Legacy Technology in Transport: More Than “Old Tech”

Introduction: More Than “Old Tech”

The word **“legacy”** is one of those terms everyone uses but rarely defines. In IT, it usually has a specific meaning: a piece of software or hardware that’s out of support, fragile, and increasingly risky to maintain. There’s even policy around it — for example, the UK Government’s Legacy IT Risk Assessment Framework, which provides a structured, risk-based method to assess aging systems. In that framework, a system is considered “legacy” when indicators like vendor support, available skills, compatibility, and security start to degrade. The typical response in IT is clear: replace it, migrate away from it, or retire it.

But in transport and operational technology, **legacy** is a more complicated story. Consider a few scenarios where the label might apply:

- A **train signalling system** built in the 1970s still governs safety-critical operations.
- An **automotive ECU** (engine control unit) running on an old 8-bit architecture continues to be deployed in new vehicles.
- A **maritime radar interface** cannot integrate with modern port IT systems.
- A **supplier’s software toolchain** disappears when their contract ends.
- A platform’s **safety certification** depends on a standard that has since been withdrawn.

In one context, *legacy* might simply mean “old.” In another, it means “unsupported” or “unmaintainable.” Elsewhere, it might just mean “in the way” of progress. This ambiguity makes it difficult to have clear conversations about risk, investment, and assurance – people may use the same word to describe very different problems.

In transport, legacy isn’t just about age or vendor support. It’s about embeddedness, certification, integration, and dependency. Many of the systems we rely on are long-lived by design. They operate inside cyber-physical environments and are bound by lifecycles that span decades, with safety and regulatory requirements to match. These systems depend on infrastructure and support tools that often predate today’s connectivity and cybersecurity assumptions. That doesn’t mean legacy risk doesn’t apply – it means we need to define it much more carefully for this sector.

This article proposes exactly that. We will start with a taxonomy of transport technologies to clarify what we are dealing with. Then we will examine the strengths and limits of the UK Government’s legacy IT risk assessment framework, and introduce a scoring approach tailored to long-lived, safety-critical sectors like transport. The goal is not to replace existing risk methods, but to create a clear, shared language for when technologies drift into fragility – even if they still function. Managing legacy risk well requires us first to be precise about what we are actually talking about.

Why Definitions Matter

In traditional IT, identifying legacy systems is relatively straightforward – for example, a database that the vendor no longer supports, or a server that cannot meet modern security standards. The fix can be comparatively straightforward too (upgrade or decommission it). In transport, however, the stakes and constraints are different. Changing or replacing technology can be far more costly and complex. Here’s why clear definitions matter before taking action:

- **Certification and assurance:** Many transport systems are safety-critical. Changing a component may invalidate existing safety cases or require extensive re-certification. For instance, updating an avionics software could mean re-doing months of safety tests. Without a clear definition of "legacy," organizations might ignore a problem technology because it's still certified – or conversely, try to rip it out without realizing the safety case impacts. A precise definition helps balance safety assurance with modernization.
- **Infrastructure coupling:** Transport platforms never operate in isolation. New trains must interface with old signalling systems; new cars should work with existing road infrastructure. If "legacy" is defined simply as "anything old," we miss the real issue: *compatibility*. A system might be decades old but still perfectly reliable – except that everything around it has moved on. Defining legacy in terms of inability to integrate or comply with current standards captures this risk, whereas a simplistic age-based definition would not.
- **Operational inertia:** Planes, trains, and automobiles are designed for service lives spanning decades. Over that time, the support tools, spare parts, and even the expertise to maintain them can fade away before the system itself retires. A component might be *functionally* reliable yet *structurally* fragile because the ecosystem around it (parts, tools, knowledge) is eroding. If we don't clearly define legacy status, organizations might underreact (continuing business-as-usual with a fragile system) until a crisis forces a change at the worst time.
- **Fragmented ownership:** In sectors like rail or aviation, multiple stakeholders own different pieces of the technology puzzle. For example, a railway's signalling hardware might be owned by one company, the trains by another, and maintenance by a third. Each has different incentives and timelines. A "legacy" system in this context is not just a technical issue – it's a contractual and organizational one. A clear definition is needed so all parties understand the risk. Otherwise, we risk overreaction (one stakeholder pushing for costly replacement) or underreaction (everyone assuming someone else will handle it).

The result of unclear terminology is a paradox: technologies that seem reliably in service can actually be fragile underneath. Without an agreed meaning of legacy, organizations may either ignore creeping fragility or push for premature, expensive replacement. Both extremes introduce unnecessary risk. Having a precise definition of legacy – tuned to the transport context – enables balanced decisions about when to maintain, when to modernize, and how to manage the transition safely.

A Taxonomy for Transport Technology

To have a meaningful conversation about legacy in transport, we first need to be clear about **what** technology we're discussing. Unlike a typical IT environment – where you can easily point to servers, databases, or applications – transport technology is a *system-of-systems*. In other words, it's a web of interdependent components: from tiny sensors and chips, to vehicles and infrastructure, all operating together to move people and goods safely.

We propose a **three-layer taxonomy** of transport technology (adapted from NIST systems engineering guidance) to clarify these layers. Every piece of transport tech can be thought of as belonging to one of three interrelated categories:

1. **Cyber-Physical Systems (CPS) – The Platforms Themselves:** This is the most visible layer: the trains, planes, automobiles, and ships – along with all the onboard technology that makes them run. These platforms tightly integrate software and hardware, creating feedback loops between the digital and physical world. For example, in an automobile, the engine control unit (ECU) reads sensor data and adjusts mechanical actuators in real time; in an aircraft, flight control software continuously modulates control surfaces based on sensor inputs.
   Examples: Automotive CPS include engine and transmission ECUs, the CAN bus network, and advanced driver-assistance system (ADAS) logic. Rail CPS include onboard signalling equipment, braking and traction control systems. Aviation CPS cover avionics networks (like ARINC or AFDX) and flight control laws. Maritime CPS encompass navigation suites and propulsion control systems.

Legacy considerations: Interestingly, CPS legacy status is shaped less by age and more by obligations to other systems. A new component might have to support decades-old protocols for diagnostic or compatibility reasons. Certification regimes can "lock in" certain technologies – for instance, a train control system certified to a 1970s standard cannot be altered easily without losing its approval. Safety case evidence (the documented proof that a system is safe) might be invalidated if you swap out a component for a newer one. Thus, backward compatibility often takes precedence over modernization – not because older tech is superior, but because breaking a working safety or compliance contract is often too costly. When assessing legacy at the CPS layer, we must consider these **external dependencies** on infrastructure and standards, not just the platform's internal condition.

2. **Operational Infrastructure – What Enables Movement:** No vehicle or vessel operates alone; it interacts with a whole infrastructure around it. This second layer includes all the fixed, distributed systems that enable coordinated, safe, continuous operation across a network.
   Examples: In rail, this means trackside signals, interlocking systems, power electrification for trains, and control center SCADA systems. In aviation: airport instrument landing systems (ILS), radar and ADS-B ground stations, and air traffic control networks. On the road: intelligent transport systems (ITS) like smart traffic lights, tolling systems, and the network of EV charging stations. In maritime: vessel traffic services, harbor radar, and port automation systems. Many of these fall under **Critical National Infrastructure (CNI)** because a failure can ripple across regions or an entire country.

   Legacy considerations: Legacy risk in infrastructure often arises from **inflexible or proprietary interfaces** and long technology lifecycles. For instance, an old railway signalling protocol might not easily interface with modern digital train systems, making upgrades complex. Replacing infrastructure is usually a massive project – long procurement cycles, regulatory hurdles, and coordination among many stakeholders. As a result, it's common to see very old infrastructure still in use (because "if it isn't broken, don't touch it"), even if it's becoming hard to maintain. However, as the world becomes more connected and demands more data (for example, predictive maintenance or real-time tracking), these legacy infrastructures become points of fragility. A single legacy subsystem – say an ancient control unit in a power substation – could limit the adoption of modern, secure monitoring across the network. And when infrastructure does fail, the impact is broad: a **single signal failure** can halt an entire rail corridor, or a radar outage can affect dozens of flights. This is why legacy fragility in operational infrastructure is especially consequential.

3. **Engineering & Support Tooling – What Builds and Maintains the Ecosystem:** The third, often overlooked layer includes all the tools and environments used to **design, build, test, deploy, and maintain** the first two layers. In other words, these are the engineering workstations, test rigs, configuration software, maintenance consoles, and so on that keep the transport system running behind the scenes.
   Examples: In automotive, this could be the laptop software and hardware rigs used to flash (update) ECU firmware in a factory or service centre, calibration tools for engine tuning, and the data logging equipment for troubleshooting. In rail, examples include depot SCADA systems that manage maintenance depots, specialized test benches for signalling systems, wheel lathes for refurbishing train wheels, and even older laptops kept around because they're the only ones that can run a particular diagnostics program. Aviation has maintenance, repair, and overhaul (MRO) rigs, avionics simulation benches to test aircraft systems, and ground service equipment software. Maritime might involve drydock automation controls, ballast tank calibration tools, and so on.

   Legacy considerations: Support tooling can become legacy even if the primary system it services is not. It's not uncommon that an aeroplane or train is fine, but the *tool required to update it or test it* is obsolete or no longer supported by any vendor. This creates an **invisible dependency risk** – you might not realize a critical piece of equipment hinges on a Windows XP laptop tucked away in a maintenance garage until that laptop dies. Other challenges include lost licensing or documentation for niche tools (perhaps the vendor went out of business), reliance on a single retiring expert who knows how a tool works, or outdated security in these support systems (they might not get the same cybersecurity attention since they aren't "operational" systems per se). Many enabling tools were created decades ago and assumed to be used on isolated networks or by trusted technicians, so they may lack modern authentication or encryption. As those assumptions become invalid in today's environment, the tools themselves pose security and reliability risks. And because these support systems are often under the radar (managed by engineering teams, not IT departments), they can be **easily overlooked** in risk assessments. They only get noticed when something goes wrong – for example, if a testing rig fails and suddenly no new trains can be certified until it's fixed.

This three-part taxonomy – CPS, Operational Infrastructure, and Engineering/Support Tools – helps us analyse legacy issues in context. It's deliberately more nuanced than the typical division of "IT vs OT (Operational Technology)." In fact, many things that would traditionally be lumped into *OT* are split between Infrastructure and Support in our model. This separation is important: a rail signaling interlocking system and a depot's calibration tool might both be considered "OT," but their roles, lifecycles, and legacy challenges are very different. By distinguishing the *platforms*, the *infrastructure* they rely on, and the *tools* that support them, we can better pinpoint where fragility lives – and how it might spread through the transport ecosystem.

IT, OT, and Why We Prefer the Term "Cyber-Physical Systems"

Many industries use the dichotomy of Information Technology (IT) vs Operational Technology (OT) to discuss digital risk. Typically, "OT" refers to systems that monitor or control physical processes (like factory control systems, power grid SCADA, building management systems, etc.), as opposed to "IT" systems that deal with data and business processes. In cybersecurity conversations, OT is often treated as a distinct realm requiring special consideration.

However, in complex sectors like transport, the label **OT** can obscure more than it clarifies. Here's why:

- **"OT" lumps together too much.** Under a broad OT umbrella, you might include a train's signaling controller, a railway station's HVAC control, and a maintenance depot's diagnostic laptop. These are very different systems, with different criticality and lifecycles, yet a high-level IT/OT inventory would bucket them all simply as OT. That makes it hard to prioritize and manage specific risks – the nuance gets lost.
- **"OT" assumes a clear boundary with IT.** In modern transport systems, that boundary is blurry. Trains and aircraft carry IT components (servers, databases, standard operating systems) on board; airports and highways have IT networks interfacing with industrial controllers. A lot of transport tech blends enterprise IT and embedded OT. If we rigidly separate them, we might miss how an IT change (say a Windows update) could affect an operational system (like an airport baggage handling PLC).
- **"OT" hides dependencies.** By treating OT as a monolithic category, we might ignore the interplay between different types of operational systems. For example, a support tool may be as critical to safety as a live control system, but if both are simply labelled OT, one could easily overlook the tool's significance. Classic OT risk management often focuses on the running systems (e.g. keep the power plant control system secure) and forgets the support systems (e.g. the laptop that configures the controller). In transport, these hidden dependencies are everywhere.

For these reasons, we emphasize the term Cyber-Physical Systems (CPS) for the platforms (layer 1 of the taxonomy) and break out the rest of what might be called OT into Infrastructure and Support layers (layers 2 and 3). By doing so, we get a clearer, more nuanced picture of where legacy-related fragility might emerge. We move beyond a simplistic IT/OT divide and instead look at distinct but interdependent subsystems: the vehicles/vessels themselves, the fixed infrastructure around them, and the tools that build and maintain both. This way, when we discuss "legacy," we can specify whether we mean a legacy issue in the vehicle, in the track/road/sky infrastructure, or in the maintenance toolkit – because each of those has different implications and possible remedies.

What the UK Legacy IT Framework Tells Us – And What It Misses

The UK's Legacy IT Risk Assessment Framework (from the Central Digital and Data Office) is a recent effort to systematically identify high-risk legacy systems across government. Notably, it defines "legacy" not simply by age, but by sustainability factors. In essence, a technology becomes legacy when it is no longer:

- **Supported by its vendor or maintainer.** (For example, the manufacturer won't provide patches or spare parts.)
- **Compatible with current standards or practices.** (It can't meet modern requirements or integrate with contemporary systems.)
- **Feasible to integrate into future architectures.** (It would hold back planned upgrades or migrations if kept in place.)
- **Safe, secure, or cost-effective to operate.** (It lacks essential security controls, poses safety risks, or is very expensive to keep running.)
- **Understandable by the available workforce.** (Perhaps it's written in an obsolete language or only a few veterans know how to maintain it.)

This framing is powerful because it highlights things like vendor support, skill availability, and ecosystem compatibility as key drivers of fragility, rather than just focusing on how old the system is. In other words, it shifts the conversation to, "Is this thing *sustainable* to keep using?" If the answer is no on multiple fronts, it's legacy.

However, when we try to apply this framework to transport technology, we hit a few challenges. Some underlying assumptions of the IT model don't cleanly map to transport/OT environments:

- **Systems aren't discrete or neatly bounded.** In government IT, you might assess a payroll system or a records database – a clearly defined "system" with ownership and boundaries. In transport, what exactly is the "system"? Is it a single locomotive? The signalling network for a whole city? The concept of an isolated system breaks down because everything is interconnected. A legacy component can span or affect multiple systems-of-systems.
- **Ownership is fragmented.** The UK framework assumes you can assign an owner to a system who will take responsibility for fixing it. Transport tech is often owned and operated by a patchwork of organizations. A railway signalling system might involve a public infrastructure manager, multiple train operating companies, and vendors. So, who answers the question "is this system legacy?" Each stakeholder might see only a piece of the puzzle.
- **Legacy risk is not only operational – it's also certification-driven.** A key impact of legacy in transport is on **assurance**. A system might be perfectly operational and safe today but considered legacy because the standards it was approved under are outdated or the evidence can't be reproduced. Traditional IT risk frameworks focus on operational risk (failures, cyber breaches, cost of downtime). In contrast, in transport you also worry about regulatory risk: "If I touch this, will I lose my certification to operate?" That factor is mostly missing in generic IT legacy assessments.
- **Ecosystem viability isn't directly measured.** The UK framework's criteria are a great start (support, integration, etc.), but in transport we often have an entire *ecosystem* that can become legacy. For example, you might score an air traffic control system as fine because it's supported and secure, but you could miss that the specialized test equipment for it is no longer made. The **support ecosystem** (tools, test rigs, suppliers, skills) can be decaying even if the main system seems OK. So, we need to measure aspects that aren't just "system X itself," but the health of everything around system X that it needs.

In summary, the UK legacy IT framework gives us a valuable lens to identify when technology is becoming unsustainable. It tells us what kinds of questions to ask. But to apply it meaningfully in transport, we must extend its dimensions and sometimes shift the unit of analysis (as we'll discuss next). Instead of looking only at a "system" in isolation, we need to look at technology types in context. The framework isn't wrong – it's just not sufficient by itself for the transport sector's complexity.

Shifting the Lens: From Systems to Technologies

One key conceptual shift we propose is moving from a system-centric view to a technology-centric view of legacy. In practice, that means instead of always asking "Is this system legacy?", we also ask a broader question: "Is this technology type approaching legacy status in one or more of its contexts?"

This shift is subtle but profound. Why do it? Because many technologies in transport are used across multiple systems and domains. A few examples:

- The CAN bus protocol (Controller Area Network) might be used inside road vehicles (cars, trucks), in depot diagnostic tools for those vehicles, and even in some rail or industrial systems. In one context, CAN bus might still be fine; in another, it could be a bottleneck (e.g., not secure or fast enough for new needs). If we only ever assess "System A on train" or "System B on car," we might miss that *CAN bus as a technology* is becoming a legacy problem across the board.
- A hardware interface like JTAG (used for low-level debugging/programming of chips) appears in avionics testing equipment, automotive manufacturing, and elsewhere. Perhaps JTAG is being replaced by newer interfaces in some industries. A particular piece of test equipment might not raise flags as legacy, but the technology it relies on (JTAG) could be on its way out, making future tools incompatible. A system view might miss that; a technology view would catch it.
- An old communication standard like ARINC 429 (widely used in aircraft) could be essential in aviation (thus maintained and supported there), but considered outdated in other domains or hindering integration with new systems (like modern networked communications). Scoring at the technology level allows us to see that dynamic.

By scoring **technologies** – not just individual systems – we reveal structural obsolescence even when specific systems still appear functional on the surface. Legacy issues often creep in not through sudden catastrophic failures, but through a gradual disconnect: the technology doesn't fail outright, but the world around it changes to the point where the old tech becomes a constraint or a vulnerability. Focusing on technology types helps us catch that "creeping legacy" earlier. It's a bit like tracking the health of an entire species in an ecosystem, rather than checking the health of one animal at a time.

Concretely, this means our legacy assessment should be able to flag something like "8-bit microcontrollers" or "Windows 7-based systems" or "CAN bus networks" as *technologies* of concern, in addition to looking at specific assets. It also means considering context: perhaps a technology is fine in one setting but problematic in another. We want to capture those differences.

This approach aligns well with the long planning horizons in transport. Systems are built to operate for decades, and a lot can change in that time. A technology-centric view lets us notice, for example, that while all our trains are individually fine today, they all rely on one navigation technology that will be phased out in 5 years – that's a legacy technology issue we need to address **before** it manifests as system failures. It's a more proactive and structural way of understanding legacy, complementing the system-by-system risk assessments.

Building a Scoring Framework for Legacy in Transport

How can we put these ideas into practice? We adapt and extend the UK IT framework into a **scoring model** that is technology-centred and context-aware for the transport sector. The idea is to evaluate a given technology (or system component) along several dimensions to determine if it's drifting into legacy territory. Below are the core categories we propose for scoring legacy status, tailored to transport and similar safety-critical fields:

1. **End of Life** – *Is the technology at or near end-of-life?* This considers whether the manufacturer or maintainer has declared an end-of-life date, whether it's no longer sold, or support (updates, spare parts) has officially ended. If a component is effectively orphaned by its maker, that's a strong legacy signal.

2. **Obsolescence** – *Is knowledge about the technology disappearing?* This looks at things like documentation and skills. If essential manuals are lost or only in one person's archive, or if few people remain who know how to service it, the technology is becoming legacy regardless of its age. A system might still run, but if no one knows its ins and outs (or the one guru retires), it's a ticking time bomb.

3. **Integration** – *How well (or poorly) does it play with others?* This category asks if the technology requires special adapters, "wrappers," or workarounds to interface with modern systems. Also, does it prevent compliance with current interoperability standards or cybersecurity standards? For example, if a sensor network can't output data in a modern format, every integration is a headache – a classic legacy trait.

4. **Operational Dependencies** – *Is the technology tightly coupled to other systems or processes in a way that complicates change?* If replacing or updating this component would break a bunch of other things, or if it's deeply embedded such that even thinking about changing it causes fear, it scores high here. Also included is **cultural or user resistance**: perhaps everyone is used to the old interface and would struggle to switch – this "inertia" is a real factor in legacy risk (people and process, not just hardware).

5. **Support Ecosystem** – *Are the external supports for the technology fading?* This means checking if the tools, suppliers, or APIs that this technology relies on are disappearing or degrading. Maybe an API it uses has been deprecated, or the one vendor that made the calibration tool went out of business. A system might be fine, but if the ecosystem around it is dying off, trouble is brewing.

6. **Security & Risk** – *Does this technology make it hard to secure our system or meet modern safety standards?* Legacy tech often lacks features like encryption, authentication, or even basic safety interlocks that are expected today. If keeping it forces you to have exceptions in your cybersecurity baseline or other risk controls, it scores as legacy. For instance, an old control system might not support strong passwords or software updates, leaving a known security hole.

7. **Financial Factors** – *Is the cost hindering decisions?* Two sides here: (a) **Replacement cost** – perhaps it would be enormously expensive to replace this component, which often leads organizations to "sweat" the asset far beyond its optimal life (and thus it lingers into legacy territory). Or (b) **Sunk cost** – "We invested so much in this 15 years ago, we hate to throw it away." If financial considerations are the primary thing keeping a technology in place, it might already be effectively legacy (we keep it not because it's great, but because we can't afford not to, which is a risk position).

8. **Practicability (ASARP)** – *Can continued use be justified as "As Secure/Safe As Reasonably Practicable"?* This last factor introduces the concept of **ASARP**, drawn from safety engineering. ASARP means that a system is as secure (or safe) as can be reasonably achieved without excessive cost or effort – any further improvement would require disproportionate cost or would impair functionality. In a legacy context, this asks: even if the technology shows legacy characteristics, can we make a strong case that keeping it is acceptable for now? For example, perhaps replacing a particular legacy subsystem would cost billions and cause years of disruption, while keeping it with some compensating security controls yields an acceptable level of risk. If such an argument holds, we mark it here – not to excuse the legacy status indefinitely, but to acknowledge that immediate replacement may not be practical or necessary if mitigations make it "secure/safe enough" for the time being.

Why this scoring matters: The aim of this framework is not to shout “Risk!” at every legacy component and mandate immediate fixes. Instead, it’s about **making legacy visible and measurable**. We can think of it as an early warning system or a health dashboard. A technology could score high on legacy factors (meaning it’s showing its age in support, integration, etc.) yet still be *operationally safe* thanks to mitigations. In that case, the scoring helps us monitor it closely and plan for eventual upgrade, rather than being caught off guard. Conversely, something might currently pose no operational issues, but if it scores high on, say, End of Life and Obsolescence, we know trouble is on the horizon even if nothing has broken yet. This nuance is important in transport: legacy is rarely a sudden binary state, it’s a *trajectory*.

Moreover, scoring at the technology level (with context) lets us spot systemic issues. If, for example, multiple systems across rail and road are scoring “3” (high risk) on “Support Ecosystem” because a certain vendor is exiting the market, that tells decision-makers something crucial: we may need a sector-wide strategy to address that vendor’s absence. This is the kind of insight a pure system-by-system risk approach might miss.

The Limits of OSINT in Spotting Legacy

Open-source intelligence (OSINT) is often the first tool people reach for when trying to understand whether a technology is becoming legacy. Vendor lifecycle announcements, standards bulletins, training catalogues, and even job adverts can provide useful indicators of a technology’s health. These signals are valuable because they are widely available, relatively low-cost to gather, and can offer early warning — for instance, when a vendor posts an end-of-support notice or when demand for a particular skill seems to vanish.

But OSINT has blind spots. A technology that looks obsolete in public may still be quietly supported under private contracts. Equally, something that appears current on paper may be fragile in practice because the skills to maintain it have drained away, or because the infrastructure around it has become brittle. Supply chain dependencies, cultural resistance to change, and the details of safety or certification obligations rarely leave a visible trail online.

The signals themselves can also be ambiguous. Marketing material often blurs the line between genuine end-of-life and a push to migrate customers to the “next generation.” Standards bodies may withdraw specifications that remain mandatory in certain operational contexts. Community forums can keep the appearance of vibrancy long after the technology has become unsustainable.

This is why OSINT should be treated as a starting point, not a conclusion. It can highlight technologies that may warrant closer examination and help organisations prioritise where to spend their limited investigative effort. But the assessment of the extent to which something is truly “legacy” almost always depends on expert domain knowledge: the system engineers who manage competing requirements, the regulators who certify it, and the operators who keep it alive, manage risk and know what would happen if it failed.

OSINT, in other words, is best understood as a wide-angle lens. It helps reveal patterns and surface gaps, but it cannot replace the close-up work of contextual analysis. For transport, where legacy often hides in the interdependencies between platforms, infrastructure, and support environments, that distinction matters.

Practical Use Cases of the Legacy Scoring

How would this approach be used in the real world of transport and critical infrastructure? Here are some practical ways it can help:

1. **Cross-Modal Comparison:** Regulators and industry groups could compare how the *same technology* is faring in different transport modes. For instance, if the railway sector and automotive sector both heavily use a certain communication protocol, a scoring framework could reveal that rail considers it highly legacy (maybe due to obsolescence of support tools) while automotive still finds it acceptable. This cross-modal view can prompt information sharing – maybe the rail industry’s mitigation strategies could help automotive or vice versa. It also helps government or national security planners see if a single point of failure technology (say, GPS for timing) is becoming legacy-vulnerable across multiple critical sectors at once.

2. **Investment Planning:** Organizations can use the scores to prioritize where to invest their limited upgrade budgets. If a transit authority has 100 different tech components across trains and infrastructure, the scoring might show, for example, that the depot support tools are hitting legacy status harder than the trains themselves. That insight could justify allocating money to modernize maintenance tooling (perhaps a less obvious choice than buying new trains, but potentially more urgent). Essentially, it helps target investments *before* a crisis hits, focusing on areas of creeping fragility.

3. **Resilience Tracking:** By mapping out dependencies and scoring them, operators can improve their overall resilience. They might discover that a seemingly minor component has high legacy scores across multiple systems – indicating a **single point of fragility**. For example, if one type of sensor is used in dozens of places and it scores poorly (no longer made, no replacement, hard to secure), that's a resilience issue. Tracking these scores over time also shows if things are getting better or worse. Are we reducing our legacy exposure, or is it accumulating unnoticed?

4. **Assurance and Certification Cases:** When preparing safety cases or cyber assurance documentation for regulators, having a structured legacy assessment can strengthen the argument. Instead of simply saying "we believe this old system is fine," an operator can show that they've scored it, identified weaknesses, and either mitigated them or have a plan. For example, a railway might continue using a legacy signaling component but demonstrate through ASARP reasoning that it's as safe as reasonably practicable – they've added monitoring, have backup units, etc., and replacement at this time would be disproportionately costly for minimal safety gain. A clear framework provides evidence that the decision to retain or replace is *rational* and *documented*.

In all these cases, the scoring framework acts as a common language between engineers, managers, and regulators. It translates a complex mix of technical and organizational factors into a format that can be discussed and reasoned about. Rather than hand-waving "this is old and risky," one can pinpoint **how** it's legacy (e.g. "It's unsupported and no one is trained on it, scoring 3 in Obsolescence and End-of-Life"), and what's being done about it or why it's acceptable for now. This leads to more persuasive justifications and more analytical, evidence-based planning – exactly the balance we need.

Broader Significance

Legacy in transport is not a simple binary of "old vs. new." It is a structural condition of industries where systems are designed to operate for decades, often under regulatory, financial, and operational constraints that make rapid replacement impossible. The real challenge is not eliminating legacy technology — that would be unrealistic — but ensuring that legacy risks are visible, understood, and managed deliberately.

By reframing the conversation:

- From **system to technology**, so we can see when whole technology types are drifting into fragility, not just isolated assets.
- From **age to sustainability**, focusing on whether something can still be supported, secured, and integrated rather than when it was built.
- From **current failures to creeping fragility**, recognising that today's safe, functioning system can still be tomorrow's bottleneck or blind spot.

…we create space for better decisions about where to tolerate legacy, where to mitigate it, and where to invest in replacement. In transport, some legacy will always be with us — by design. The key is to manage it as a normal, expected part of the lifecycle, rather than treating it as an afterthought or surprise.

Handled well, legacy can be a managed risk. Handled poorly, it becomes a hidden vulnerability that undermines resilience across safety, security, and operational performance.

Conclusion

Legacy technologies in transport and critical infrastructure are not simply outdated artefacts to be discarded. They are often the backbone of long-lived systems, embedded in wider ecosystems that still depend on them. The challenge is not their existence, but how we recognise, track, and manage them.

A structured approach helps. By defining legacy in context, mapping it across platforms, infrastructure, and support tooling, and applying principles like As Secure/Safe As Reasonably Practicable (ASARP), we can distinguish between technologies that can be tolerated for now and those that demand urgent attention. This makes decisions about investment, assurance, and resilience clearer, more transparent, and easier to justify.

Ultimately, legacy is less about what sits in the past and more about how we prepare for the future. Every system in transport will one day become legacy; the real question is whether we arrive at that point by surprise or by design. With clarity, evidence, and foresight, legacy stops being a liability and becomes something we can navigate — deliberately, sustainably, and securely — for the decades still ahead.