DRAKE
MARITIME SERVICES

ANALYTICAL METHODOLOGY

From Signals to Risk: How Maritime Threat Scoring Works

How vessel movement, behaviour, and open-source reporting become a structured risk picture, and why calibrated, explainable scoring matters.

ANALYTICAL METHODS DESK | REV. A · JUN 2026

A modern picture of a single ship is made of fragments. There is where it says it is, and where it actually appears to be. There is its declared identity, its flag, its registered owner, and the longer story of who has owned and operated it over time. There is its track over the past hours and days, the ports it has called at, the speed it has held, the places it has paused. Around all of this sits a wider context: the waters it is transiting, the reporting coming out of that region, and the patterns of behaviour seen in ships like it. On their own, these fragments say very little. The work of threat scoring is turning them into something a person can act on, a structured sense of how much attention a particular vessel or situation deserves and why.

METHODOLOGY NOTE

This is a general account of methodology drawn from long-standing intelligence tradecraft. It does not describe any specific or proprietary system, and it deliberately avoids any formula, weighting, or threshold, because the point is to explain the reasoning, not to provide a recipe. Throughout, one ordinary bulk carrier with years of unremarkable history is followed as the ideas build.

BOTTOM LINE UP FRONT

1. Threat scoring turns scattered signals, position, identity, behaviour, and context, into a structured and defensible sense of how much attention a vessel warrants and why.
2. It works by reading behaviour against a baseline, weighing discriminating indicators rather than counting anomalies, joining them to context, and updating in both directions as things change.
3. The output is only worth anything if it is calibrated to the strength of the evidence and explainable, so a person can see, question, and correct the reasoning behind it.

01 From raw signals to indicators

A raw signal is just a fact about the world. A vessel slowed down. A position report stopped arriving for a few hours. An ownership record changed. None of these facts is, by itself, a sign of risk. Ships slow for weather, traffic, and engine trouble. Position reports drop out because equipment fails. Ownership changes

for entirely ordinary commercial reasons. A fact becomes useful only when it is read against an expectation.

That expectation is a baseline. Every vessel, route, and vessel type has a normal: a usual speed band, a customary track, a typical pattern of ports and pauses. An indicator is a departure from that normal that carries possible meaning. Return to the bulk carrier. A few weeks ago its registered ownership changed, passing through a structure opaque enough that the ultimate owner is hard to name. By itself that is a weak signal, common enough in legitimate trade. Then the ship slows almost to a stop in open water, well off any normal sea lane, in a place where there is no commercial reason to pause. The slowdown means nothing in the abstract. Measured against the carrier's own baseline, it is an indicator, because the ship is doing something its own history says it does not do.

This is why behavioural indicators matter more than any single status fact. A flag or an owner is a static attribute. Behaviour is what a vessel actually does over time, and behaviour against a baseline is where intent and anomaly tend to show themselves. A good indicator is also discriminating, which is worth dwelling on. Some departures from normal are consistent with many explanations and therefore tell you little. Others are consistent with only a few, and those are the ones that move a judgement, because they help separate one explanation from another rather than adding to a pile of vaguely interesting facts. A ship drifting with a genuine engine fault and a ship loitering for some other reason can look similar at first glance, so the useful indicators are the ones that begin to tell those cases apart.

One kind of indicator is easy to overlook: the absence of something that should be present. Reporting tends to fixate on events, on things that happen, while a silence where there should be noise can be just as telling. As the bulk carrier slows, its position reporting also becomes intermittent. That gap could be ordinary equipment failure, or it could be a choice. The distinction matters, because a signal lost to a broken antenna and a signal switched off deliberately point in very different directions, and separating the two is part of the reasoning rather than a footnote to it. A structured approach has to make room for the meaning of what is missing, not only the meaning of what appears.

02 Why context changes everything

An indicator never stands alone, because the same behaviour means different things in different settings. The bulk carrier's pause in a calm, low-risk sea would be mildly curious. The same pause in waters where vessels have recently been approached, or close to infrastructure that has been a target, is something else entirely. Context is what converts a raw anomaly into a risk-relevant one.

Context has several layers. There is the geography and threat environment of the area being transited. There is the recent reporting about what is happening there. And there is the profile of the vessel itself, which in some situations is the single biggest driver of whether it is at risk at all, since who owns a ship, what flag it flies, and where it has recently been can matter more than anything it does on a given afternoon. Hold the carrier's two facts together: an opaque recent ownership change, and a slowdown with intermittent reporting in a stretch of water with a recent history of approaches. Either fact alone is thin. Together, and placed in that setting, they describe a situation that warrants a closer look. A sound method keeps the indicator and the context joined, because an alarming behaviour in a benign setting and a benign behaviour in an alarming setting are different problems, and each deserves to be seen for what it is.

03 What open-source reporting adds

Much of that context comes from open-source reporting, which deserves treatment as a genuine input rather than only as a hazard. It is where a great deal of the picture actually originates: regional incident and advisory reporting that tells you which waters are dangerous this month, ownership and sanctions information that fills in a vessel's profile, local and trade reporting that surfaces things no sensor will show. For the bulk carrier, open sources might reveal that ships under similar opaque ownership have recently been used in a particular trade, or that the area it paused in saw an approach days earlier. That is real signal, and a risk picture built without it would be blind to most of its own context.

Open-source reporting also has to be handled with care, because it is uneven in quality. Some of it is confirmed, some is rumour, and some is deliberately misleading. Long-standing analytic practice keeps two separate questions in view: how reliable is the source of a piece of information, and how credible is the information itself. A claim from a consistently trustworthy source still has to be weighed against everything else known, and a single unverified report should not, on its own, swing a picture. The discipline is to draw fully on open sources for what they contribute while never treating any one of them as settled fact.

04 Building the structured picture

Bringing indicators and context together is an exercise in disciplined judgement rather than arithmetic, and a few principles hold it together.

The first is the separation of evidence from inference. What the data actually shows, that the carrier slowed here and changed ownership then, is different

from what an analyst concludes it might mean. A sound method keeps the observed facts distinct from the interpretation laid over them, so a reader can see both and can disagree with the interpretation without losing the facts.

The second is weighing rather than counting. A risk picture is not improved by tallying how many unusual things a vessel has done, because ten weakly meaningful anomalies can matter less than one strongly meaningful one. The discriminating power of each indicator, and how much each actually shifts the picture, is a matter of reasoned weight. The carrier's single opaque ownership change plus one well-placed, well-timed slowdown can outweigh a long list of minor oddities on another ship that add up to nothing.

The third is updating as things change. A risk picture is a snapshot of a moving situation, and new information should revise it in a consistent direction, raising the assessed concern when corroborating signals arrive and lowering it when innocent explanations are confirmed. If the carrier resumes a normal track and its reporting comes back cleanly with a plausible account, the concern should fall. Equally, old signals should decay, since an anomaly from two weeks ago should not carry the same force as one from two hours ago. A picture that only ever ratchets upward, and never relaxes when the evidence warrants, is not tracking reality.

The fourth is the discipline against bias. The mind reaches for the familiar explanation, fixes on a first impression, and sees confirmation where it expects to. Structured approaches exist partly to slow that down: to make assumptions explicit, to ask what evidence would point the other way, and to keep alternative explanations alive rather than collapsing onto the first plausible story. The value of a structure is not that it removes judgement but that it makes judgement visible and checkable.

05 Where scoring goes wrong

The case for doing this well is clearest in the recognisable ways it fails.

PRECISION	A false sense of precision, where a confident-looking score rests on thin evidence and carries more authority than it has earned.
CIRCULAR	The self-confirming picture, where an early read quietly filters later evidence to fit, so the carrier, once flagged, is only ever seen through suspicious eyes.
PROXY	Proxy bias, where a method scores what is easy to measure rather than what matters, so the loudest or most-reported vessel looks riskiest simply because there is more noise about it.

STALE	Stale weighting, where old anomalies never lose their force and ships accumulate permanent suspicion they cannot shed.
OPAQUE	The opaque output, where a score arrives without its reasoning, so a wrong call goes unchallenged because no one can see what produced it.

Each of these is a failure of calibration or explainability, which is why those two qualities are the point rather than an afterthought.

06 Why calibrated, explainable scoring matters

The output of all this is, in some form, a sense of how much concern a vessel or situation warrants. Whether expressed as a level, a band, or a phrase, two qualities decide whether it is worth anything.

Calibration means the expression of concern matches the actual strength of the evidence. A picture that reads as highly alarming on thin evidence, or as reassuring on serious evidence, misleads precisely when it matters most. Part of calibration is honesty about uncertainty, including a frank account of what is not known and what would change the assessment. There is a well-known temptation to attach more precision to a judgement than the underlying information can bear, and resisting it is part of getting calibration right. For the bulk carrier, the honest read is not “this ship is hostile.” It is “this vessel’s behaviour and profile, in this place, warrant a closer look, and here is exactly why,” held with appropriate uncertainty.

> **A figure presented without calibration or explanation asks for trust it has not earned.**

Explainability means a person can see why the picture reads the way it does. A score that arrives without its reasoning is hard to trust, hard to challenge, and hard to act on, because the decision-maker cannot tell which indicators drove it or whether the logic holds. A judgement tied clearly to the indicators and context behind it can be examined, questioned, and corrected. It also makes the picture auditable after the fact, so that when an assessment turns out right or wrong, the reasoning can be reviewed and the method improved.

This is also the right place to be clear about the division of labour. A structured approach is good at what people are bad at: watching many vessels consistently, day and night, and surfacing the few that deserve attention without tiring or playing favourites. A person is good at what a structure cannot do: bringing in context the data never captured, judging when an indicator is misleading, and making the actual decision. Scoring done well plays to the first and defers to the

second. It exists to help a person decide, often under time pressure and with incomplete information, by giving them the clearest and most honest picture the available information can support. A figure presented without calibration or explanation asks for trust it has not earned, and in a domain where the cost of being wrong is measured in ships, cargo, and lives, that is a poor trade. The aim of turning signals into risk is not to replace human judgement with a number. It is to give human judgement something solid to work with.

CLASSIFICATION	DESK	DATE	BASIS
ILLUSTRATIVE ANALYSIS SAMPLE	ANALYTICAL METHODS DESK	REV. A · JUN 2026	GENERAL METHODOLOGY