Safe Harbours Australia
Expert Maritime Advice

NAVIGATION & RESILIENCE

Redundancy is not resilience

JULY 2026 • BY CAPTAIN MATT SHIRLEY • 13 MIN READ

In port pilotage, what happens when GPS lies? Here is how we assess what a port actually depends on, and why the answer is worth more than it costs.

A modern container port is a positioning and timing system with wharves attached. Image: iStock.

Consider a manoeuvre performed daily in constrained ports around the world.

A container vessel, many times the size the channel was originally cut for, is worked into a berth the pilot cannot see, in a basin the pilot cannot see, past infrastructure that was designed to withstand a collision from a ship a fraction of its mass. There are no visual references available at the decision points that matter. The margins are measured in metres. The only thing making the transit possible is a portable pilot unit taking a position from a satellite constellation.

The organisations conducting these manoeuvres are, in general, the most sophisticated operators in the industry. They are the ones who have invested. Two units rather than one. Inertial motion sensors integrated into the units. Two pilots on the bridge instead of one.

Procedures and training built to a standard that receives nothing like the recognition it deserves.

And that investment, which represents the current state of the art, still contains the assumption that undoes almost every port resilience plan we have seen.

Two portable pilot units are not redundancy. They are duplication. If both take their position from the same constellation, both fail at the same moment, in the same direction, for the same reason. Redundancy only counts when the backup fails differently. Lose the signal systemically and everything you were relying on goes together, in step, and quietly.

This is not a criticism of the operators who have gone furthest. It is the opposite. It is an observation that the leading edge of current practice has reached the limit of what duplication can buy, and that the next step is a different kind of question entirely.

Ports ask *do we have a backup?*

The question that matters is *does the backup fail for a different reason than the primary?*

We have written elsewhere about why GNSS disruption is a present exposure for ports rather than a projected one, and we will not repeat that ground here. The purpose of this piece is narrower and more useful: to set out, in operational terms, how we assess a port's positioning, navigation, and timing dependency, what we look for when we do, and what a port stands to gain from asking.

The port is not a navigation problem

The first structural error is one of ownership. In most ports, GNSS is understood as a marine department issue. It sits with the Harbour Master, alongside charts and aids to navigation, and it is treated as a question of how ships find their way in.

It is not a navigation problem. It is a whole-of-port dependency, and it does not respect the organisational chart.

A modern container port is a positioning and timing system with wharves attached. Satellite-derived position and time underpin:

- vessel navigation, ECDIS, AIS, portable pilot units, and dynamic under-keel clearance
- vessel traffic services and the integrity of the traffic picture
- tug positioning and escort towage
- automated and semi-automated container handling, including straddle carriers, automated stacking cranes, crane anti-collision, and container position tracking
- yard inventory accuracy and terminal operating system spatial data
- truck and rail gate systems, geofencing, and appointment scheduling

- network time synchronisation across terminal systems, CCTV, access control, SCADA, and security logging
- supply chain visibility, arrival prediction, and berth scheduling

Note what happens across that list. It begins in the marine department and it ends in IT, security, and commercial scheduling. There is no single manager in any Australian port whose remit covers all of it. Which is why nobody owns it, and why the dependency has been allowed to accumulate without anyone counting it.

The timing half of the problem is the more neglected of the two. Position failure is loud. Someone notices the ship is in the wrong place. Timing failure is quiet. Systems drift out of synchronisation, logs disagree, transactions fail intermittently, and the diagnosis takes days. Most ports have never traced where their time actually comes from.

The failure mode that should worry you is the one that looks fine

Systems fail in three ways, and only one of them is genuinely dangerous.

Systems that **fail safe** stop and say so. A unit that loses lock and blanks the screen has done its job.

Systems that **fail visibly** degrade in a way an operator can see and account for. Accuracy drops, an alarm sounds, confidence is reduced, and the human adjusts.

Systems that **fail silently** continue to produce an output that looks entirely plausible and is wrong. A spoofed position does not announce itself. The display is clean. The prediction vector is smooth. The ship's shape on the screen is lying, and it is lying convincingly.

Automated systems have no scepticism. They do not know they are being deceived. And the people standing next to them, in port environments where the visual cues that would betray the error are frequently unavailable, increasingly do not have the reference points to catch it either. A generation of mariners has now grown up navigating in an environment where the screen has always been right. They are highly capable users of the system. They are also, entirely rationally, disposed to believe it.

Silent failure is the dominant risk in a modern port, and almost no port resilience plan addresses it, because almost every plan is built around the question *what do we do when we lose the signal?* rather than *how would we know if we had not lost it, but it had started lying to us?*

The six things we assess

We assess a port's PNT resilience against six principles. They are deliberately not technology categories. They are properties a resilient port possesses, and each one can be evidenced, tested, and found wanting.

1. Layering

No single source of truth for position or time.

We look for how many genuinely distinct sources of position and time the port has, and how many it can actually use. Not how many receivers, how many sources. A port with four receivers on one constellation has one source.

What we ask: How many independent means exist to establish a vessel's position in the channel, the swing basin, and alongside? Where does port time originate, and how many hops from a satellite is it? If the primary source is removed, what is the second, and is it real or is it a line in a plan?

2. Independence

Backups that fail differently.

The one that catches everybody, including the operators who have invested most. We test whether the fallback shares a failure mode with the primary. A backup that fails for the same reason, at the same moment, is not a backup. It is a second copy of the problem.

What we ask: Do the primary and secondary systems share a constellation, an antenna, a receiver chipset, a power supply, a network path, or a time source? If a wideband jammer is operating at the port boundary, which of the two survives, and by what mechanism? What is the common-mode failure that takes both?

3. Integrity

So we know when positioning or timing is lying.

Detection, not just backup. A port that can survive a signal loss but cannot detect a spoof is defended against the lesser of the two threats.

What we ask: Does the port have any independent means of cross-checking that the position it is being shown is true? Does anything monitor the RF environment? Would an interference event be recorded, or merely experienced? What cross-check is available to the pilot at the moment of the manoeuvre, when there are no visual references and no time to consult anybody?

4. Human capability

People who can question the system.

Detection is ultimately a human act, and it depends on somebody having a reason to doubt. This is not a training tick-box. It is a question of whether the port has deliberately preserved the capacity for scepticism.

What we ask: When did a pilot last conduct a transit without the unit, in the simulator or on the water? What independent cues remain available in this port, and does anyone still

use them? Is degraded-mode operation practised, or merely documented? Would a pilot who said “I do not believe the position” be supported, or would they be asked what the delay would cost?

5. Local control

A port that can act without waiting.

Resilience that depends on somebody else’s decision is not resilience. If the port cannot detect, decide, and act inside its own boundary, it is exposed to the response time of a party with different priorities.

What we ask: Who inside the port can call a halt, and on what trigger? What positioning or timing capability does the port own and physically control, as opposed to consume as a service? If the national or vendor response takes 48 hours, what is the port doing for those 48 hours?

6. Governance

Somebody has to own resilience.

The one that determines whether the other five ever get built. In our experience this is the finding that recurs most consistently: PNT dependency is a whole-of-port risk for which no single role is accountable, and so it sits, uncounted, on the boundary between marine, terminal, IT, and security.

What we ask: Which named role owns PNT resilience? Does it appear on the corporate risk register, and if so, how is it described? Who would be accountable, at board level, if positioning integrity were lost during a transit? Has anyone quantified the exposure, or has it simply never been asked?

How we find out

The six principles are the criteria. The following is the process that produces the evidence to test them against.

Step one: PNT dependency mapping. Systematic identification of every system, process, and design element in the port that depends on satellite-derived position or time. Marine, terminal, landside, and security. This is deliberately exhaustive, because the second-order dependencies are the ones that hurt. The port knows the pilot unit needs GNSS. It usually does not know that its access control audit log does.

Step two: failure mode and consequence analysis. What actually happens, operationally and commercially, under degradation, denial, and spoofing, across durations from minutes to days. Each dependency is classified as fail-safe, fail-visible, or fail-silent. The fail-silent register is, almost always, the document that changes the conversation.

Step three: control identification. Where the dependency can be removed, diversified, layered, or detected. This is where the six principles turn into a control set: what would give this port a genuinely independent second source, what would give it integrity monitoring, what would give it local control.

Step four: degraded-mode operating concept. The port defines, in advance and in writing, how it operates at reduced PNT availability. What movements continue, what is suspended, what throughput is retained, who holds the authority to stop, and on what trigger they exercise it. Written before the event, not improvised during it, and not written by the party whose schedule depends on the answer being “carry on”.

Step five: human detection and response. Scenario-based simulation in which the position is wrong and does not say so. This is the step ports skip, and it is the one that reveals the most. It is straightforward to write a procedure for the loss of a signal. It is considerably harder to build the professional instinct that says *something here is not right* when every instrument agrees and the screen looks perfect.

The output is a resilience profile: the port scored against six principles, with a control set, a degraded-mode concept, and a governance recommendation naming who owns it. Not a technology recommendation. We hold no product and we sell no mitigation, which is why we are able to tell a port that the answer is a procedure and a person rather than a purchase, on the occasions when that is true.

The upside, and the trap inside it

Everything above is framed as risk. It should not be, entirely. Positioning resilience is also an enabler, and for a port being designed now, it is a strategic one.

Surety of supply. A nation with a single container port has no fallback. There is no alternative berth, no second terminal, no other channel. In that configuration, a positioning or timing failure is not a port incident. It is a national supply chain event, and the exposure is measured in the daily value of the trade that stops. Resilience purchased at the design stage is, in that light, not a safety expense. It is continuity of supply, and it is the cheapest form of it available.

Design headroom. Assured, independently verifiable positioning changes what a designer is entitled to assume. Channel geometry, swing basin dimensions, under-keel clearance policy, and vessel size envelopes are all set against positional uncertainty. Reduce the uncertainty and the design has room it did not have before.

Automation viability. Terminal automation is a positioning and timing bet. The more automated the concept, the more it depends on continuous, accurate position and synchronised time, and the more expensive an undetected failure becomes. A port that intends to automate and has not resolved its PNT integrity has bought the dependency without the assurance.

Operational consistency. Delays and suspensions driven by degraded positioning confidence are avoidable. Fewer of them means more predictable arrival windows, berth utilisation, and turnaround, and that predictability propagates outward into terminal and landside scheduling.

Now the trap, and it is one we would rather name than sell around.

The dividend can be banked as margin, or it can be spent as throughput. It cannot be both.

If a port invests in resilient positioning and then immediately uses the improved confidence to tighten its under-keel clearance buffers, narrow its weather windows, and squeeze larger vessels into the same channel, it has converted its new safety headroom into productivity. It is back where it started, standing on the same thin margin, with a new dependency underneath it and a resilience investment that has already been consumed.

That is not an argument against the investment. It is an argument for being explicit about what the investment is for. The decision to bank the dividend as margin or spend it as capacity is a board decision, not an operational one, and it should be made deliberately, once, and in writing, rather than absorbed quietly into a series of individually reasonable operational adjustments.

We are a safety consultancy. We will tell a port the efficiency case is real. We will also tell it that an efficiency gain purchased with a safety margin is not a gain.

Why the design stage is different

Everything above can be done to an existing port, and should be. But there is a category of port where this work has an entirely different value, and a hard expiry date.

A port being designed now is, at this moment, deciding its PNT dependency profile for the whole of its operating life. Decisions being taken about VTS architecture, terminal automation levels, timing infrastructure, network design, aids to navigation, and the positional accuracy assumed in the channel and the yard are, cumulatively, an unexamined bet on the continuous availability of a signal broadcast from twenty thousand kilometres away at roughly the received power of a distant car headlight.

Resilience designed in is an order of magnitude cheaper than resilience retrofitted. Once a terminal operating concept assumes centimetre-level accuracy, degraded-mode operation is no longer a design choice. It is an unfunded liability that somebody will discover at commissioning, and it will be absorbed, as these things always are, into a procedural restriction rather than a design feature.

The window for asking these questions closes at design freeze. After that, the answers cost concrete.

The question we are actually asking

Not *will GNSS fail?* It already has, repeatedly, and the record is public.

The question is whether the port would know, whether it could operate through it, whether it could act without waiting, and whether anybody in the organisation is accountable for the answer.

Most ports we speak to have not been asked. That is not a criticism. Nobody has been asking. It is a question that has fallen between the marine department, the terminal, and the IT function, and remained there quietly accumulating consequence.

The technology exists. The skills exist. The knowledge exists. What is generally missing is somebody prepared to map the dependency, count the exposure, and put a name against it.

That is the work.

Captain Matt Shirley is CEO of Safe Harbours Australia and a former marine pilot with over 5,000 pilotages, including fourteen years at Port Hedland. Safe Harbours Australia is an independent maritime consultancy led by three Master Mariners with over 110 years of combined frontline experience. We hold no products and have no commercial relationships with equipment suppliers, system integrators, or terminal operators.

Related service: Navigation Resilience & GNSS Risk Assessment

WRITTEN BY

Matt Shirley

CEO & Co-founder, Safe Harbours Australia

Safe Harbours Australia

Independent maritime consultancy delivering trusted judgement and real-world impact across port safety, risk, and operations.

CONTACT

consultants@safeharbours.com.au

+61 438 752 500

LinkedIn

www.safeharbours.com.au