Why the Way You Maintain Assets No Longer Matches the Risk You Are Carrying

Dynamic asset criticality and maintenance decision-making under climate stress

Martin Boettcher[1]

Abstract: Climate change is exposing a growing mismatch between how infrastructure assets were originally designed and how they are now required to operate. While physical hazards such as extreme heat, flooding and severe weather remain important drivers, many of the most significant changes arise from society's response to climate change as well as changing technologies. Renewable energy, electrification, desalination, digital infrastructure growth and increasing interdependence between critical services are reshaping operating environments and altering where consequence sits within infrastructure systems. As a result, many assets now support operating conditions, demand profiles and system dependencies that differ significantly from those assumed when traditional maintenance strategies and risk frameworks were developed. This paper argues that maintenance strategy has become a primary resilience control. In an environment where capital is constrained and asset replacement is slow, the choice between static and dynamic maintenance approaches can materially influence reliability, resilience and downside risk. The paper presents a practical maintenance-led approach based on dynamic asset criticality, where asset importance changes according to operating context, condition, redundancy, weather exposure and system demand.

The paper highlights a key limitation of conventional risk matrices: their tendency to under-represent low-probability, high-consequence events and the dependency pathways through which relatively minor failures can escalate into major operational and financial impacts. Drawing on examples from power, water and other infrastructure sectors, it demonstrates how condition-informed maintenance and dynamic criticality can help organisations focus effort on assets that become temporarily "super-critical" under stressed conditions

Rather than replacing established frameworks such as ISO 55001, dynamic criticality enhances them by translating climate and operational risk into day-to-day maintenance, inspection and investment decisions, providing a clearer line of sight between asset condition, operational resilience and organisational risk.

Keywords: dynamic criticality; climate resilience; maintenance strategy; infrastructure risk; condition-based intervention

1 Introduction

Infrastructure assets are often designed for one duty, operated in another and maintained under a third.

Climate change is one reason. More frequent extreme weather is changing the conditions under which infrastructure operates. At the same time, technology, economics and changing community expectations are reshaping infrastructure systems. Renewable generation, electrification, desalination and rapid growth in digital infrastructure are changing demand patterns, operating regimes and system dependencies. As a result, many assets now support operating conditions that differ significantly from those assumed when they were originally designed.

In many cases the asset itself has changed little, but the consequence of failure has changed. Assets that once operated within well understood systems now support increasingly interconnected services. Equipment that was once peripheral can become operationally critical. Traditional maintenance strategies can therefore remain technically compliant while becoming progressively disconnected from actual risk. For asset-intensive organisations, maintenance strategy is increasingly a resilience strategy. Where capital replacement is constrained and asset lives are measured in decades, maintenance decisions often provide the fastest practical means of managing risk.

This paper examines the use of dynamic asset criticality and other improved maintenance approaches as a mechanism for aligning maintenance effort with changing operational risk. Rather than treating criticality as fixed, the approach recognises that asset importance changes as operating conditions, demand, redundancy, condition and environmental exposure change.

[1] Martin Boettcher CPEng RPEV RPEQ NER
Principal Asset Engineer, Ventia Engineering Services, Melbourne Australia
e-mail: martin.boettcher@ventia.com.au, telephone +61 437762123

2 Literature Review

Reliability Centred Maintenance (RCM) remains one of the most widely adopted maintenance frameworks within asset-intensive industries. Moubray (1991) developed a structured process for selecting maintenance activities based on asset function, failure modes and failure consequences. Similar principles are reflected in IEC 60300-3-11 (IEC, 2011), which provides guidance for the application of RCM, and the ISO 55000 Series (ISO, 2024), which incorporates risk-based decision making within asset management systems. Standards Australia (2023) provides a framework for the collection and exchange of reliability and maintenance data through AS ISO 14224:2023.

The impact of climate change on infrastructure assets has been examined extensively within the resilience literature. Abdelaal (2026) reported increasing infrastructure exposure to extreme temperatures and weather-related deterioration mechanisms. The UNDRR (2025) reported increasing populations affected by extreme weather events. The World Economic Forum (2023) identified climate-related risks affecting infrastructure reliability, maintenance planning and service continuity. Similar observations have been reported within the water sector by the American Water Works Association (2025). In Australia, the Insurance Council of Australia (2026) reported that extreme weather events caused $3.5 billion in insured losses in 2025, showing that climate-related disruption is already a material financial exposure.

Research into maintenance optimisation has increasingly incorporated changing operating conditions and system risk. Karar et al. (2025) developed a resilience-based maintenance optimisation framework using multiple criteria decision making and scenario-based prioritisation. Pinciroli, Baraldi and Zio (2023) examined maintenance optimisation using condition monitoring and digital technologies within Industry 4.0 environments. Ruiz-Rodríguez et al. (2024) examined how maintenance schedules can be changed as conditions change and uncertainty increases.

Asset criticality is widely used to support maintenance planning, capital investment and risk assessment. Traditional approaches typically combine likelihood and consequence to establish criticality rankings. Boettcher (2021) presented a rule-based criticality methodology incorporating operational mode, redundancy status and environmental conditions. Karar and Labib (2020) developed an Agile Criticality Matrix that adjusts asset criticality according to changing operating conditions. The Callide C4 incident provides a practical example of dynamic criticality. A battery charger that would normally be considered a medium-criticality asset became temporarily super-critical because the generating unit depended on it at that point in time. When it failed, the consequences were far greater than the apparent importance of the asset would have suggested (Brady Heywood, 2024).

Risk matrices continue to be widely applied within asset management and corporate risk frameworks. Cox (2008) examined limitations associated with risk matrix ranking methods. Hubbard (2020) examined limitations associated with qualitative risk scoring approaches. Aven (2021) explored the relationship between risk and resilience within organisational decision making. Liu and McNeil (2020) examined the use of resilience concepts within risk-based asset management planning. Condition-based and predictive maintenance have been supported by advances in sensing, data analytics and digital systems. Bokrantz et al. (2020) defined the concept of Smart Maintenance and examined the integration of data-driven technologies into maintenance management. Cao, Yu and Duan (2025) reviewed condition-based maintenance approaches for complex degradation systems, including degradation modelling, multi-component systems and maintenance strategy optimisation. Brundage et al. (2018) examined the derivation of maintenance performance indicators from work order systems.

The Institute of Asset Management (2025) published guidance on climate resilience within asset management. The guidance examined sector-specific approaches to incorporating climate resilience considerations into asset management activities. The Task Force on Climate-related Financial Disclosures (TCFD, 2017) published guidance for the identification and assessment of climate-related risks and opportunities. Lloyd's (2022) reported industry approaches for incorporating climate-related risks into asset management and operational planning. Increasing attention has also been given to infrastructure interdependency and cascading failure mechanisms. Recent resilience literature has recognised that infrastructure failures increasingly emerge from interactions between systems rather than failures within individual assets (UNDRR, 2025). These observations are relevant to maintenance decision-making because assets that appear to have modest consequences in isolation may play a critical role in preventing wider system disruption when dependencies are considered.

The literature provides strong guidance on RCM, asset management, resilience, climate risk and data-driven maintenance. What is less developed is how these ideas are translated into day-to-day maintenance decisions when asset criticality changes with weather exposure, redundancy, condition, demand and system configuration. This is the gap addressed in this paper.

3 Why Static Maintenance Assumptions Are Failing

3.1 Infrastructure Systems Are Changing

Many infrastructure assets operating today were designed under a different set of assumptions to those that exist now. Design assumptions are typically based on expected demand, operating duty, consequence of failure and available redundancy. Operations and maintenance strategies are then developed around those assumptions and often remain largely unchanged for many years. Climate change is altering those assumptions through more frequent extreme heat, flooding, drought, bushfire and severe weather. At the same time, broader technological, economic and societal changes are reshaping infrastructure networks.

The issue is not only asset age. It is that design assumptions, operating reality and maintenance strategy are drifting apart. An asset may still be in acceptable condition against its original duty but no longer be operating in that duty. A standby asset may now run routinely. A baseload asset may now cycle. A local support system may now determine wider service continuity. In these cases, the maintenance strategy can remain unchanged while the risk it is controlling has changed materially.

The electricity sector provides a clear example. Much of the existing thermal generation fleet was designed for predictable demand and baseload operation. Climate policy, cheaper renewable generation and energy storage have changed dispatch across the sector. Fossil generators that were not designed for frequent cycling are now asked to load follow, start and stop more often, or operate in two-shift operation. This changes the duty on the original plant. Boilers, turbines, valves, pipework, electrical systems and controls experience more thermal movement, fatigue, starts, transients and off-design operation than their original maintenance strategies assumed.

Demand is also changing in ways that affect asset duty. In developed markets, data centres, cooling loads and industrial electrification are creating large and concentrated electrical loads in some locations. In some developing markets, unreliable grid supply is increasing reliance on embedded generation, batteries, UPS systems, standby generators and local fuel logistics. These changes do not affect all assets equally. Transformers, switchgear, backup power systems and cooling equipment may experience higher utilisation, greater thermal loading, reduced redundancy and fewer practical outage windows.

Similar changes are occurring in the water sector. Extended droughts, growth corridors and changing rainfall patterns are increasing reliance on desalination plants, major transfer pipelines, recycled water schemes and high-lift pumping systems. Assets such as desalination plants that were once treated as contingency or drought-response assets may now be essential to routine supply security. In Perth, average streamflow to drinking water dams has fallen from around 420 GL per year before 1975 to approximately 25 GL annually today, forcing increasing reliance on desalination and groundwater sources (Government of Western Australia, 2026). This changes both the duty of these assets and the consequences of their failure.

Communications and digital systems are also becoming increasingly important. The 2024 CrowdStrike outage demonstrated how a failure within a widely used software platform could simultaneously disrupt airlines, transport systems, healthcare providers, financial institutions and government services. In many cases, affected organisations experienced significant operational disruption despite no failure of their own physical assets. The event highlighted how modern infrastructure increasingly depends on shared digital platforms and services that may sit outside the direct control of asset owners (Congressional Research Service, 2024; CrowdStrike, 2024).

Across these examples, the repeated pattern is that support assets become primary risk controls. Backup power, cooling, communications, switchboards, control systems, drainage, access roads and transfer assets may not be the main production assets, but they can determine whether the system continues to operate under stress.

Infrastructure sectors are becoming increasingly interdependent. Power depends on water and communications and IT systems. Water depends on power and communications and IT systems. Communications and IT systems depend on both power and water. As these dependencies increase, the consequences of failure can change even when asset condition does not. Assets that once played a supporting role can become critical as redundancy is lost, demand rises or alternative supply options narrow. Maintenance strategies developed under one set of assumptions may therefore be managing a very different level of risk today.

3.2 Corporate and Financial Planning Constraints

Asset management decisions do not occur in isolation. They are influenced by financial, operational and organisational constraints that affect how risk is managed and resources are allocated.

In some organisations, maintenance budgets are reduced using fixed percentage targets. The approach is simple. The risk is not. When inspections, planned maintenance or defect rectification are deferred to meet short-term cost targets,

redundancy is often reduced and known defects remain in service for longer. Contingency can be consumed before demand, weather or operating conditions are fully understood. The savings are usually visible in the budget. The increased risk often is not.

A related problem is replacement-first thinking. Ageing or underperforming assets are sometimes moved quickly into capital planning before repair, refurbishment, optimisation or duty management options are properly tested. Replacement is necessary for some assets, but it is not always the fastest or most risk-effective response. A replacement-first culture can consume scarce engineering and planning effort while leaving practical maintenance interventions underused.

These pressures leave maintenance teams managing the gap between long-term plans and short-term exposure. Some activities can be deferred with little consequence. Others reduce redundancy or extend known defects. The difficulty is that the difference is rarely visible in budgets or work plans. The same issue can occur with replacement-first planning. A capital project may be the right long-term solution, but it can distract from maintenance and optimisation activities that would reduce risk today.

The point is not that maintenance should replace capital investment, or that all deferred work is unsafe. The point is that corporate and financial decisions change the risk environment in which maintenance is performed. If those changes are not recognised, the organisation may believe it is only managing cost, when it is also changing redundancy, exposure and recovery capacity.

3.3 Corporate risk matrices hide important failure pathways

Corporate risk matrices remain useful for routine prioritisation, but they are weakest where organisations most need clarity: the bottom-right corner. This is the area of very low-likelihood, extreme-consequence events. It is where process safety failures, prolonged service loss, regulatory intervention, major environmental harm, loss of public trust and organisational survival events sit.

These events should not be treated as ordinary high risks with a very low frequency score. A consequence that can overwhelm recovery capacity, remove licence to operate or cause irreversible harm requires separate attention. For these events, consequence class should dominate the discussion, not the arithmetic of the matrix. Reducing the likelihood score does not reduce the consequence. It only makes the event easier to hide in a ranking table.

A related weakness is poor boundary setting. Industry examples are useful, but a risk register that tries to capture every imaginable event becomes clutter, not control. The test should not be whether an event can be imagined, but whether it is credible, material and connected to the organisation's assets, controls and obligations. Corporate risk matrices should focus on credible pathways to material harm, not speculative scenarios that cannot be meaningfully managed.

A further weakness is the treatment of residual risk. An "after controls" score can make an organisation believe a control exists simply because it has been written into a risk assessment. That is unsafe. Controls that remain as intentions, projects, procedures, budget bids or assumptions do not reduce risk. Only implemented, tested and maintained controls should change the score. A risk matrix should reflect the current verified control state, not the hoped-for future state.

This matters for maintenance because many severe pathways begin with assets that look unremarkable in isolation. A battery charger, switchboard, communications link, drainage asset, access road, cooling system, fire system or backup generator may not score as critical under normal conditions. Under stress, its failure can remove redundancy, disable response, prevent recovery or allow a minor defect to escalate into a major event.

Risk matrices can therefore create a dangerous form of comfort. They give the appearance of structured decision making while masking the pathways that could do the greatest damage. In climate-stressed and interdependent systems, maintenance priority should not be based only on the average position of an asset in a static matrix. It must also test whether failure of that asset can credibly contribute to a bottom-right-corner event.

4 From Static to Dynamic Maintenance

Before considering dynamic maintenance, it is useful to examine how maintenance priorities are traditionally established. Asset identification and criticality assessment remain fundamental to effective asset management.

4.1 Not Everything Is an Asset

One of the most damaging ideas introduced into asset management was that every physical item must be captured in the asset management system. In many organisations, asset registers gradually became inventory lists. Every valve, fitting, instrument and component was assigned an asset number whether it influenced organisational outcomes or not.

This creates noise. Engineering attention is finite. Every hour spent classifying, analysing and maintaining low-

consequence items is an hour not spent managing the assets that actually matter. A 50 mm valve that can be purchased from a local hardware supplier should not compete for attention with a major transformer, treatment train, pump station, communications hub or transmission asset. It may require replacement and inventory control. That does not automatically make it an asset.

A common criticism is that the asset register is incomplete. In practice, most organisations already know where the important assets are. The larger problem is often the opposite. Too much engineering effort is consumed managing information about thousands of low-consequence components that contribute little to decision quality. Before maintenance can be prioritised, the clutter must be separated from the assets that genuinely influence safety, service, resilience, compliance and cost.

4.2 Establishing Static Asset Criticality

Once the clutter has been removed, the remaining assets can be prioritised. The objective is not to calculate the perfect answer. It is to develop a consistent answer that directs maintenance, operations and investment effort towards the assets that matter most. In practice, criticality is determined by understanding the consequences of failure and expressing those consequences in business terms. Safety, service delivery, environmental impact, regulatory exposure, recovery effort and financial loss should ultimately translate into dollars, whether directly or indirectly. Criticality is therefore not an engineering exercise. It is a business consequence exercise.

The process should be simple. A small group of experienced operations, maintenance and engineering personnel can often develop a practical criticality model in a matter of days. Rule-based approaches are generally sufficient. If failure of an asset could expose the organisation to a $10 million consequence, it belongs above an asset that presents a $100,000 consequence. If the consequence is measured in a few hundred dollars and a short maintenance call-out, it belongs lower again. The exact number matters less than placing assets into the right consequence group.

Many organisations make the process far more complicated than it needs to be. Large workshops, detailed scoring systems and highly granular FMECA studies can consume enormous effort while adding little value to decisions. If someone proposes a FMECA across tens of thousands of assets, it is usually a good indication the exercise has lost sight of its purpose. Near enough is usually good enough. The objective is not to calculate the perfect consequence value. It is to separate the assets associated with million-dollar consequences from those associated with thousand-dollar consequences.

The output of a static criticality assessment is a small number of consequence groups. Maintenance frequencies, inspection programs, spares strategies and investment decisions can then be aligned to those groups. For many years this approach has served industry well. The challenge is that it assumes consequence remains relatively stable over time.

4.3 Where Static Criticality Stops Being Static

For most organisations, static criticality works surprisingly well. The consequences of failure for many assets do not change significantly from day to day. A major transformer, treatment plant, pump station or transmission asset will usually remain important regardless of the operating conditions. For these assets, static criticality provides a practical and effective basis for maintenance planning, spares strategies and investment decisions.

The difficulty arises when the consequence of failure depends on operating context. Many assets are neither permanently critical nor permanently non-critical. Their importance changes as redundancy, demand, condition, weather exposure and system configuration change. The asset itself may not have changed at all. The consequence of failure has changed.

A common example is a water pumping station with three duty pumps. Under normal conditions, the loss of a single pump may have little impact because two pumps remain available and the station retains redundancy. The pumps may therefore be assessed as having a moderate criticality. If one pump is removed from service for maintenance, the remaining pumps immediately become more important because the resilience margin of the system has been reduced. If a second pump then becomes unavailable, the final remaining pump effectively becomes the single barrier preventing loss of supply. The condition of the pump may be unchanged, but its consequence of failure has increased dramatically. At that point it can become temporarily super-critical.

This pattern occurs throughout infrastructure systems. Backup generators become critical when grid reliability deteriorates. Communications systems become critical during emergencies. Transfer assets become critical when alternative supply paths are unavailable. Static criticality remains a valuable starting point, but it does not always reflect the consequence that exists today. The challenge is recognising when consequence has moved and adjusting maintenance priorities before failure occurs.

4.4 Making Criticality Dynamic

Static criticality establishes which assets are important. Dynamic criticality determines when that importance changes. The objective is not to recalculate the entire asset register every day. It is to identify the small number of assets whose consequences have changed because operating conditions have changed.

The process begins by monitoring the factors that influence consequence. Typical examples include condition, known defects, redundancy status, demand, outages, weather exposure, access constraints and supply chain issues. Most of this information already exists within organisations. It is often scattered across maintenance systems, work orders, defect lists, operating logs, outage plans and weather services. Dynamic criticality brings these inputs together and asks a simple question: if this asset failed today, would the consequence be different to the consequence assumed during the original criticality assessment?

A simple visual representation can be highly effective. Assets can retain their base criticality while being temporarily elevated when predefined conditions occur. A dashboard showing critical assets, major defects and operating constraints often provides more value than a detailed mathematical model.

Scenario thinking can then be used to test the consequences of changing conditions. If one pump is unavailable, what happens if a second pump fails? If access is lost due to flooding, which defects become unacceptable? If a backup generator is already out of service, which electrical assets become critical to maintaining operations? These assessments do not need to be complicated. Simple "if this, then what?" checks often identify the assets whose importance has changed most.

Artificial intelligence can assist by monitoring defect records, operating conditions, asset availability, weather forecasts and maintenance backlogs and highlighting situations where criticality may have changed. The value of AI is not that it determines criticality. The value is that it can identify patterns and combinations of events that may otherwise be overlooked. Engineering judgement remains responsible for the final decision.

The output is a list of assets whose current consequence differs materially from their normal consequence. These assets become the focus of maintenance attention until conditions return to normal. Table 1 shows a simple dynamic criticality framework that links changing consequence to maintenance decisions. In this way, maintenance effort follows risk as it moves through the system rather than remaining fixed to assumptions established months or years earlier.

Table 1. Practical dynamic criticality logic

Element	Purpose	Typical inputs	Decision output
Base criticality	Establish the normal consequence of failure in business terms.	Safety consequences, service loss, environmental impact, regulatory exposure, recovery effort all defined in risk dollar terms.	Normal criticality groups are used for default maintenance, inspection, spares and outage planning.
Current asset state	Identify whether the asset is currently more vulnerable than normal.	Known defects, condition monitoring, inspection findings, work history, overdue maintenance, temporary repairs and reliability issues.	Escalate maintenance priority, increase inspection, repair earlier or apply temporary controls.
System context	Identify whether the consequence of failure has changed.	Redundancy status, parallel assets out of service, system demand, operating mode, outage plans, weather exposure, access constraints and supply constraints.	Temporarily uplift or downgrade criticality based on current operating risk.
Scenario checks	Test credible “if this, then what?” pathways.	Loss of another pump, loss of backup power, flood access loss, heatwave demand, storm exposure, communications failure or known weak points.	Identify assets that could become temporarily super-critical.
Visibility and triggers	Make changing criticality visible to operations, maintenance and management.	Dashboard, defect list, outage status, weather forecast, AI-assisted monitoring, operating constraints and agreed rule triggers.	Issue alerts, reprioritise work, delay outages, mobilise spares or focus engineering attention.
Action rules	Convert changing criticality into controlled maintenance decisions.	Agreed authority levels, risk tolerance, maintenance rules, escalation paths and work management links.	Targeted inspection, defect escalation, outage deferral, temporary monitoring, contingency preparation or repair.

4.5 Following Risk as it Moves

The objective of dynamic criticality is not to create a more complicated criticality model. The objective is to direct maintenance effort towards the assets that are currently carrying the greatest consequence of failure. Most organisations already know which assets are generally important. The challenge is recognising when consequence has shifted and making sure maintenance priorities shift with it.

Risk rarely remains stationary within an infrastructure system. As assets are removed from service, defects emerge, demand changes, weather conditions deteriorate or operating configurations change, the consequences associated with failure also change. An asset that was assessed as moderately critical six months ago may become one of the most important assets in the system today. Equally, an asset that normally attracts significant attention may become less important when redundancy is restored or operating conditions improve.

The practical value of dynamic criticality is that maintenance effort follows these changing conditions. Inspection resources can be redirected towards assets whose consequences have increased. Defects that would normally be tolerated can be escalated when resilience margins decrease. Planned outages can be delayed when the system becomes vulnerable. Spares, contractor resources and engineering attention can be concentrated where they provide the greatest reduction in risk.

Importantly, this does not mean maintaining every asset to a higher standard. In many cases it means the opposite. By understanding where consequence currently sits, organisations can safely reduce effort on low-consequence assets while concentrating scarce maintenance resources on the assets that are temporarily protecting service continuity. The result is not more maintenance. It is better-targeted maintenance.

This shift is particularly important in climate-stressed and resource-constrained environments. Most organisations cannot replace everything, inspect everything or maintain everything to the same standard. Dynamic criticality provides a mechanism for making those decisions more visible, more defensible and better aligned with the risks the organisation is carrying today rather than the risks it assumed years ago.

4.6 Resilience Is a Maintenance Decision

Many discussions about resilience assume the answer is simply replacing ageing infrastructure. In reality, many critical assets will remain in service for decades. Dams, transmission corridors, major pipelines, water networks and treatment facilities often have operational lives measured in generations rather than years. Even where replacement is technically justified, capital funding, approvals, outages and construction programs can delay implementation for many years.

At the same time, new infrastructure will continue to be added. Renewable generation, batteries, desalination plants, data centres and digital systems are becoming increasingly important parts of modern infrastructure networks. These assets create opportunities, but they also introduce new failure modes, new dependencies and new operational risks. Newer infrastructure is not automatically more resilient infrastructure.

The future is therefore not a choice between old infrastructure and new infrastructure. Most organisations will operate a combination of both for the foreseeable future. The challenge is ensuring that ageing assets, new assets and supporting systems continue to operate together as a resilient system.

This is where maintenance becomes important. Resilience is not determined solely by strategy documents, climate adaptation plans or future capital projects. It is also determined by day-to-day decisions about inspections, defect management, outages, repairs and asset condition. A well-maintained asset can remain a reliable part of the system for many years beyond its original expectations. Conversely, a poorly maintained asset can become the starting point for a major failure pathway regardless of its age.

The practical consequence is that resilience often depends less on which assets an organisation owns and more on how effectively it maintains them. Dynamic criticality provides a mechanism for identifying the assets that are currently protecting system resilience and directing maintenance effort towards them before failure occurs. In this way, maintenance acts as the bridge between the infrastructure inherited from the past and the infrastructure being built for the future.

5 Climate Change and Engineering Assumptions

5.1 Engineering Design Is Built on Assumptions

Every engineered system is built on assumptions. Designers must decide what temperatures are likely to occur, how much rain will fall, what flood levels are credible, how strong winds may become, what wave heights should be considered, how much demand the system will experience and how often equipment is expected to operate. Similar assumptions are made regarding power availability, access for maintenance, emergency response capability, supply chains and the

availability of redundancy. Without these assumptions, nothing could ever be designed because there would be no practical limit to what might need to be considered.

These assumptions are not mistakes. They are engineering compromises. Every design represents a balance between performance, risk and cost. It is always possible to build something stronger, larger or more resilient. A house can be built like a fortress. A pumping station can be built with more redundancy. A power station can be designed for more extreme ambient conditions. The question is not whether it is technically possible. The question is whether the additional cost is justified by the reduction in risk. Good engineering is rarely about creating the perfect design. It is about creating a design that is good enough for the expected conditions.

For this reason, design standards and engineering practices typically focus on credible conditions rather than worst-case conditions. Infrastructure is designed to withstand events that are considered reasonably foreseeable and economically justifiable. Exceeding those design assumptions is expected to be rare. For most of an asset's life this approach works well. Most infrastructure performs as intended because the conditions experienced remain broadly aligned with the assumptions used during design.

Problems arise when those assumptions change. An asset may remain in good physical condition while the environment around it changes significantly. Temperatures may exceed historic design ranges. Rainfall patterns may shift. Flood frequencies may increase. Assets may operate harder or more frequently than originally intended. Support systems that were once considered reliable may become less reliable. Under these circumstances, the original design may still be functioning exactly as intended, but the assumptions that justified the design are no longer valid.

This distinction is important. Infrastructure failures are not always the result of poor maintenance, poor operation or poor engineering. In many cases they occur because the conditions for which the system was designed no longer reflect the conditions under which it is operating. The challenge for asset managers is therefore not only understanding asset condition. It is understanding which assumptions continue to hold and which assumptions are becoming increasingly difficult to defend.

5.2 When Assumptions Become Wrong

Engineering history contains many examples where infrastructure performed broadly in accordance with its design, but the assumptions underpinning the design proved incorrect, incomplete or no longer applicable.

The Fukushima Daiichi accident is one example. The facility incorporated multiple layers of protection, backup power and emergency systems. These protections were designed around assumptions regarding credible tsunami heights and the availability of emergency systems following an extreme external event. When those assumptions were exceeded, multiple protective barriers were affected simultaneously (IAEA, 2015).

The failure of the New Orleans levee system during Hurricane Katrina provides a different example. The flood protection system was developed around assumptions regarding storm intensity, storm surge height and loading conditions. When those assumptions were exceeded, large sections of the protection system were compromised, exposing the surrounding city to extensive flooding (Select Bipartisan Committee, 2006).

Major flood events in Queensland and elsewhere in Australia have highlighted similar issues. Bridges, roads, drainage systems, pump stations and flood mitigation assets are commonly designed using historic rainfall records and expected flood frequencies. When rainfall intensity, flood extent or event duration exceeds those assumptions, infrastructure that was considered adequately protected can become vulnerable (AIDR, n.d.).

Extreme heat provides another simple example. Roads, bridges, runways and rail systems are designed for expected temperature ranges and thermal movement. When temperatures exceed those assumptions, rail deformation, pavement deterioration, bridge-joint expansion and signalling issues can emerge (UNECE, 2026).

The California wildfire seasons provide another example. Electricity networks, communications systems and communities were developed around assumptions regarding vegetation growth, fuel loads, fire behaviour and weather conditions. Increasing fire intensity and duration have challenged many of those assumptions, forcing changes to operating practices, maintenance strategies and asset design (CPUC, n.d.).

Coastal infrastructure around the world is facing similar pressures. Ports, seawalls, coastal roads and marine facilities are commonly designed around assumptions regarding wave height, storm surge levels, sea level and erosion rates. As these conditions change, assets that remain structurally sound can experience operating environments significantly different from those considered during design.

These examples occur in different industries and under different conditions, but they share a common characteristic. The assets themselves were often not the primary problem. The underlying assumptions used to justify the design no longer reflected the conditions the infrastructure eventually encountered. In many cases, multiple layers of protection relied on the same underlying assumption. Redundancy only works if the redundant asset does not share the same vulnerability. A

backup server located in the same flood zone, a backup generator dependent on the same fuel supply, or a standby pump exposed to the same inundation level may appear resilient on paper while remaining vulnerable to the same event. The vulnerability is often not hidden within the asset itself, but within the assumptions surrounding it.

5.3 Recognising Emerging Vulnerabilities

The challenge for asset managers is not predicting every possible climate scenario. The challenge is recognising when the assumptions underpinning resilience are beginning to weaken. In many cases, vulnerabilities emerge gradually and remain largely invisible within traditional maintenance reports, condition assessments and asset registers.

One useful question is whether assets are still operating within the assumptions used when maintenance strategies were originally developed. Higher utilisation, increasing temperatures, reduced outage windows, changing demand profiles and more frequent use of contingency equipment can all indicate that the operating environment has changed materially. An asset does not need to be in poor condition for this to become a concern.

A second question is whether redundancy is genuinely independent. Backup systems are often assumed to improve resilience, but their value depends on whether they remain available during the same event affecting the primary system. A standby asset located in the same flood zone, dependent on the same power supply, connected to the same communications network or reliant on the same access route may provide less protection than expected.

A third question is whether the organisation understands the consequences of losing supporting services. What happens if power is unavailable? What happens if communications are lost? What happens if cooling water is restricted, access roads are blocked or fuel deliveries are interrupted? These dependencies are often treated as services rather than resilience assumptions. They can determine whether recovery remains possible following an event.

Another useful question is whether support assets have become more important than the primary assets they support. Cooling systems, backup power supplies, drainage systems, communications networks, control systems and access infrastructure are often treated as secondary assets. Under stressed conditions they may become the assets that determine whether the wider system continues to operate.

Questions should also be asked about recurring workarounds. Long-term temporary repairs, repeated maintenance deferrals, increasing defect backlogs, contingency equipment in routine service and assets operating outside their intended duty can all indicate that resilience margins are being consumed. These conditions do not necessarily predict imminent failure. They indicate that the system is becoming less tolerant of failure.

Finally, organisations should consider which assumptions have never been revisited. Infrastructure systems do not operate in isolation. Power depends on communications. Communications depend on power. Water depends on both. Transport, fuel supply, cloud services, contractors and external service providers can all become part of the failure pathway. Some of the most significant vulnerabilities exist not within the organisation itself, but within the organisations and services on which it now depends.

The purpose of these questions is not to identify every possible failure. It is to identify where assumptions may no longer be valid and where hidden dependencies may be quietly increasing the consequences of failure. Climate-related vulnerabilities often emerge long before any asset actually fails. The warning signs are frequently visible, provided the organisation knows where to look

6 Conclusion

Climate change does not mean infrastructure failure is inevitable. The problem can be managed, but only if organisations stop treating maintenance as a residual cost and start treating it as a resilience capability. The conditions around many assets are changing. Design assumptions are being tested. Redundancy is being consumed. Support systems are becoming critical. Consequences are moving through infrastructure networks in ways that are not always visible in traditional asset registers, maintenance plans or corporate risk matrices.

Capital replacement, redesign and climate adaptation programs remain important. However, they are slow, expensive and constrained by funding, approvals, outage windows and delivery capacity. In the meantime, organisations must continue operating the assets they already own. Many of those assets will remain in service for decades. Resilience will therefore depend heavily on how well those assets are understood, maintained and prioritised.

The positive message is that organisations are not powerless. Most already hold much of the information needed to act. They know where defects exist. They know where redundancy has been lost. They know which assets are difficult to access, operating harder than intended or have become critical to continuity. The challenge is to bring that information together in a way that makes changing risk visible and usable.

That visibility must be practical. A dynamic criticality system should not disappear into complex models, unreadable spreadsheets or dashboards that look impressive but do not change decisions. It should show clearly which assets are carrying increased consequence today, why that consequence has changed, and what action is required. If the output does not help maintenance, operations and management decide what to inspect, repair, defer, monitor or escalate, it is not doing its job.

This requires investment in maintenance capability. It requires reliable asset information, clear criticality rules, defect visibility, operating context, weather awareness and practical decision triggers. It also requires enough skilled people to interpret the information and act on it. Dynamic criticality is not a dashboard exercise. It is an engineering and management discipline that must be resourced.

Boards and executives have an important role in this. If they want resilience, they must give asset managers the ability to manage the risks that threaten resilience. That means moving beyond fixed percentage maintenance cuts, replacement-first thinking and risk registers that assume controls exist because they have been written down. It means asking where the organisation could lose redundancy, where support systems could fail, where adjacent services could be unavailable, and where a low-profile asset could contribute to an organisation-level event.

The aim is not to maintain everything to the highest standard. That is unaffordable and unnecessary. The aim is to understand where failure would hurt most under current conditions, then direct scarce maintenance effort accordingly. Some assets can safely wait. Others cannot. The difference depends on consequence, context and timing.

There is a practical path forward. Separate the critical from the clutter. Express consequence in business terms. Test whether redundancy is real. Look for hidden dependencies. Make changing criticality visible. Move maintenance effort as risk moves. These are not abstract ideas. They are practical steps that help organisations reduce exposure to severe events before those events occur.

Climate change will continue to challenge infrastructure systems and some assumptions will fail. However, organisations that invest in maintenance capability, understand their vulnerabilities and maintain visibility of changing risk are far better positioned to respond. The objective is not to prevent every failure. It is to ensure that when failures occur, they remain manageable.

7 References

1. Abdelaal, A. (2026) Climate impacts on rail infrastructure and maintenance adaptation. Journal of Infrastructure Systems (in press).
2. Australian Institute for Disaster Resilience (AIDR) (n.d.) *Queensland Flood, 2010/2011*. Australian Disaster Resilience Knowledge Hub, Link.
3. Aven, T. (2021) 'On some foundational issues concerning the relationship between risk and resilience', Risk Analysis, 42(9), pp. 2062–2074, Link.
4. AWWA (2025) Incorporating Climate Change Impacts Into Demand Forecasting. American Water Works Association, Link.
5. Boettcher, M. (2021a) 'High Speed Rule Based Criticality Analysis', Proceedings of the AMPEAK Conference 2021, Link.
6. Bokrantz, J., Skoogh, A., Berlin, C. and Stahre, J. (2020) 'Smart Maintenance: an empirically grounded conceptualization', International Journal of Production Economics, 223, 107534, Link.
7. Brady Heywood (2024) Technical and Organisational Investigation of the Callide C4 Incident. Brisbane: Brady Heywood, Animation Link. Main report Link
8. Brundage, M.P., Sexton, T.B., Morris, K.C., Moccozet, S.A. and Hoffman, M.L. (2018) 'Developing Maintenance Key Performance Indicators from Maintenance Work Order Data', 46th SME North American Manufacturing Research Conference (NAMRC) and ASME Manufacturing Science and Engineering Conference (MSEC 2018), College Station, TX, Link.
9. California Public Utilities Commission (CPUC) (n.d.) *Public Safety Power Shutoffs, Link*.
10. Cao, H., Yu, J. and Duan, F. (2025) 'Condition-Based Maintenance in Complex Degradation Systems: A Review of Modeling Evolution, Multi-Component Systems, and Maintenance Strategies', *Machines*, 13(8), 714, Link.
11. Congressional Research Service (2024) *CrowdStrike IT Outage: Impacts to Public Safety Systems and Considerations for Congress*. Washington, DC: Congressional Research Service, Link.
12. Cox, L.A. (2008) 'What's wrong with risk matrices?', Risk Analysis, 28(2), pp. 497–512, Link.
13. CrowdStrike (2024) External Technical Root Cause Analysis — Channel File 291. Austin: CrowdStrike, Link.
14. Government of Western Australia (2026) Rebalancing our groundwater, Link.
15. Hubbard, D. (2020) The Failure of Risk Management: Why It's Broken and How to Fix It. 2nd edn. Hoboken: Wiley, Link.

16. Institute of Asset Management (IAM) (2025) Assessment Tools and Methodologies for Climate Resilience Across Sectors. Bristol: IAM, Link.
17. Insurance Council of Australia (2026) Extreme weather cost $3.5 billion in 2025, Link.
18. International Atomic Energy Agency (IAEA) (2015) The Fukushima Daiichi Accident: Report by the Director General. Vienna: IAEA, Link.
19. Karar, A. and Labib, A. (2020) 'An agile criticality assessment framework for dynamic asset management', Journal of Quality in Maintenance Engineering, 26(4), pp. 567–585, Link.
20. Karar, A., Labib, A. and Jones, D. (2025) 'A resilience-based maintenance optimisation framework using multiple criteria and knapsack methods', Reliability Engineering & System Safety (in press), Link.
21. Liu, Y. and McNeil, S. (2020) 'Using resilience in risk-based asset management plans', Transportation Research Record, 2674(4), pp. 178–192, Link.
22. Lloyd's (2022) Catastrophe modelling & climate change report. London: Lloyd's, Link.
23. Moubray, J. (1991) Reliability-Centered Maintenance. Oxford: Butterworth-Heinemann, Link.
24. Pinciroli, L., Baraldi, P. and Zio, E. (2023) 'Maintenance optimization in Industry 4.0', Reliability Engineering & System Safety, 234, 109204, Link.
25. Ruiz-Rodríguez, M.L., Kubler, S., Robert, J. and Le Traon, Y. (2024) 'Dynamic maintenance scheduling approach under uncertainty: Comparison between reinforcement learning, genetic algorithm simheuristic, dispatching rules', *Expert Systems with Applications*, 248, 123404, Link.
26. Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina (2006) A Failure of Initiative: Final Report of the Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina. Washington, DC: U.S. House of Representatives, Link.
27. Standards Australia (2011) AS IEC 60300.3.11—2011: Dependability management, Part 3.11: Application guide — Reliability centred maintenance. Sydney: Standards Australia, Link.
28. Standards Australia (2023) AS ISO 14224:2023: Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment. Sydney: Standards Australia, Link.
29. Task Force on Climate-related Financial Disclosures (TCFD) (2017) *Recommendations of the Task Force on Climate-related Financial Disclosures*. Basel: Financial Stability Board, Link.
30. The *ISO 55000 series (ISP, 2024): Asset management*. International Organization for Standardization, Link.
31. United Nations Economic Commission for Europe (UNECE) (2026) Assessment of Climate Change Impacts and Adaptation for Inland Transport: Towards climate resilient transport systems, reported in UN News, 30 June, Link.
32. United Nations Office for Disaster Risk Reduction (UNDRR) (2025) Global Assessment Report on Disaster Risk Reduction 2025: Resilience Pays — Financing and Investing for Our Future. Geneva: United Nations, Link
33. World Economic Forum (2023) *The Global Risks Report 2023*. 18th edn. Geneva: World Economic Forum, Link.