

Far-Field Signature

Setting it up, running it, and reading what it says

SEG-D Monitor 2.26.3 · guide revision 2.26.3 · 2026-10-11

Contents

1. What this is	3
1.1 What it gives you	3
1.2 What it costs	3
2. The files you need	4
2.1 A near-field channel set	4
2.2 The Source Configuration	5
2.3 Gun positions (gun,x,y,z)	6
2.4 The array model report	6
3. Setting it up	7
3.1 Notional sources, or summation?	8
4. Choosing the channel set	9
4.1 Why the wider-band set is worse	9
4.2 What to choose	9
5. Which hydrophone, and how far above the guns	10
5.1 The setting and the per-cluster override	10
5.2 The two heights	10
5.3 Single hydrophones and dead ones	10
6. Targets: what to set, and what not to	11
6.1 Levels carry no verdict against a model	11
6.2 The ghost notch does carry one	11
6.3 Peak frequency and bandwidth, where they exist	11
6.4 Sources are not alike: targets per source	11
6.5 Ripple: the model has more of it than the sea	12
6.6 A sensible starting set of targets	13
6.7 Depth from the ghost notch	13
6.8 The depth baseline	14
6.9 Notional QC: PASS, WARNING, FAIL	14
7. Reading it while the line runs	16
8. Reading the End of Line pages	17
9. When something is wrong	21
9.1 Numbers that look wrong	21
10. A worked example	22
10.1 What to expect	22
10.2 When it does not agree	22
Appendix A. The gun position file	23
Appendix B. Checklist	24

1. What this is

The far-field signature is what the source sounds like far below the array: the one waveform the whole source radiates downwards, rather than what any individual gun does. It is the thing a processor deconvolves with, and the thing an array model such as Gundalf predicts.

SEG-D Monitor rebuilds it **from each shot's own near-field hydrophones**, so it can be followed along a line the way P/B and gun timing are. Nothing extra is recorded and nothing extra is read from disk unless you ask for it: by default the far field is computed from the traces the app has already read.

This guide is self-contained and can be handed to a crew on its own. It covers only the far field. Everything else — monitoring, the spec check, reports — is in the main *SEG-D Monitor User Manual*, and the mathematics is set out in that manual's Appendix F.

1.1 What it gives you

WHERE	WHAT
Graphs panel	A  Far field page: peak-to-peak per shot point, per source, live as the line runs.
End of Line report	Pages per line: every metric against its references, the levels along the line, the median signature with its spectrum, the depth check — and, if you switch it on, notional QC (§6.9).
Project database	One row per shot per source, including that shot's 2 ms signature (about 1 KB), so a report can be re-made later without re-reading the SEG-D files.

Everything is measured at the channel set's own sample rate. The signature is decimated to 2 ms only for storage and for the waveform page of the report. Measuring the decimated copy — which the app did before revision 2.10.0 — reports the anti-alias filter's peak rather than the source's, and reads peak-to-peak low by a margin that depends on the filter.

1.2 What it costs

Tens of milliseconds of processing per shot. No extra disk read at all when it uses the traces already read for P/B; when you point it at a different channel set ([§4](#)), it reads only the window it needs — a 512 ms window of a 10-second record is a twentieth of that record's bytes.

It runs in **normal mode only**. Light Mode never reads the traces it needs, and switches it off with everything else.

2. The files you need

Five things, of which two are optional. Gather them before you start: the feature cannot be configured halfway.

#	WHAT	NEEDED?	WHERE IT COMES FROM
1	SEG-D files with a near-field channel set	Required	The acquisition system. See §2.1 for how to confirm yours has one.
2	Source Configuration — which guns share which NFH channels	Required	Already in the app, from the Source Designer. See §2.2 .
3	Gun positions , a gun, x, y, z CSV	Required for the notional method	The array model report's gun table. Format in §2.3 , example in Appendix A .
4	Array model report (Gundalf or similar)	Optional	The source's design paperwork. Only needed if you want targets to compare against — §2.4 .

2.1 A near-field channel set

The source's hydrophones must be recorded in the SEG-D file as a channel set of their own. Open  Settings → **Channel Sets** and confirm the sets are described correctly: the app takes each set's **channel count, record length and sample interval from this table**, never from the file, so an error here is an error everywhere downstream.

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 1. Settings → Channel Sets. Every number the far field uses to find and read its traces comes from this table.

A recording system usually writes several sets: time breaks, the near-field hydrophones, and sometimes its own computed notional and far-field signatures. Which set holds what differs from system to system and from job to job — read it from the recording system's setup. Which set to compute from is §4, and it matters more than anything else in this guide.

Check the arithmetic once. Multiply out $\text{channels} \times (20 + 4 \times \text{record} \div \text{dt})$ for every set, add the file header, and compare with the actual file size. If it does not match to the byte, at least one set is described wrongly and the app will read from the wrong offset. The check proves the sample *counts*, not the sample interval: a wrong interval with the right count adds up just as well, and stretches every time in the file. Take the interval from the recording system's setup.

2.2 The Source Configuration

The far field needs to know which guns are recorded by which hydrophone. That is already in  Settings → **Source Configuration**: each gun carries its **NFH channels**.



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 2. Guns that share the same NFH channels form one cluster, and a cluster is one notional source in the computation.

Two things about this that catch people out:

- **A sub-array switched off here produces no far field for its source**, even if that source is firing. The app says so once per session in the log; see §9.
- **Guns sharing NFH channels are treated as one source**. That is correct — a clustered pair's bubbles merge, so the pair radiates as one — but it means the cluster list, not the gun list, is what the method works with.

2.3 Gun positions (gun, x, y, z)

The notional method needs to know where the guns are relative to each other, so that it can work out how much of each hydrophone's recording came from its own cluster and how much from the others. One line per gun, metres throughout:

COLUMN	MEANING
gun	The gun number as it appears in the Source Configuration.
x	Inline, increasing towards the stern; 0 at the front gun.
y	Crossline, negative to port, positive to starboard, measured from the array centre line.
z	Gun depth, positive down.

Separators may be commas, semicolons, tabs or spaces. A header row, blank lines and # comments are ignored. Only the guns you list are used, so a source numbered 25–48 lists those numbers. The complete example file is in [Appendix A](#) and ships with the app as docs/examples/gun_positions_example.csv.

The z in this file is the nominal modelled depth and is used only for the geometry between guns. Each shot's far field is rebuilt at the depths the guns actually measured on that shot, so the ghost notch follows the real depth rather than this number.

2.4 The array model report

Optional, and only for targets. From a Gundalf report's summary page you can read peak-to-peak, zero-to-peak, RMS, the spectral figures and the ghost notch. Before you type any of them in as a target, read [§6](#) — most of them should not be one.

The gun table of the same report is where the position file of [§2.3](#) comes from. Copy it; do not measure anything by hand.

3. Setting it up

Settings → Far Field is five pages:

PAGE	WHAT IS ON IT
Signal	Method, the channel set to read, polarity, the window, the NFH scale.
Channels	Which hydrophone of each cluster to use, and the per-cluster table.
Geometry	The two NFH heights, the depth baseline, speed of sound, sea-surface reflection, how the direct arrivals are summed, and the gun positions.
Targets	The values to check against, for all sources or for one, and which metrics the report carries at all.
QC	Notional QC — optional, off until ticked: the validation hydrophones and two report heat maps, each its own choice (§6.9).

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

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 3. The Channels page, on a four-source array whose clusters each carry a single hydrophone.

From nothing to a running far field

1. Open  Settings → **Far Field** and tick **Compute each shot's far-field signature**. The tick applies to all five pages below.
2. *Signal*. Leave **Signature polarity** on *SEG marine* — the Bubble Test's convention — unless you are laying the picture beside a model plot (§8). Choose the **Method**. *Notional sources* is the real one and needs gun positions; *Summation* needs no geometry but its output is not in bar·m — see §3.1.
3. Set the **Channel set**. This is the most consequential choice on the page: §4.
4. Set the **Window** — how much of the record to use, from its start. 512 ms is a sensible default and is the whole record on a typical NFH channel set.
5. Set the **NFH scale**: bar per recorded unit. 1.0 when the SEG-D already carries bar, which is the usual case. Everything the app reports in bar·m is wrong by this factor if it is wrong.
6. *Channels*. Set **Hydrophone to use**, and any cluster that differs, in the table below it — §5.
7. *Geometry*. Set the two **NFH height** boxes (§5.2), then **Import gun positions (CSV)** (§2.3) — the label confirms how many guns were read.
8. *Geometry*. Set **Speed of sound** and **Sea-surface reflection** to the values the recording system computes its far field with — they are in its setup. A coefficient quoted there as a positive number for an inverting surface (for example 0.7) is entered negative (−0.7). The defaults, 1500 m/s and −1, are nominal seawater and a perfect mirror: placeholders, not measurements. The coefficient takes any value from −1 to +1; see §6.2 for what its sign does to the ghost notch. Leave **Direct arrivals** on *Aligned* — how the recording system sums its clusters — unless you are comparing with an array model, which delays each direct arrival by its depth.
9. *Targets*. Fill in only the targets you actually want checked, and give a source its own where it differs from the rest — §6. Untick **Report** for any metric you do not want in the report at all.
10. QC. Optional: tick notional QC, the validation hydrophones and the heat maps you want (§6.9).
11. *Geometry*. Once a line has been shot on depth sensors you trust, take a **depth baseline** from it (§6.8).
12. Press OK and start a line. The first shot that produces a far field writes a line to the log.

3.1 Notional sources, or summation?

	NOTIONAL SOURCES	SUMMATION
What it does	Works out what each cluster alone radiated, removing the other clusters' contributions and the sea-surface reflection, then rebuilds the downward far field with the ghost at that shot's measured depths.	Adds the hydrophone traces together.
Needs	Gun positions.	Nothing beyond the traces.
Units	bar·m — comparable with a model.	Not bar·m. Comparable only with itself.
Use it when	Always, if you have positions.	You have no geometry and want a shot-to-shot consistency check only.

In Summation mode the app **offers no level target at all**, because there is no meaningful number to compare with. It is not being unhelpful: an earlier version did allow it, and marked every shot of a line as failing.

4. Choosing the channel set

This is the decision that changes the answer most, and the one most easily got wrong. The same shots, computed from two channel sets that differ only in their band, can differ in peak-to-peak by tens of per cent.

4.1 Why the wider-band set is worse

The method solves for one unknown per cluster, and the hydrophone array cannot resolve the sources above roughly

$$f_{\max} = c / (2 \times \text{cluster spacing})$$

— for example $1500 / (2 \times 3 \text{ m}) = 250 \text{ Hz}$ for clusters 3 m apart. Give the solve a recording that carries energy far above that and it has to explain detail it cannot represent: what comes back is not a better signature but ringing after the primary.

The app therefore solves only up to a smooth taper — full weight to 600 Hz, falling to nothing at 900 Hz, or over the top third of the set's own band when its Nyquist is lower — and not to a hard cut-off, which would ring in its own right. A set that has already been anti-alias filtered below the array's limit inverts cleanly and needs no rescue.

4.2 What to choose

Picking the set

1. List the sets that contain the **near-field hydrophones** (not a computed signature, not the seismic data).
2. Of those, discard any whose record is shorter than your window.
3. Of what remains, take the one with the **largest sample interval** — the narrowest band.

The app does this for you: once gun positions are loaded, the line under the settings grid names the recommended set, computed from your own cluster spacing rather than from any particular vessel's numbers. It is a recommendation, not a default — nothing is stored until you choose.

Leaving the setting on **Active set** keeps the behaviour of earlier versions: the far field uses whatever set the main window is monitoring, and costs no extra read.

The window is taken from the *start* of the chosen set's record. A 512 ms window of a 10-second set, for example, reads a twentieth of each trace, not the whole thing.

5. Which hydrophone, and how far above the guns

A cluster usually carries **two** near-field hydrophones, vertically separated — commonly an internal and an external unit — and the far field is computed from **one** of them, not from their average. That is how a GunLink system computes it, and it is what the app does.

5.1 The setting and the per-cluster override

Hydrophone to use sets the default for every cluster of every source: the *odd* channel of each pair, the *even* one, or both averaged. The table below it overrides that one cluster at a time, naming each cluster's real channel numbers so there is no counting involved.

A cluster is remembered by its **channels**, not by its position in the list. Switching a gun off rennumbers every cluster after it; your choices survive that.

5.2 The two heights

Because the two hydrophones hang at different heights, each has its own **NFH height above cluster** box, and a cluster is solved with the height of the hydrophone it actually reads. The computed level scales almost linearly with this distance, so it is the single most important number on the page after the NFH scale.

If every cluster carries one hydrophone, set both boxes to the same value. With one hydrophone per gun, consecutive channels are neighbouring guns rather than an internal/external pair, and 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 your array is in.

5.3 Single hydrophones and dead ones

- **One hydrophone on a cluster:** its row is locked and simply names that channel. There is nothing to choose, and it is used whatever the default says.
- **The selected hydrophone 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.

This is not an edge case: on many arrays only one hydrophone of each pair is live, and not always the same one of the pair.

6. Targets: what to set, and what not to

The app deliberately refuses to fail a shot on some metrics. This chapter is why, and it is the part worth reading before arguing with a number.

6.1 Levels carry no verdict against a model

Peak-to-peak, zero-to-peak, RMS and the band levels all scale with the distance from a gun to its hydrophone — a mounting dimension few vessels have measured. Moving the assumed hydrophone from half a metre to three metres above its cluster changes the level several-fold.

So the app **counts every target you set** — "59 of 60 shots within", and the count outside when there is one — and prints the distance from the target beside the median, but it does not present a level that disagrees with a model as a proven source fault. Judge it as a calibration question first.

Look at the line's own shot-to-shot spread before choosing a tolerance: it should sit well inside it. The shot counts are what show whether the line held it, since a line can sit on target with a third of its shots outside.

6.2 The ghost notch does carry one

The sign of the sea-surface reflection decides where the notch is. A pressure-release surface inverts the reflection ($R = -1$, the usual case) and the first notch falls at $c/(2z)$ — for example about 94 Hz for guns at 8 m in water at 1500 m/s. A non-inverting coefficient ($R > 0$, which is the right entry under the opposite polarity convention) reinforces at that frequency and notches at $c/(4z)$ instead — about 47 Hz at the same depth. At $R = 0$ there is no ghost and the app reports no notch at all. Set the coefficient to match your data's polarity and the notch follows it; the app does not assume.

Its magnitude moves peak-to-peak and nothing else. The primary arrives before the ghost, so zero-to-peak barely moves with it, while peak-to-peak — whose trough *is* the ghost — falls as the magnitude falls. If your peak-to-peak disagrees with the recording system's and your zero-to-peak agrees, the reflection is the first thing to compare, not the source.

Set a target on **ghost notch** with confidence. It is geometry, not calibration: it depends on the depth of the guns and the speed of sound, both of which are known. A notch that moves is a source that has changed depth.

6.3 Peak frequency and bandwidth, where they exist

A tuned array's far-field spectrum is flat: it can sit within 1 dB of its maximum over several tens of hertz. There is no peak in it to find, and a single-bin maximum picked from such a spectrum jumps from shot to shot.

Where the app detects this, it reports **no peak frequency and no bandwidth** rather than a number that means nothing, and gives the *width of the flat top* instead. A model of the same source may well quote a peak frequency, because a modelled signature has more structure than the real one; that is a difference between the model and the sea, not an error in either.

6.4 Sources are not alike: targets per source

A vessel's sources are often different sizes: an array can carry one source several times the volume of the others, with a far field correspondingly larger. A single peak-to-peak target cannot judge both: whichever number you set, one group is wrong.

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

Targets for: All sources ☐ this source has its own

Editing the shared list: 6 target(s). Every source follows it unless given its own.
Theoretical values to check each shot against (leave Target empty to skip a metric):

Metric	Target	Tolerance	in %	Report
Peak-to-peak (bar-m)	101.6	5	☒	☒
Zero-to-peak (bar-m)	61.5	5	☒	☒
RMS, 2 s window (bar-m)	3.3	10	☒	☒
Peak frequency (Hz)	55.18	3	☐	☒
-3 dB bandwidth (Hz)	31.25	5	☐	☒
-6 dB bandwidth (Hz)		0	☒	☒
10-70 Hz maximum (dB)		0	☐	☒
10-70 Hz average (dB)		0	☐	☒
10-70 Hz ripple (dB)		0	☐	☒
Ghost notch (Hz)	95.1	5	☐	☒

Target peak-to-peak, zero-to-peak and the ghost notch. The three 10-70 dB rows are measured but are NOT model-comparable: an array model carries in-band bubble energy a tuned array does not produce, so a healthy source sits several dB under the model and a target taken from it fails every shot. P/B and bubble period are left out entirely for the same reason: such an array has no clear bubble and the model does not define them.

OK Cancel

Figure 4. The Targets page. **Targets for** selects whose targets the table is editing; a source with none of its own follows "All sources".

Giving one source its own targets

1. Fill in the table with **All sources** selected. Every source follows these unless told otherwise.
2. Select the source that differs in **Targets for**. The table greys out: it is still following the shared list.
3. Tick **this source has its own**. The table comes alive, pre-filled with a copy of the shared values.
4. Change what differs. Untick to send that source back to the shared list.

In the report, a metric whose sources agree prints one figure in the **Target** column; where they differ the column says *per source* and each source's cell carries its own number. Each source is judged only by the targets that apply to it.

6.5 The spectral dB figures are not model-comparable

Do not set a target on **10-70 Hz maximum, average or ripple** from an array model. The model's bubble puts energy — and structure — into that band that a tuned, clustered array does not produce, so a healthy source measures several decibels under the model's level and far less ripple, and the app and the recording system agree with each other while both disagree with the model. A target taken from the model fails every shot of every line and tells you nothing.

A percentage tolerance on a decibel is meaningless. A dB is already logarithmic, so "210 dB ± 10 %", for example, is ± 21 dB — a factor of eleven in amplitude, which nothing can fail. The dB rows therefore start in their own units; leave them that way and give a tolerance in dB.

6.6 A sensible starting set of targets

METRIC	SET A TARGET?	WHY
Ghost notch	Yes	Geometry; provable; a real depth check.
Peak-to-peak	Enter the model value	Printed for comparison; judged against the line's own median, never failed against the model.
Zero-to-peak, RMS	Optional	Same treatment as peak-to-peak.
Peak frequency, bandwidths	No	Undefined on a flat spectrum; the app will report nothing to check.
10–70 Hz maximum, average, ripple	No	None of the three is model-comparable (§6.5).
P/B, bubble period	Not offered	A tuned, clustered array has no clear bubble, and models do not define them for one.

6.7 Depth from the ghost notch

Every other depth figure in this app comes from the gun's own depth transducer, so nothing can contradict one. This can. The direct arrival at a hydrophone and its surface reflection are $2 \cdot z/c$ apart, so the first spectral notch of the hydrophone's own trace gives the depth it hangs at — and adding its height above the guns gives the gun depth, measured acoustically.

It is measured on the **raw hydrophone trace**, not on a notional. A notional is ghost-free by construction, and reading the notch off the synthesised far field is worse than useless: that ghost is placed using the depth sensor's own number.

But the notch is not an absolute depth. Each hydrophone hears every cluster of the array and every cluster's ghost, not just its own, and no single trace can separate them: the notch it shows is pulled by its neighbours — by up to a metre on the big guns of an array and hardly at all at the string ends. The pull is steady for as long as the array fires the same guns, so the app does not show the notch as a depth until a **depth baseline** has been taken (§6.8). Until then the report page and the live window show the depth sensors and say *Not calibrated*.

It is information, not a judgement. The page raises no problem, reaches no verdict on the cover and never counts a shot point bad. With a baseline, a cluster whose depth has moved from its sensor by more than the depth tolerance since the baseline is highlighted in **light yellow**, and that is the whole of it. The table lists **every cluster with its gun numbers**, so every gun is accounted for.

The Spec Check ignores all of this. A far-field result — the signature, the comparison or the depth — never counts a shot point bad. The Spec Check judges timing, depth, pressure, nav mask, missing shots and a dead source, and nothing else.

Per cluster, not per gun. The notch gives the hydrophone's depth, so a cluster carrying two guns yields one depth, compared against the mean of its guns' sensors.

Every shot is read. On each shot, every hydrophone of the cluster that carries a signal is read, and the shot's depth is their mean — so it does not matter which hydrophone Hydrophone to use selects; that setting is the far field's alone. The figure is a **line median**, never a per-shot alarm — which suits the failure it hunts, since a transducer drifts slowly.

6.8 The depth baseline

A baseline is taken from a line whose depth sensors you trust: Settings → Far Field → Geometry → Take depth baseline.... For each hydrophone it measures the offset between what its notch says and what the sensors say on that line — the mounting height and the pull of the neighbouring clusters together — and it records which guns each source fired. From then on every reading, on every line (the baseline line included, and lines shot before it was taken), is converted through that offset, so the baseline line reads its own sensors and a later line shows **drift since the baseline**.

It takes that line's depth sensors as truth, so choosing the line is the assertion. The measured offsets are shown first, hydrophone by hydrophone, with Apply and Cancel — and Cancel is the default. Clear beside the button discards the baseline, and the depth check goes back to showing the sensors only. A hydrophone read on fewer than 20 shots of the line gets no offset. A line processed before version 2.11.4 does not record which hydrophone each depth came from, so re-run it before taking a baseline from it; one processed before 2.14.0 does not record which guns fired, so its baseline cannot tell firing configurations apart until the line is re-run. A calibration taken with an earlier version is kept as a baseline, and the Geometry page says its line was not recorded.

A shot fired with a different set of guns is not judged. The pull of the neighbours depends on which of them fire, so a shot whose source fired a different set from the baseline's is shown — marked *firing differs from baseline* — and counted, but kept out of the medians and the highlight. Take a new baseline when the array's normal firing configuration changes.

What it can and cannot catch. A baseline finds a cluster that drifts away from its sensor after the baseline was taken. It cannot find an offset that was already there on the baseline line — nothing acoustic can. That needs a tape measure. The baseline is used by the depth check only; the far-field solve uses the two NFH heights.

6.9 Notional QC: PASS, WARNING, FAIL

Optional. Tick it on Settings → Far Field → QC and every notional shot is judged **PASS**, **WARNING** or **FAIL**; a WARNING or a FAIL always names its reasons. A shot is judged on what it can be judged on alone:

CHECK	WHAT IT ASKS	GIVES
Validation hydrophones	The solve reads one hydrophone per cluster, so reproducing that one proves nothing: there are as many unknowns as recordings. The other hydrophone of each pair is an independent test. It is predicted from the notionals at its own height and compared with what it recorded, over 10 ms before to 250 ms after source time. Judged on its shape — the correlation. The normalised RMS difference is shown but judges nothing: it also carries the level, and so the hydrophone height.	WARNING, or FAIL below the lower limit. Its own tick box.
The solve	How much the solve could magnify an error in the recordings — the condition number of the propagation matrix over the band the solve uses at full weight — and whether it lost a cluster. It depends only on the geometry: positions, heights, depths, speed of sound and reflection. So it catches a wrong setup or a bad depth reading, never a gun.	FAIL above 100, or when a cluster is lost. Always checked.
Inputs	A cluster with no live hydrophone at all.	FAIL. Always checked.

The heat maps. The End of Line report adds two heat maps per source, each with its own tick box, one row per cluster and one column per shot point:

- **Energy against the cluster's own line median** — its own, not a sub-array's: a sub-array mixes gun volumes, so against a shared median a healthy small cluster reads weak and a healthy large one strong.
- **Bubble period against the clusters of the same volume** — every cluster on the line, of every source, whose guns that fired add up to the same volume. A cluster steadily off its peers shows here where its own median would hide it.

White is no difference; blue is below, amber above; full colour at the scale you set, deeper past it. Grey means there is nothing to show: the cluster did not fire, no bubble was found, or the shot failed and is left out. A healthy line is pale speckle; a cluster going wrong is a horizontal stripe, a bad stretch of line a vertical one. The heat maps are pictures, never a verdict: they change no shot's PASS, WARNING or FAIL.

Where there is no second hydrophone — one hydrophone per cluster, or Hydrophone to use set to *Both, averaged* — there is nothing independent to validate against, and the report says *no validation hydrophones*.

Every limit and scale is in Settings → Far Field → QC. The defaults are a starting point: set them from your own lines. By default a FAILED shot is left out of the line's far-field statistics, median signature and heat maps (Leave FAILED shots out of the line's far-field statistics).

Information only. Like everything in the far field, notional QC reaches no problem row, no cover verdict and never the Shot Point Spec Check.

7. Reading it while the line runs




Figure 5. The Far field page of the Graphs panel: peak-to-peak per shot point, one panel per source, target band in green.

Select **Far field** in the Graphs panel's metric list. Each source gets its own panel; each point is one shot's peak-to-peak. Where a peak-to-peak target is set, its band is drawn in green and shots failing a target are drawn in red.

What to look for, in order of how often it matters:

- **A step.** A sudden change in level on one source and not the others is a source event — a gun out, a change in depth, a volume change. Cross-check with the Depth, Pressure and Timing graphs.
- **A drift.** A slow trend over a line usually follows depth; check the Depth graph over the same shot points.
- **Scatter.** Healthy scatter is small and steady — learn your own source's normal. Sudden scatter with no change in median often means a hydrophone is dropping in and out — the log will say so.
- **A flat line at a level you did not expect.** Not a source problem. Check the NFH scale and the hydrophone height before anything else.

8. Reading the End of Line pages




Figure 6. Page one: the metrics, per source, with their references and peak-to-peak along the line.

The table gives each metric's **line median** with its 5th–95th percentile in brackets — so the spread is visible, not just the middle. Beside it:

- the **Target** column, the model figure, where one is set;
- **how many shots met the target**, for every target you set ([§6.1](#)).


A row is shaded when a median falls outside a target it is actually judged on. The note under the table repeats which is which, so the page can be read without this guide beside it.

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



Far-field zero-to-peak along the line



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 7. Page two: the line's median signature and its amplitude spectrum.

Polarity. The signature is published in the **Bubble Test's convention**, and the two windows agree by construction — one crew reads both reports for one source. The SEG standard for marine seismic has a compression recording *negative* on a hydrophone, so the primary is the strongest negative excursion, the surface ghost — inverted by the free surface — is positive, and a bubble is negative again. What As *recorded* gives depends on the hydrophones' and the recording system's conventions: look at it once on your own data.

The one reason to change it: a model plot is usually drawn the other way up. An array model plotted with its own *SEG polarity* option **disabled** shows the primary positive. Set *Signature polarity* to **Positive** when the two pictures are to be compared side by side. No measured value changes under any setting. That is not a convenience, it is what the metrics ARE: zero-to-peak is the largest *absolute* excursion, peak-to-peak a span between two extremes, the spectrum a magnitude and the notch a frequency. None of them carries a sign, which is why a model report quotes zero-to-peak unsigned and why your targets never need one. Only the stored waveform is re-signed — and an inverted signature used as a deconvolution operator would flip the polarity of a section. A line processed before the setting was changed keeps its old picture until it is re-processed; its figures are comparable as they stand.

The median signature is taken sample by sample across the line, so a single bad shot does not move it. Read the spectrum for the **ghost notch** — a sharp minimum whose frequency gives the depth as $c / (2 \times \text{notch})$ — and for the general shape. A missing notch means the depths used were wrong, or the shots were not aligned; a notch at the wrong frequency means the guns were not at the depth you think.

8.1 Page three: the shots that missed

Page one's *shots failing a target* figure is a count. When it is not zero, a third page follows the graphs listing the shots behind it: one row per target missed, giving 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 for each miss.

It exists because a count on its own cannot be acted on — it names neither the shot points nor what they missed, so there is nothing to go and look at. The page is absent when no target is set or nothing fell outside one, like every other block in this report that would only say "nothing is wrong".

Read a level miss differently from a shape miss. A level — peak-to-peak, zero-to-peak, RMS — scales with the hydrophone mounting distance ([§6.1](#)), so a whole line sitting below its target is a calibration question before it is a source fault. Four shots out of four hundred sitting below it is not: that is a real excursion, and worth the SP numbers this page gives you. The ghost notch has no such excuse at all — it is geometry ([§6.2](#)).

The far-field pages appear only for a line that has far-field rows, so a project that never switched the feature on gets the report it always got.

8.2 Notional QC

Where notional QC was on while the line was shot, the *Notional QC* page gives each source's PASS, WARNING and FAIL counts, a tally of the reasons and the validation hydrophones' correlation (median and worst) with their normalised RMS difference shown beside it. The checks in force, the heat-map scales and how source time was taken are printed under the table; validation reads *not checked* when it was not ticked.

The heat-map pages follow, one per source ([§6.9](#)), then the clusters one row each: its guns and their volume, the bubble ratio, validation, and how many shots raised each flag. A row with a flag is light yellow; a cluster that lost every hydrophone is red. When FAILED shots are left out of the statistics, the first far-field page says how many.

9. When something is wrong

Everything the far field complains about is said **once per source per session** in the app log, so a bad shot cannot flood it. Open the log with  Logs.

WHAT THE LOG SAYS	WHAT IT MEANS	WHAT TO DO
no gun positions configured	The notional method has no geometry, so nothing is computed.	Import the CSV (§2.3), or switch to Summation if you have no positions.
source N fired but has no clusters	That source's sub-arrays are switched off in the Source Configuration, or its guns carry no NFH channels — so it produces no far field while firing perfectly well.	Tick the sub-arrays back on, or give the guns their NFH channels.
channel set N could not be read	The chosen set's channel count, record length or sample interval does not match the files.	Check Settings → Channel Sets against the file size (§2.1).
the selected hydrophone is dead on cluster(s) ...	Normal on many arrays: only one hydrophone of the pair is live. The other is used instead, with its own height.	Nothing, unless you expected both live. Consider setting that cluster explicitly (§5.1).
the depth reading of gun(s) ... is at or above its hydrophone	A gun cannot be shallower than the hydrophone hanging above it, so the reading is left out and the cluster's other reading, or its nominal depth, is used for that shot.	Check that gun's depth sensor.
notional QC FAIL – ...	A shot failed a notional QC check; the reasons are named. Said once per source; the report counts them all.	Read the Notional QC pages (§8.2).
Far field not computed for source N: ...	The geometry cannot be solved — usually two clusters at the same position, or a hydrophone nearer another cluster than its own.	Check the position file: duplicated coordinates, a wrong sign on y, or a height larger than the gun spacing.

9.1 Numbers that look wrong

SYMPTOM	MOST LIKELY CAUSE
Level far from the model	The NFH scale or the hydrophone height. Both scale the answer directly. This is a calibration question before it is a source problem — see §6.1 .
Level low and the waveform rings	Too wide a channel set (§4).
No ghost notch, or one at the wrong frequency	The gun depths in the headers, or the speed of sound. The notch should be $c / (2 \times \text{depth})$.
Peak frequency blank	Not a fault: the spectrum is flat and the app declines to name a peak (§6.3).
One source reads differently from the others	Compare their hydrophone selections and heights first; a genuine difference between sources shows in the depth and pressure graphs too.
Nothing at all on one source	The no clusters case above.

10. A worked example

An illustration of the order of work, not a record of any survey: a three-source array recorded by GunLink. Every value comes from somewhere you can look up — the right-hand column says where.

SETTING	EXAMPLE	WHERE IT CAME FROM
Method	Notional sources	Gun positions were available.
Channel set	the hydrophone set	The narrowest-band set carrying the hydrophones (§4).
NFH scale	1.0	The data is already in bar.
Speed of sound / reflection	1500 m/s, −0.7	GunLink's setup, which quotes the coefficient as 0.7.
Direct arrivals	Aligned	GunLink's convention.
Hydrophone heights	as mounted	The rigging drawings, or measured on deck (§5.2).
Gun positions	every gun	The array model's gun table — Appendix A .
Depth baseline	the first line	Shot on sensors the crew trusts (§6.8).

10.1 What to expect

- **The validation hydrophones agree in shape.** With notional QC on, their correlation close to one on every cluster, shot after shot.
- **The heat maps are pale.** Scattered cells at most — no row or column standing out.
- **The notch sits where the depths put it:** near $c / (2 \times \text{depth})$ for an inverting surface.
- **The spread is small.** Shot-to-shot scatter of a few per cent at most is what makes the far field useful: it shows a real change long before anyone could see it in the absolute number.
- **The depth page reads the sensors on the baseline line** and, on later lines, shows only drift.

10.2 When it does not agree

Work down the settings the far field depends on, in this order: the speed of sound and reflection against the recording system's setup; the direct-arrival convention; the NFH scale; the heights, against the rigging as mounted; and last the gun positions. The measured values are the evidence — change a setting only to match the rig and the recording system's setup, never to make the numbers agree.

Appendix A. The gun position file

This is the complete example that ships with the app as `docs/examples/gun_positions_example.csv`. Its numbers are a real 3120 cu.in two-string source at 8 m: twelve guns per string, 3 m inline spacing, 1 m between the two guns of a cluster, strings 10 m apart.

```
# SEG-D Monitor - example gun position file
# Settings -> Far Field -> "Import gun positions (CSV)..."
#
# Format:  gun,x,y,z   - one line per gun, all values in METRES.
#   gun   the gun number as it appears in the Source Configuration
#   x     inline, increasing towards the stern (0 at the front gun)
#   y     crossline, negative to port and positive to starboard,
#         measured from the array centre line
#   z     gun DEPTH, positive down
```

```
gun,x,y,z
```

```
# ---- String 1 (port) ----
```

```
1,0.000,-5.500,8.000      # 40 cu.in
2,0.000,-4.500,8.000      # 40 cu.in
3,3.000,-5.500,8.000      # 100 cu.in
4,3.000,-4.500,8.000      # 100 cu.in
5,6.000,-5.500,8.000      # 250 cu.in
6,6.000,-4.500,8.000      # 250 cu.in
7,9.000,-5.500,8.000      # 250 cu.in
8,9.000,-4.500,8.000      # 250 cu.in
9,12.000,-5.500,8.000     # 100 cu.in
10,12.000,-4.500,8.000    # 100 cu.in
11,15.000,-5.500,8.000    # 40 cu.in
12,15.000,-4.500,8.000    # 40 cu.in
```

```
# ---- String 2 (starboard) ----
```

```
13,0.000,4.500,8.000      # 40 cu.in
14,0.000,5.500,8.000      # 40 cu.in
15,3.000,4.500,8.000      # 100 cu.in
16,3.000,5.500,8.000      # 100 cu.in
17,6.000,4.500,8.000      # 250 cu.in
18,6.000,5.500,8.000      # 250 cu.in
19,9.000,4.500,8.000      # 250 cu.in
20,9.000,5.500,8.000      # 250 cu.in
21,12.000,4.500,8.000     # 100 cu.in
22,12.000,5.500,8.000     # 100 cu.in
23,15.000,4.500,8.000     # 40 cu.in
24,15.000,5.500,8.000     # 40 cu.in
```

For a source whose guns are numbered 25–48 or 49–72, list those numbers. Guns you do not list are simply not used.

Appendix B. Checklist

One page to work through before a first line.

Before you start

1. Channel Sets describe the file correctly — the arithmetic adds up to the file size.
2. Every gun in the Source Configuration carries its NFH channels.
3. Every sub-array that will fire is ticked active.
4. Gun positions imported; the label confirms the expected count.
5. Channel set chosen — the narrowest-band hydrophone set that covers the window.
6. NFH scale correct for the units in the file.
7. Hydrophone heights entered; both the same if there is one hydrophone per cluster.
8. Speed of sound and reflection from the recording system's setup; direct arrivals Aligned.
9. Notional QC switched on, with the validation and the heat maps you want, if you want it.
10. A ghost-notch target set; level targets entered only as references.

After the first ten shots

1. The  Far field graph shows points for every source that is firing.
2. The log carries no far-field warning you have not read.
3. The ghost notch is near $c / (2 \times \text{depth})$ for your depth.
4. The scatter is a few per cent, not tens.
5. A depth baseline taken, once a line has been shot on sensors you trust.

At the end of the line

1. The far-field pages are in the report.
2. The median signature has a visible primary and a visible notch.
3. Level differences from the model are understood as calibration, not logged as source faults.