



SONARPOINT · RELIABILITY

Built to come back: redundancy in SonarPoint arrays

How six independent safeguards protect the dataset and the gear. Field results from U.S. Navy LMR Deployments 1.3 and 2.

6 / 6

LMR2 moorings commanded to release came up and were recovered, in under four hours

0

float failures with pressure-tested syntactic foam at 750–1450 m in LMR2

1 of 4

floats left after an implosion at 1000 m, yet the mooring and its data were recovered (LMR1.3)

1 Why redundancy matters

A SonarPoint array spends weeks to months on the seafloor, out of sight and out of contact. Its recordings stay inside each recorder until that recorder is back on deck. A problem that goes unnoticed, such as a release that will not fire, a float that collapses or a clock that cannot be synchronized, may only be discovered at the end of the project, when it is too late to fix.

And a localizing array is more than the sum of its moorings: every position depends on precise timing from several recorders at once. A lost mooring or an unsynchronized recorder costs data across the whole array.

SonarPoint is therefore built in layers. No single part has to work perfectly; each critical function is covered by at least two independent mechanisms, so that one failure does not become a lost dataset.

2 Six layers at a glance

#	Safeguard	Protects against
1	Dual acoustic releases	A release that fails to fire
2	Secure-release protocol	Replies from a deep release not heard at the surface
3	Pressure-tested syntactic foam floats	Float implosion at depth
4	Reserve buoyancy	Loss of one or more floats
5	SeaTag satellite transmitter + radio direction finding	Early or unplanned surfacing; finding a drifting mooring
6	Two independent time-sync methods	Pinger failure, terrain shadowing, long baselines, early battery depletion




Figure 1. SonarPoint mooring string, about 9 m (30 ft) tall. Gold numbers mark where each safeguard sits.

3 Layer 1: Dual acoustic releases

Each mooring is anchored through **two ARC1800-1XD acoustic releases** that share one release cord. Either release alone frees the cord and lets the mooring rise. If one release fails to receive the command, has a mechanical problem or runs low on battery, the other still brings the mooring home.

Dual releases were also a condition of the Monterey Bay National Marine Sanctuary permit for these deployments: the sanctuary required assurance that gear would not be left behind.

In the field. During the LMR2 recovery on October 2, 2026, one release of mooring M7 did not trigger when commanded. Its partner did, and M7 surfaced and was recovered normally. On three other moorings a single command to one release was enough.




Figure 2. (a) Recovered LMR2 mooring M3: the two acoustic releases with their shared release cord, and floats. (b) Recovered M2: SonarPoint recorder and hydrophone.

4 Layer 2: The secure-release protocol

A conventional acoustic release recovery is a conversation. The boat sends a link request and the release must reply before the boat can range on it, send an arm command, receive the acknowledgment and finally send the release command. Every step depends on the boat hearing a reply from a small transducer on the seafloor, often 1000 m or more below.

Analysis of the LMR1.3 recordings showed where this breaks down. Commands from the boat reached releases at about 1000 m reliably, but the releases' replies were often not detected at the surface. Each missed reply stalled the sequence and made recovery difficult.

Conventional sequence

↓ Link request → ↑ Reply → ⇅ Ranging → ↓ Arm → ↑ Acknowledge → ↓ Release

Three steps (orange) need a reply heard at the surface before the mooring can leave the seafloor.

Secure release

↓ Link request with release ID → Release fires; mooring rises → ⇅ Reply and ranging while rising

Only the downlink, which is the reliable direction, is needed to start the recovery.

With the secure-release firmware, introduced for LMR2, a release that has been in the water for more than 24 hours triggers as soon as it hears the first link request carrying its own ID. The 24-hour lockout prevents accidental release during deployment operations, and the ID requirement means only the intended mooring responds. Because the release starts rising immediately, the acoustic path shortens with every minute, and the replies that were hard to hear from the seafloor come through clearly.

In the field. In LMR2 the first link attempt often got no reply, but contact followed within a few minutes, typically with the release already a few hundred meters off the seafloor and climbing. All six commanded moorings, at 750–1450 m, released at once.

5 Lesson from LMR1.3: a 1000 m implosion

In LMR1.3, five moorings were deployed at 755–1102 m depth in October 2025. Each was lifted by four 9-inch plastic hard floats: the top float rated to 1800 m, the three below it to 1500 m. On recovery day two moorings came up normally. Two acknowledged the release command but stayed on the seafloor, and the deepest did not respond.

The array diagnosed its own failure. The two recovered recorders contained six implosion events. Because the recorders are time-synchronized, each event could be traced to one of the three stuck moorings, two per mooring and in order of depth, deepest first. The 1500 m rated floats had failed by creep at less than 1000 m.




Figure 3. Mooring M4 found by MBARI's ROV *Doc Ricketts* at 1004 m, February 2026. Both releases had fired; the line still rises toward the one surviving float.




Figure 4. The recovered top section: three 1500 m rated floats imploded; the 1800 m rated top float with its SeaTag survived.

Mooring M4 shows how the remaining layers carried the recovery. Both releases had fired as commanded, but too little buoyancy was left to lift the anchor-free string. After the ROV cut the line above the releases, the one surviving float brought the recorder up, barely awash. The SeaTag on that float reported within minutes of surfacing, its positions led the search over two days of drift, and a boat-mounted radio direction finder closed the last 1.5 nautical miles. **The recorder had kept recording through the implosions, and its data were recovered intact.**

6 Layer 3: Pressure-tested syntactic foam floats

For LMR2 the plastic hard floats were replaced by custom **syntactic foam floats** (FLOAT-1800: 10 in diameter, about 9 lbf buoyancy each, 1800 m service depth). Syntactic foam is a solid material made of microscopic glass spheres in resin, so there is no large hollow volume to implode. Material samples from a production float were pressure tested to confirm the 1800 m rating before deployment.

In the field. LMR2 placed seven moorings at 500–1450 m for a month. Every float on the six recovered moorings came back intact, and no implosions were recorded.

7 Layer 4: Reserve buoyancy

A SonarPoint mooring carries more flotation than it needs. With four syntactic foam floats, the bottom float alone cannot lift the string, the second float holds it up at the surface, and the top two are reserve. In the 1800 m configuration, two or three floats are enough to bring the string up; the rest are margin. Spreading the buoyancy over several separate floats means no single float is critical. The margin also keeps the top float and its antenna riding clear of the water after release, which matters for the next layer.

8 Layer 5: SeaTag satellite transmitter and radio direction finding

Every mooring carries a solar-powered Desert Star **SeaTag** on its top float. At depth it is silent and dark. As soon as the float reaches the surface, it transmits to the Argos and Kineis satellites. This serves two purposes:

- **Early-release warning.** A tag on a seafloor mooring should never transmit. Any message means the mooring is at the surface, whether by command or not.
- **Tracking.** Satellite Doppler positions, typically within 250 m to 1 km for the best fixes, bring a boat to the right area. The tag also transmits about once a minute, so a boat-mounted radio direction finder (RDF) can home in over the last few miles.

The tag's sensors add context: the SeaTag-MOD on LMR1.3 M4 logged a 2 m drop in depth at the moment a float imploded, confirming the acoustic diagnosis.




Figure 5. (a) An LMR2 mooring at the surface after release, SeaTag antenna up. (b) The release deck unit during recovery, showing range to the connected release and distances to the others.



Figure 6. Doppler radio direction finder connected to a VHF/UHF scanner tuned to the tag frequency; the LED ring shows the bearing.




Figure 7. Boat track on the LMR1.3 M4 search, February 14, 2026: course changes follow the RDF bearings to the recovery location.

In the field: an honest account. In LMR2, mooring M6, the shallowest at about 500 m, left the seafloor on its own twelve days into the deployment, for a reason not yet determined. Its SeaTag reported the surfacing within minutes. The hardware did its job; our response did not. The messages went unnoticed for two days because the satellite portal was checked by hand, and we then misread the tag's temperature data as a tag floating free while the mooring stayed down. By the time the scheduled recovery showed otherwise, M6 had drifted beyond the range of our boat. It has not been recovered. The full account is in our application note *Satellite Tracking of SonarPoint Mooring Strings*.

What we changed: a locking pin now fixes the tag to its float so a transmitting tag always marks its mooring; tags are deployed in on-fish mode, which adds sensor summaries that help diagnose a mooring's state; automatic alerts on the first message are being set up so a response can start the same day; and ocean-model current forecasts are used to judge where a drifting mooring is headed.

9 Layer 6: Two independent time-synchronization methods

Localization depends on knowing when each recorder heard a sound, to a millisecond or better. Each recorder keeps time with its own crystal oscillator, which drifts slowly. SonarPoint measures that drift in two independent ways:

A. Continuous acoustic sync. A sync pinger on one mooring emits a ping every recording segment, typically every 10 minutes. Every recorder that hears it gets a fresh timing reference throughout the deployment.

B. GPS-timed sync before and after. Just before deployment and right after recovery, a deck unit beside each recorder emits GPS-timed pings. These two anchor points define a straight-line drift correction for the whole deployment.

The two methods fail in different ways, so each covers the other's weak spot:

If this happens...	...this method still works
Recorder battery runs out early: no post-recovery GPS anchor	A: continuous pings up to the last recorded file
Sync pinger fails	B: GPS anchors before and after
Terrain shadows the pinger from some moorings	B
Long baselines: pings fade beyond roughly 100× the pinger's height above the seafloor	B
Moorings sway in tidal currents, adding false drift to A	B; A becomes a cross-check




Figure 8. LMR 1.2: drift measured by the continuous pinger at a remote mooring (blue) against the GPS-anchored straight line (orange). The wiggles are the mooring swaying with the tide, not clock error.

Measurements with a pinger co-located with its recorder show that SonarPoint oscillators drift on a straight line in the stable deep-water environment, with residuals under 0.5 ms over two weeks. GPS interpolation therefore delivers about 1 ms accuracy on its own, while the continuous pinger confirms that nothing unusual happened in between. On top of both, Pipeline Manager's TDOA closure test checks timing consistency across the array on every localized call.

10 How the layers work together

Failure	First line	Backup
One release does not fire	Second release (1)	SeaTag + RDF if the string surfaces later (5)
Release replies not heard at the surface	Secure release fires on the downlink alone (2)	Ranging once the release is rising
Float implodes at depth	Syntactic foam, pressure tested (3)	Reserve floats still lift the string (4)
Mooring surfaces early or unplanned	SeaTag alert and positions (5)	RDF for the final approach; reward markings on floats
Mooring held on the seafloor	Acoustic ranging locates it; ROV can free it	SeaTag tracks it once cut free (5)
Timing reference lost on one recorder	Other sync method (6)	TDOA closure test flags inconsistent timing

11 Track record

Deployment	Configuration	Outcome
LMR1.3 Oct–Nov 2025 5 moorings, 755–1102 m	Plastic hard floats (1500 m rated, top float 1800 m); dual releases; SeaTag on top float; conventional release protocol	2 recovered on command. 3 held down by imploded floats. M4 recovered with its data in February 2026 via ROV, SeaTag and RDF. M2 and M5 remain on the seafloor. Recordings from 3 of 5 moorings recovered.
LMR2 Sep–Oct 2026 7 moorings, 500–1450 m	Syntactic foam floats (1800 m); dual releases; secure-release firmware; SeaTag on top float	All 6 commanded moorings released and were recovered in under four hours, with all floats intact. M6 surfaced uncommanded on day 12; its SeaTag reported it and still tracks it, but it drifted beyond boat range before recovery. Recordings from 6 of 7 moorings recovered.

These were engineering deployments, run under the Navy LMR program specifically to find weaknesses before routine use. LMR1.3 exposed the float and release-communication weaknesses; both were fixed and proven in LMR2. LMR2 exposed weaknesses in how a tag was attached and how its alerts were monitored; those are fixed in our hardware and procedures for the next deployment.

12 Recommendations for SonarPoint users

- **Keep both releases** on every mooring and test each on deck before deployment.
- **Use the secure-release firmware** for deep deployments, and plan recovery so the boat can range on the release as it rises.
- **Stay well within the flotation rating**, and do not remove reserve floats to save weight.
- **Lock a SeaTag to every top float**, set up automatic alerts before deployment, and agree in advance who responds to an alert and how. Carry an RDF receiver on every recovery trip.
- **Have an early-surfacing plan**: a vessel available on short notice and its working range from shore. When in doubt, go: an unnecessary trip costs a day, a lost mooring costs its dataset.
- **Use both sync methods**: a sync pinger in the array, plus GPS-timed pings at every recorder before deployment and immediately after recovery.

Bottom line

- Every critical function of a SonarPoint mooring, from release and flotation to location and timing, is covered by at least two independent mechanisms.
- The layers have been tested by real failures at 1000 m: a mooring held down by imploded floats still returned its data, and with the LMR2 upgrades every commanded mooring came back.
- Redundancy in hardware works best with readiness in procedures: alerts watched, a boat on call and a plan agreed before deployment.

Related application notes: *How accurate is a sparse SonarPoint array?* · *Satellite Tracking of SonarPoint Mooring Strings* · *Clock Drift Compensation in SonarPoint Acoustic Arrays*. Contact: Marco Flagg, m.flagg@desertstar.com.