


Know which sensor reports supported the decision, and what was left out

Several reports arrive inside one processing cycle. The system takes the newest one from each sensor, produces a result, and moves on. A week later something is disputed and nobody can say what was used, what was overwritten, or what arrived too late to matter. The data was collected, stored and searchable the whole time, and none of that answers the question.

SENTINEL is an integrity layer for that question. It authenticates each report against a key belonging to that sensor, or to the gateway that signs for it at the junction box; it compares reports of the same quantity under their declared uncertainty and excludes the ones that disagree beyond the configured bounds; and it gives every accepted report a recorded status, including those deferred past the budget or expired, chaining each decision into a digest that a replay must reproduce.



Existing brokers, buses, historians and SIEM platforms keep their role: they carry the signed envelopes unchanged, or receive copies of the evidence. Installed instruments are not replaced; where a device cannot hold a key, a signing gateway at its junction box signs what arrives on its port, and the white paper is explicit about what such a gateway can and cannot attest.

What has been measured

Controlled simulations of a maritime reference application: 1,299 accelerated cycles, 100 ms logical slots, eight signed feeds. Every figure is computed by the benchmark driver and published with its method in the evidence appendix.

Accepted reports with no ledger entry, in the baseline and 50 Hz runs (3,627 and 8,823 accepted)	0
Still open inside the reconciliation window when the run ended (baseline / 50 Hz inertial)	4 / 8
Inertial reports the per-cycle core alone replaced without record, at 50 Hz, in this fixture	67.7 %
The same recorded stream re-executed by a separate C implementation on 1, 2 and 8 real OS threads	1,298 of 1,298
Along-track bias from arrival-time aging, 350 ms link at 6.17 m/s, paired over 391 cycles (2.16 m predicted)	2.17 m

Worst cycle, authentication to analytics inclusive of the accountable layer: baseline /
overload run

7.7 / 13.5 ms

Simulations, not sea trials or hardware qualification. Reproducible replay shows the outcome was determined by the data and the configuration: it is not navigation accuracy, not detection of every attack, and not end-to-end latency. Provenance and agreement are not truth. The white paper marks which behaviour is demonstrated, which is a deployment requirement, and which is proposed and not yet implemented.

Three ways to start

Telemetry replay assessment

two weeks • fixed scope

Send recorded multi-sensor logs. We return what the record shows: reports a per-cycle consumer would have dropped without trace, sources disagreeing beyond their declared uncertainty, and what arrival-time aging costs on your slowest links.

Inline pilot

four to six weeks

The adapter and core run against your testbed, simulator or recorded feed, with a signing gateway on one legacy bus. You get the measurements on your own scenarios and the harness that reproduces them.

Production and OEM integration

by agreement

Deployment on the platform, or the core and SDK embedded in an equipment vendor's own product, with the operator interface built for the use case.

Book a discovery call

Bring the awkward case: what you measure, what already carries the data, and whether this architecture would earn its place. A first call is a technical conversation, not a pitch.



rvtechconsulting.com/contact

The full white paper (25 pages) carries the platform-by-platform questions to ask of your own estate, ten industry cards with their sector standards, the deployment checklist and the limits of every claim. Related technology is described in three U.S. provisional patent applications; these are not granted patents.