selected criterion. A nav-mask mismatch always marks the shot bad.

RULE	EVALUATED
Maximum consecutive bad shots per source	Live during the line, and in the report
Window 1 — at most n bad within m shots	Live during the line, and in the report
Window 2 — a second, wider window	Live during the line, and in the report
Maximum bad shots overall (%)	Live during the line, and in the report — but see the note

The percentage rule behaves differently from the other three. They latch: a breach somewhere on the line stays a breach, because it is a local fact that later good shots do not undo. The percentage is a whole-line figure and **recovers** — a line that starts badly and then runs clean is not a failing line, so the warning clears again.

It also stays silent early. A partial line's percentage is not the line's: one bad shot in the first 32 is 3.1%, over a 3% limit, on a line that finishes at 2.0%. The rule says nothing until the line is long enough that a single bad shot can no longer exceed the limit by itself — 34 shots at a 3% limit — which is derived from the limit you set, not a fixed number.

A gun that is switched off is **not** a bad shot. It is reported among the report's findings, but it fails no contractual criterion: a line once failed its spec because one 250 cu.in gun of a 3120 cu.in source stopped firing. What does count is the SOURCE firing nothing at all.

A shot that breaks several rules at once is counted **once**. The report's Violations table lists one row per span with every rule that span broke, so the "# Bad" column adds up.

Missing shot points are injected as bad shots on the full shot-point grid, so the sliding windows advance with the acquisition one SP interval at a time rather than skipping over gaps.

Appendix D. Keyboard and Mouse Reference

WHERE	ACTION
Anywhere	F1 — open the built-in manual
Manual window	Ctrl + F — find; Esc — close
Wiggle / stack viewers	F11 — full screen; Esc — exit full screen
Heatmap	Hover — value read-out; right-click — open the traces or headers of that shot and channel
Seismic Viewer (shots)	 /  or  /  — previous / next shot; gain slider — display amplitude
BBT trace plot	Click — manually re-pick the bubble; snaps to the nearest local maximum
Database Viewer — line replay	Play / pause, speed control, scrub slider
Sub-array stacks	Drag panel title — re-dock; drag out — float as its own window; drop on a tab bar — tab panels together;  Fit — return to the full view

Appendix E. Glossary

TERM	MEANING
SEG-D 1.0 (Gunlink)	The field-recording format written by the gun controller; carries trace data plus G_FIRE gun headers (depth, pressure, timing, volume) per shot.
NFH	Near-field hydrophone — the sensor recorded per gun or cluster and processed by this application.
Reference NPZ	The stored "golden" signature the monitor correlates every shot against.
Autocorrelation window	The millisecond window in which similarity is evaluated.
Nav mask	A bit field in the SEG-D external header saying which sub-arrays fired. Sub-array N is bit $(N-1)$ (§6.4).
P/B ratio	Peak-to-Bubble ratio of the source signature.
Primary / bubble	The first (strongest) peak of the airgun signature and the following bubble oscillation peak; their spacing is the bubble period.
Rayleigh–Willis	The relation linking bubble period to pressure, volume and depth; the basis of the bubble-test evaluation.
Theoretical k	Per-gun constant of the Rayleigh–Willis relation, fixed on the gun's first test from ideal conditions (§11.4).
Reconciliation sweep	The 30-second re-scan of the watched folder that re-queues anything on disk but not yet ingested. It, not the folder watcher, is what guarantees no file is dropped (§7.3).
Unreadable file	A file present on disk that could not be read or could not be attributed to a source. Distinct from a missing shot (§12.5).
Spec check	Contractual bad-shot acceptance checking (Appendix C).
Machine ID	16-character code identifying a computer for licensing; shown in the License dialog.
Project	One folder = one survey: database, NPZ files, stacks, logs and reports.

Appendix F. Far-Field Signature Mathematics

This appendix sets out what the far-field QC ([§6.9](#), [§8.4](#), [§12.7](#)) computes. The method is the notional-source method of Ziolkowski, Parkes, Hatton and Haugland (1982) and Parkes, Ziolkowski, Hatton and Haugland (1984).

F.1 What a near-field hydrophone actually records

A hydrophone about a metre above its own cluster does not record that cluster alone. It records every cluster in the source, each weakened by the distance it travelled and delayed by that distance divided by the speed of sound — and it records each one twice, because the sea surface reflects a near-perfect mirror image back down. Writing q_j for the signature cluster j alone radiated (its **notional signature**), hydrophone i records

$$h_i(t) = \sum_j q_j(t - r_{ij}/c) / r_{ij} + R \cdot \sum_j q_j(t - g_{ij}/c) / g_{ij}$$

SYMBOL	MEANING
r_{ij}	Distance from hydrophone i to cluster j [m]
g_{ij}	Distance from hydrophone i to cluster j 's mirror image above the sea surface — the ghost path [m]
c	Speed of sound in water [m/s], from the settings
R	Sea-surface reflection coefficient, -1 for a flat sea
q_j	Cluster j 's notional signature [bar·m] — the unknown

Guns that share NFH channels form one cluster, and a cluster is treated as **one** notional source: a clustered pair's bubbles merge, so they do not radiate independently.

A cluster often carries **two** hydrophones, vertically separated — usually an internal and an external unit — and the far field is computed from **one** of them, as Gunlink computes it, not from their average ([§6.9](#)). The choice is part of the geometry, not only of the bookkeeping: the two hang at different heights, so h_i for a cluster reading its odd hydrophone and for one reading its even hydrophone are at different depths, and every r_{ij} and g_{ij} in that row of the matrix changes with it. If the selected hydrophone is dead on a shot, the other one of the pair is used — with *its* height — and the log says so once per source. Where a cluster carries only one hydrophone, or the array carries one per gun, that hydrophone is used and its own height applies; nothing above changes, because the equations are written per cluster and never assume a pair.

F.2 Solving for the notional signatures

With n clusters there are n equations and n unknowns at every instant. Taking the Fourier transform turns each delay into a phase, and the system becomes one small matrix equation per frequency:

$$H(f) = A(f) \cdot Q(f), \quad A_{ij}(f) = e^{-i2\pi f \cdot r_{ij}/c} / r_{ij} + R \cdot e^{-i2\pi f \cdot g_{ij}/c} / g_{ij}$$

$$Q(f) = A(f)^{-1} H(f)$$

which is solved frequency by frequency and transformed back. Only the band the array can resolve is solved, under a smooth taper — full weight to 600 Hz, nothing above 900 Hz (or the top third of the band of a set whose Nyquist is lower) — so the upper bins, where the system is least well conditioned, are skipped. A hard cut-off would ring; the taper does not.

Why not solve it in the time domain? The same equations can be marched sample by sample, feeding each cluster's recovered past back into its neighbours, and that is the classical formulation. On a real array it can be unstable: for a source with 3 m cluster spacing, 10 m between strings and hydrophones a metre up, the feedback loop has a gain above one near 48 Hz, and measured data — which always carries noise and a little geometry error — then grows without bound. Solving per frequency has no feedback loop in time, so there is nothing to grow.

F.3 From notional signatures to the far field

The far field is what the whole array radiates as heard from far below it, scaled to a distance of 1 m. Each cluster contributes its own signature and its ghost, the ghost $2 \cdot z_j / c$ after it:

$$F(t) = \sum_j [q_j(t - d_j) + R \cdot q_j(t - d_j - 2 \cdot z_j / c)]$$

where z_j is cluster j 's depth **as measured on that shot**, taken from the gun headers. With *Direct arrivals* on **Aligned** — the default, and the recording system's convention — $d_j = 0$: every direct arrival at its own time zero. On **Delayed by depth**, an array model's convention, $d_j = (z_{\max} - z_j) / c$, with time zero on the deepest cluster.

Because the direct arrival and its ghost are separated by the two-way travel to the surface, they cancel at

$$f_{\text{notch}} = c / (2 \cdot z)$$

— for example about 94 Hz for guns at 8 m in water at 1500 m/s. This **ghost notch** is the most useful self-check in the whole calculation: it is fixed by physics, so a computed signature whose notch sits where the measured gun depths put it has its geometry, its depths and its ghost right.

F.4 The figures the app reports

FIGURE	DEFINITION
Peak-to-peak	$\max F - \min F$ [bar·m]
Zero-to-peak	$\max F $ [bar·m] — the absolute peak, so the hydrophone polarity convention cannot change it
RMS	$\sqrt{(\text{mean}(F^2))}$ over the window
RMS (2 s window)	$\text{RMS} \cdot \sqrt{(T_{\text{window}} / 2 \text{ s})}$ — RMS depends on the window it is taken over, and models quote it over 2 s. Exact while all the energy falls inside the window, as it does
Amplitude spectrum	$X(f) = \text{FFT}(F) \cdot \Delta t$ [bar·m/Hz], shown as $20 \cdot \log_{10}(X \cdot 10^{11})$ dB re 1 $\mu\text{Pa}/\text{Hz}$ at 1 m (1 bar = 10^{11} μPa), on a 2 s frequency grid
Peak frequency	Where that spectrum is highest
–3 / –6 dB bandwidth	Width of the contiguous band around the peak that stays within 3 or 6 dB of it
10–70 Hz max / average / ripple	Highest, mean and (max – min) level in that band
Ghost notch	The deepest minimum near $c / (2 \cdot z)$, reported beside the value that depth predicts

F.5 What the numbers depend on

Three things set the absolute level, and all three must match the model you compare against:

ASSUMPTION	WHY IT MATTERS
Hydrophone distance	A notional signature is essentially the recorded pressure multiplied by the distance it travelled, so the level scales with it. The setting is the height above the cluster centre — one for the odd channel of each pair and one for the even, since the two hydrophones are vertically separated, and a cluster is solved with the height of the one it reads.
NFH scale	Bar per recorded unit. If the recording is not already in bar, every amplitude is out by that factor.
Filtering	Peak values depend on the filter applied. The app applies an anti-alias filter for the 2 ms output; a model report may also apply a recording-system response, which changes both the peak and the shape of the spectrum.

Primary-to-bubble and bubble period are deliberately not offered as model targets. A tuned, clustered array suppresses its bubble: there is no clear bubble peak to measure, and a model asked for one may print a figure on a summary page while its own signature plot reports that no bubble was found. Use them, as the app does, by comparing a line with itself.

What does not depend on any of this is shot-to-shot comparison: the same source, the same geometry and the same filter on every shot of the line. A source losing output shows up as a step in peak-to-peak whatever the absolute calibration, which is what the live page and the report's median and percentile columns are for.