

Which source do you believe?

When your instruments disagree and one of them is confidently wrong

AT A GLANCE

The problem. A failed sensor usually looks failed. An interfered-with sensor can look healthy. Most alarm design addresses the first case.

What this does. Compares your sources against each other as well as one at a time, sets a declared trust state from the result, and ties a defined response to that state — so that a source which is wrong while appearing healthy can be caught by its inconsistency with the others, under the conditions set out below.

Who it is for. Operators and owners who act on multi-source information, and the people who answer for what the system did.

Maturity. Implemented in a company-funded reference implementation and exercised in simulation. **Not yet proven in the field:** no sea trial, no hardware-in-the-loop, no live interference, no installed base.

Status. Covered by US provisional patent application **64/144,562**, filed 30 August 2026. One of four applications filed 2026; the fourth, on consensus over categorical judgements, has its own paper. Patent pending; none examined or granted.

A case that made the problem concrete

On 10 May 2025 the container ship MSC ANTONIA ran aground on Eliza Shoals near Jeddah. Salvage took over five weeks.

According to published analyses, the grounding is attributed to false GNSS signals. Pole Star and Windward each reported reaching that conclusion; one of the published assessments characterised the interference as unsophisticated. We have not examined the vessel, its equipment or its voyage data, and we are not in a position to say what the bridge team saw, what any receiver displayed, or who was responsible.

What the public reporting supports is narrower than the story usually told about it, and still worth the attention: an attribution of a major grounding to GNSS interference, with at least one analyst describing the technique as not advanced.

Sources: gCaptain (<https://gcaptain.com/pole-star-confirms-gps-interference-caused-msc-antonia-grounding/>) · Splash247 (<https://splash247.com/msc-ship-aground-off-jeddah-likely-victim-of-gps-spoofing/>) · Lockton (<https://global.lockton.com/us/en/news-insights/cyber-physical-risk-in-the-marine-sector-a-wake-up-call-from-the-msc-antonia>)

The gap conventional alarms leave

Instrument design has made *failure* legible. Equipment drifts, drops out, throws a code, goes out of range, and alarms catch that.

Interference can present differently. The receiver may be healthy, the signal strong, the self-tests passing — and the reported position wrong, reported with the same confidence as a correct one.

Detecting that from inside the affected source alone is difficult, because the reading may be internally consistent. Other approaches exist and work: signal-authentication schemes, multi-antenna and direction-of-arrival techniques, receiver-autonomous integrity monitoring, and inertial coasting all address parts of this, and a well-found bridge may carry several of them. We are not claiming they cannot detect interference.

What cross-source comparison adds is a different kind of evidence — an **additional consistency check** that does not depend on the suspect source being able to diagnose itself. A reading can be individually plausible and still irreconcilable with everything else at the same moment.

Where several sources are combined by averaging or a filter that assumes independent errors, that irreconcilability can be absorbed rather than surfaced: the answer moves toward the outlier by some amount instead of the disagreement being declared.

What this capability does

It compares sources against each other. Each cycle, every source is evaluated against what the others are reporting, in addition to any checks it performs on itself.

It sets a declared trust state. The result is a state the system is in, with a response defined before anything happens, rather than a number in a log or a dialogue box dismissed at three in the morning.

It keeps a tamper-evident record of the decision. What was observed, what was concluded, and what followed, hash-chained so that a later alteration is detectable by recomputing the chain against a separately held anchor. This makes editing **detectable**; it does not make it impossible, and its strength depends on the anchor being held somewhere the record's host cannot rewrite.

What this would be for, on a bridge

This is illustrative. It is not an account of the MSC ANTONIA, which had none of this fitted, and it is not a claim that this capability would have changed that outcome.

A position arrives that does not reconcile with the radar picture, the inertial solution and the chart. Rather than being averaged into a track, the disagreement changes the declared trust state, with a response already attached to that state.

No machine here identifies the interference. What changes is that **the disagreement becomes a declared fact somebody has to deal with**, rather than resolving itself quietly into an average.

The one condition, stated up front

This needs at least one source that does not depend on the one under suspicion.

If the inertial unit is GNSS-aided, the radar is GNSS-referenced and the clock is GNSS-disciplined, one interference event may move them together. Where it does, they agree, they are all wrong, and comparison has nothing to work with. How far a shared reference actually couples them depends on the installation, and this capability does not detect that dependency for you.

So the first question in a deployment is not about our software. It is: **on this platform, what is genuinely independent, and what only looks independent?** That answer is usually more interesting than expected, and better established before a pilot than during one.

What this is not

- **Not a better receiver, and not an anti-spoofing antenna.** It does nothing at the radio layer and does not compete with authenticated signals or multi-antenna techniques. It sits above them and decides what to do when sources disagree.
- **Not a detector of interference as such.** It detects disagreement. Disagreement has many causes, and the capability does not establish which one applies.
- **Not yet proven in the field.** Results to date are from simulation. No sea trial, no hardware-in-the-loop, no live interference, no installed base.
- **Not a granted patent.** Four US provisional applications are on file. None has been examined and none is granted.
- **Not a product you can buy today.** A reference implementation and a demonstration; no certified equipment and no type approval.

We would rather you read that list here than discover it in a meeting.

A first conversation

Bring a nonconfidential example — a case where your sources disagreed and somebody had to decide what to believe. We will work through whether this architecture would earn its place on your platform, what independence you actually have, and what a joint evaluation would need to establish. Recorded data, if it ever helps, comes later and under terms you set.

A first call is a technical discussion, not a pitch, and we will say when the answer is that this would not have helped.

info@rvtechconsulting.com · RV Technology Consulting, Las Vegas, Nevada

Covered by US provisional patent application 64/144,562, filed 30 August 2026. Four US provisional patent applications filed 2026. Patent pending; none examined or granted. The bridge scenario above is illustrative and describes no actual deployment. This document describes capability only and is not a technical specification.